User's Guide
Converged Network Adapter

QMD8262-k, QLE8262, QME8262-k

This document is provided for informational purposes only and may contain errors. QLogic reserves the right, without notice, to make changes to this document or in product design or specifications. QLogic disclaims any warranty of any kind, expressed or implied, and does not guarantee that any results or performance described in the document will be achieved by you. All statements regarding QLogic's future direction and intent are subject to change or withdrawal without notice and represent goals and objectives only.

Document Revision History	
Revision A, September 2011	
Revision B, March 30, 2012	
Revision C, July 18, 2012	
Revision D, August 30, 2012	
Revision E, November 20, 2012	
Revision F, April 8, 2013	
Revision G, September 6, 2013	
Revision H, June 9, 2014	
Revision J, January 22, 2015	
Revision K, June 24, 2015	
Revision L, March 24, 2016	
Revision M, January 27, 2017	
Changes	**Sections Affected**
Updated QLogic header logo, and Dell logo.	All
Updated Windows content to include Windows Server 2016 Nano, 2008 SP2 and x64 (12G Only), 2008 R2 with SP1, 2012, 2012 R2,	"Supported Operating Systems" on page xv
Windows PE 5.0 64-bit, and Windows PE 10.0 64-bit	
Updated RHEL content to include 7.3, 7.2, 6.8, 6.7, 12G and 13G, CentOS 7.2, and SLES 12 SP2, and 11 SP4	
Updated VMware to include vSphere ESXi 6.0 U2, 6.5 U1	
Added XenServer 7.0, and 6.5	
Added UEFI 2.3, 2.3.1, and 2.5	
Removed "Running the DUP from the Command Line" topic.	"Windows Driver Installation and Configuration" on page 5
Removed redundant options.	"Options" on page 12

Table of Contents

5   Boot Configuration

Introduction

Overview

This user's guide covers the following products:

- QLogic QMD8262-k blade network daughter card
- QLogic QLE8262 monolithic server standup card
- QLogic QME8262-k blade mezzanine card

> **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.

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

The *QLogic QMD8262-k/QLE8262/QME8262-k User's Guide* includes the following sections:

- Hardware Installation covers the hardware and software requirements, safety precautions, a pre-installation checklist, and adapter installation.

- Driver Installation and Configuration covers the installation of the three drivers—NIC, iSCSI, and Fibre Channel over Ethernet (FCoE)—included with the adapter on Windows, Linux, and VMware operating systems.

- Adapter Management Applications covers how to use QConvergeConsole®, as well as operating-system-specific applications for Windows, Linux, and VMware.

- ■ Switch Independent Partitioning covers how to configure Switch Independent Partitioning using utilities such as QConvergeConsole, as well as configuring iSCSI over data center bridging exchange (DCBX) using a Brocade® Series 8000 FCoE switch and a QLogic iSCSI Host Bus Adapter.

- ■ Boot Configuration provides information on booting from SAN, pre-execution environment (PXE) boot setup, and iSCSI boot configuration using *Fast!*UTIL, iSCSI Boot Firmware Table (iBFT), DHCP, and QConvergeConsole.

- ■ Troubleshooting provides troubleshooting flowcharts of steps for diagnosing adapter problems specific to NIC, iSCSI, FCoE, and ESX®.

- ■ Specifications defines the physical characteristics and power requirements, and lists supported standards, interface specifications, and environmental specifications.

- ■ QConvergeConsole GUI provides an overview of the QConvergeConsole Web management interface.

- ■ Regulatory Information provides warranty, regulatory, and compliance information.

Related Materials

For additional information, refer to the following:

- ■ *QConvergeConsole GUI Help System*, available through the QConvergeConsole GUI, provides help topics on configuring and managing host servers and adapters using the QConvergeConsole GUI.

- ■ *QConvergeConsole GUI Installation Guide* contains instructions for installing and starting the QConvergeConsole GUI.

- ■ *QConvergeConsole CLI User's Guide* provides reference material on using the QConvergeConsole CLI.

- ■ *QLogic QConvergeConsole Plug-ins for vSphere User's Guide* provides reference material on using the QConvergeConsole VMware vCenter Server Plug-in and the QConvergeConsole VMware vSphere Web Client Plug-in.

> **NOTE**
>
> To access QLogic documents online, go to www.qlogic.com and click **Downloads**.

Functionality and Features

This section provides the following information:

- Functional Description
- Features
- Supported Operating Systems

Functional Description

Functional descriptions for the adapters are as follows:

- **QMD8262-k:** This a network daughter card with FCoE and iSCSI offload for the blade server environment.

- **QLE8262:** This is a standard form-factor adapter with FCoE and iSCSI offload for the rack and tower server environment.

- **QME8262-k:** This is a mezzanine card with FCoE and iSCSI offload for the blade server environment.

Features

The adapters provide the following features:

- Switch Independent Partitioning
- Message signaled interrupts (MSI-X)
- Device management for power and SAN
- Multi-boot capability including:
  - ❑ PXE
  - ❑ iSCSI
  - ❑ Fibre Channel
  - ❑ Unified extensible firmware interface (UEFI)
- PCIe® 2.0 x8
- User diagnostics that can be run from the CLI and the GUI
- Ethernet functions include:
  - ❑ 2x10 gigabit Ethernet (GbE) with KR (copper backplane) (does not apply to QLE8262)
  - ❑ Priority and virtual LAN (VLAN) tagging
  - ❑ Jumbo frames up to 9618 bytes
- Enhanced Ethernet functions include:
  - ❑ Priority-based flow control
  - ❑ Enhanced transmission selection
  - ❑ Advanced teaming
  - ❑ VLAN configuration and management
  - ❑ Preservation of teaming and VLAN configuration information during driver upgrade

- Advanced stateless offload features include:
    - IP, TCP, and user datagram protocol (UDP) checksums
    - Large segment offload (LSO)
    - Large receive offload (LRO)
- Stateful offload features include:
    - iSCSI offload
    - Fibre Channel and FCoE offload
- Advanced management features for Converged Network Adapters and Fibre Channel adapters, including QConvergeConsole (GUI and CLI)

- Interrupt management and scalability features including:
    - Receive side scaling (RSS)
    - Interrupt moderation
    - Flow control
    - Locally administered address (LAA)
- Enhanced optimization with MSI, MSI-X, and NetQueue

Supported Operating Systems

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

Windows

- Windows Server® 2016 Nano
- Windows Server 2012
- Windows Server 2012 R2
- Windows Server 2008 SP2 and x64 (12G Only)
- Windows Server 2008 R2 with SP1
- Windows PE 50 64-bit
- Windows PE 10.0 64-bit

Linux

- Red Hat® Enterprise Linux (RHEL®) 7.3
- Red Hat Enterprise Linux (RHEL) 7.2
- Red Hat Enterprise Linux (RHEL) 6.8
- Red Hat Enterprise Linux (RHEL) 6.7 12G and 13G
- CentOS® 7.2 and later
- SUSE® Linux Enterprise Server 12 SP2
- SUSE Linux Enterprise Server 11 SP4

VMware

- vSphere®: VMwareESXi 6.0 U2
- vSphere: VMwareESXi 6.5 U1
- vSphere: VMwareESXi 5.1 U3

XenServer™

- XenServer 7.0
- XenServer 6.5
- UEFI 2.3, 2.3.1, and 2.5

NOTE

For the most current versions of the OS and drivers supported by the adapter, refer to the release notes. The release notes are supplied in the release.txt file.

---

1 Hardware Installation

Overview

This section provides the hardware and software requirements, safety precautions, a pre-installation checklist, and a procedure for installing the adapter.

Hardware and Software Requirements

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

- Hardware

  - ❑ For port and slot assignments for the QMD8262-k or QME8262-k adapter, refer to the blade and M1000e chassis diagram in the *Dell PowerEdge M1000e Systems Configuration Guide*.

  - ❑ For QLE8262 adapter port and slot assignments, refer to the "Expansion Cards" section of the *Hardware Owner's Manual* for your Dell PowerEdge server.

- Software: For information on the supported operating systems, firmware versions, adapter drivers, and utilities, refer to the product release notes.

Safety Precautions

⚠ CAUTION

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.

- 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 antistatic devices and an antistatic 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.

Installing the Adapter

Follow the instructions for your adapter.

QMD8262-k, QME8262-k

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*:

ftp://ftp.dell.com/Manuals/all-products/esuprt_ser_stor_net/esuprt_poweredge/poweredge-m610x_Owner%27s%20Manual_en-us.pdf

QLE8262

To install the QLE8262 adapter, follow these steps:

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 x8 bus slot.

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 Ethernet 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.

Connecting to the Network

Follow the instructions for your adapter.

QMD8262-k, QME8262-k

Refer to the "Guidelines for Installing I/O Modules" section of the *Dell PowerEdge Modular Systems Hardware Owner's Manual*:

ftp://ftp.dell.com/Manuals/all-products/esuprt_ser_stor_net/esuprt_poweredge/poweredge-m610x_Owner%27s%20Manual_en-us.pdf

QLE8262

Refer to the *Hardware Owner's Manual* for your Dell PowerEdge server.

2 Driver Installation and Configuration

Overview

NOTE

If you need to update the Flash memory of multiple adapters simultaneously:

- For the QConvergeConsole GUI, refer to the "Update the Flash Using the Flash Update Wizard" topic in the QConvergeConsole Help System.
- For the QConvergeConsole CLI, use the `-flashsupport` command to update the Flash memory for all cards supported by the specified file (for example, `qaucli -pr nic -flashsupport -i ALL -a p3p11179.bin`).

This section provides links to the following information about the three drivers—NIC, iSCSI, and FCoE—included with the adapter:

- Windows Driver Installation and Configuration
- Linux Driver Installation and Configuration
- VMware Driver Installation and Configuration

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

Running the DUP in the GUI

To run the DUP in the GUI:

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

> **NOTE**
>
> The actual file name of the DUP varies.

The Update Package window appears, as shown in Figure 2-1.

Figure 2-1. Update Package Window

2. Click **Install** to continue.

The QLogic Super Installer—InstallShield® Wizard appears, as shown in Figure 2-2.
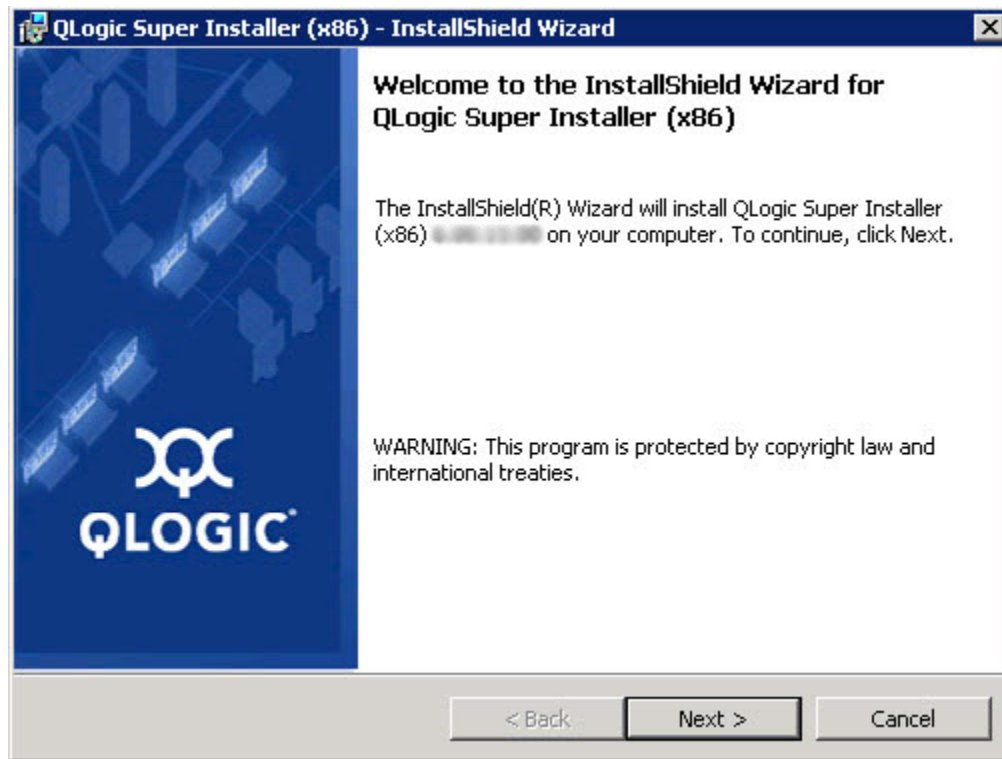
Figure 2-2. QLogic Super Installer—InstallShield Wizard

3.   Click **Next** to continue.

The License Agreement dialog box appears, as shown in Figure 2-3.
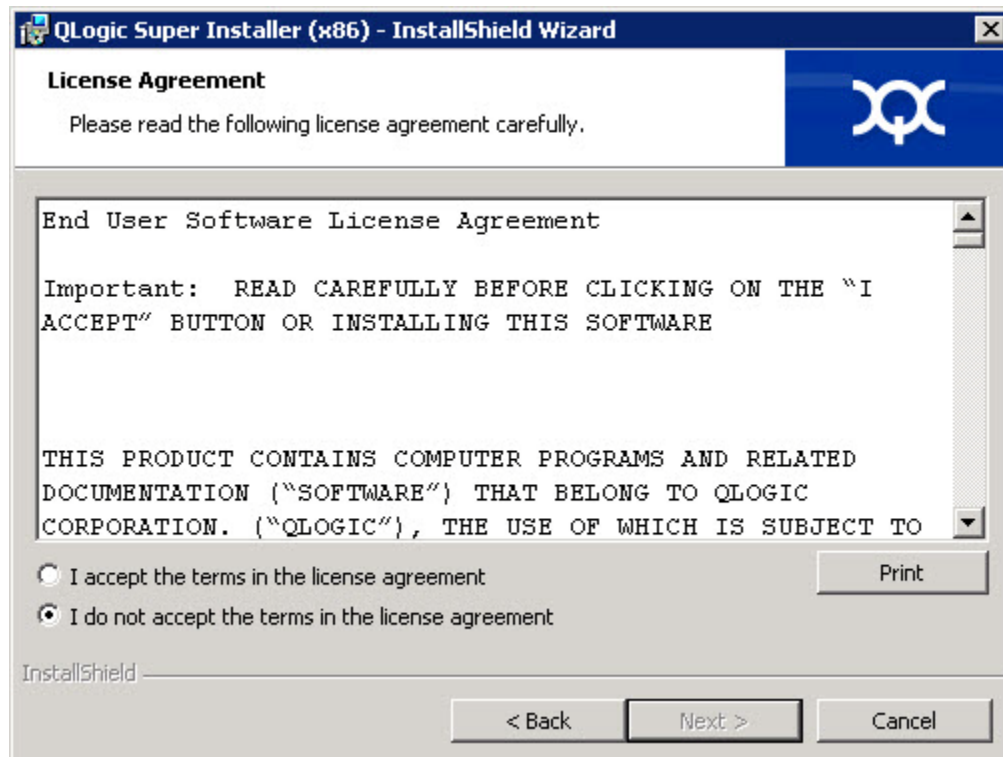
Figure 2-3. License Agreement Dialog Box

4. Select **I accept the terms of the license agreement** and click **Next**.

The Setup Type dialog box appears, as shown in Figure 2-4.
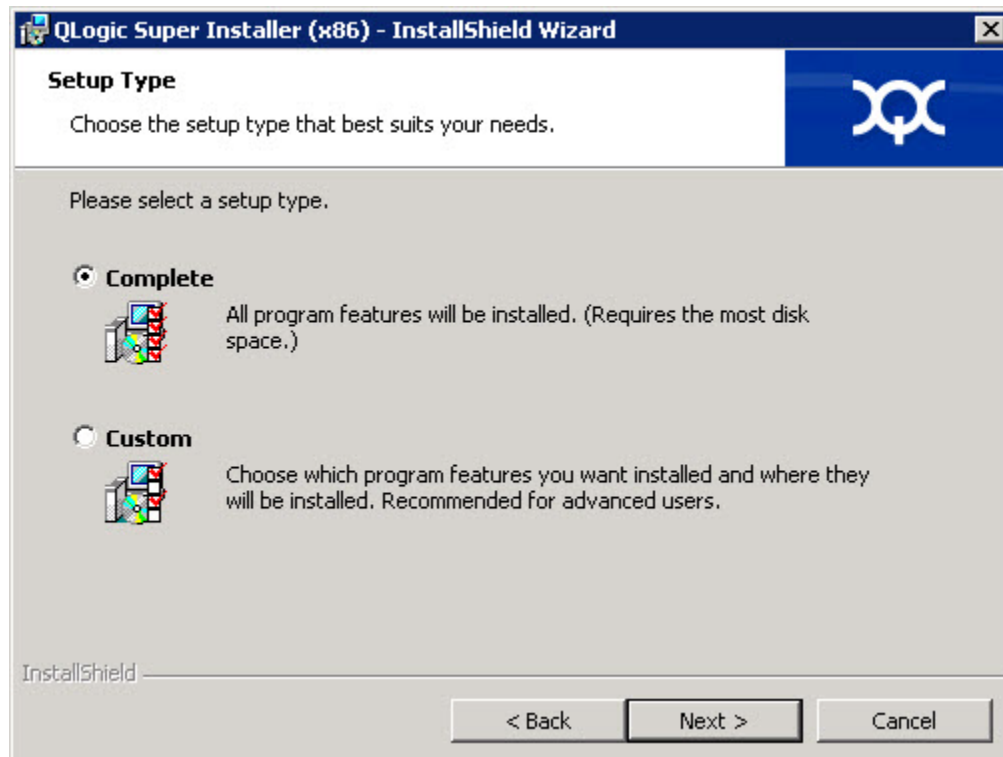
Figure 2-4. Setup Type Dialog Box

a. Select a setup type as follows:

- ❑ Select **Complete** to install *all* program features.
- ❑ Select **Custom** to manually select the features to be installed.

b. Click **Next** to continue.

If you selected **Complete**, proceed directly to Step 5.

c. The Custom Setup dialog box appears, as shown in Figure 2-5.
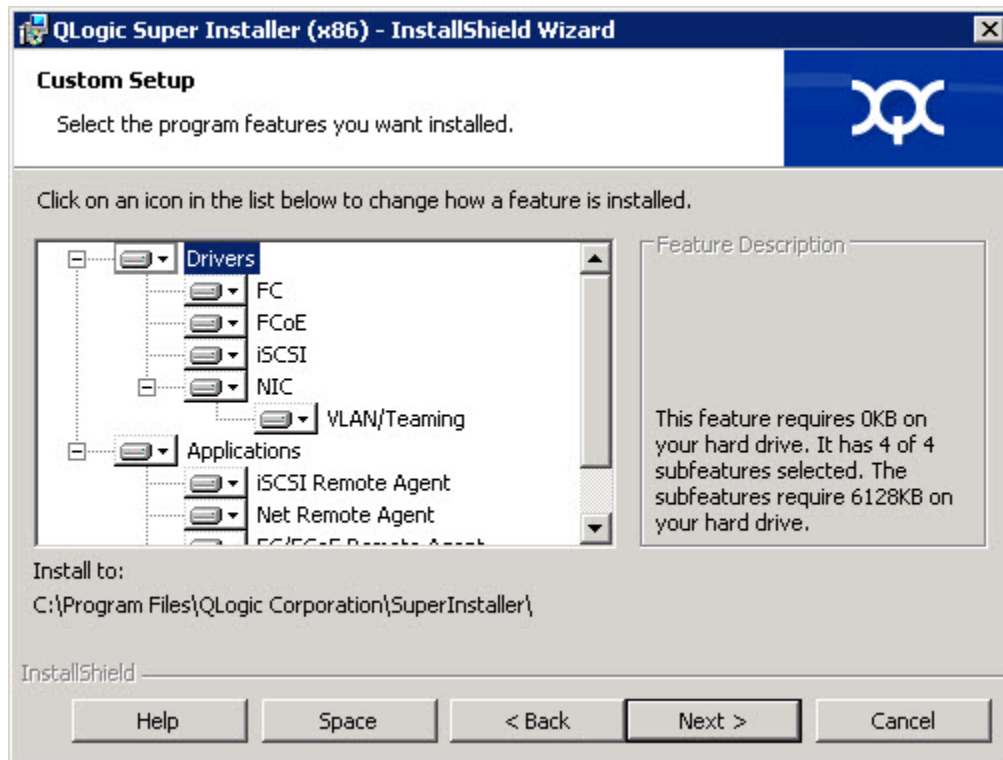
Figure 2-5. Custom Setup Dialog Box

d. 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 select one of the following:

❑ **This feature will be installed on the local hard drive**—This setting marks the feature for installation

❑ **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.

e. Click **Next** to continue.

5. The Ready to Install the Program dialog box appears, as shown in Figure 2-6.
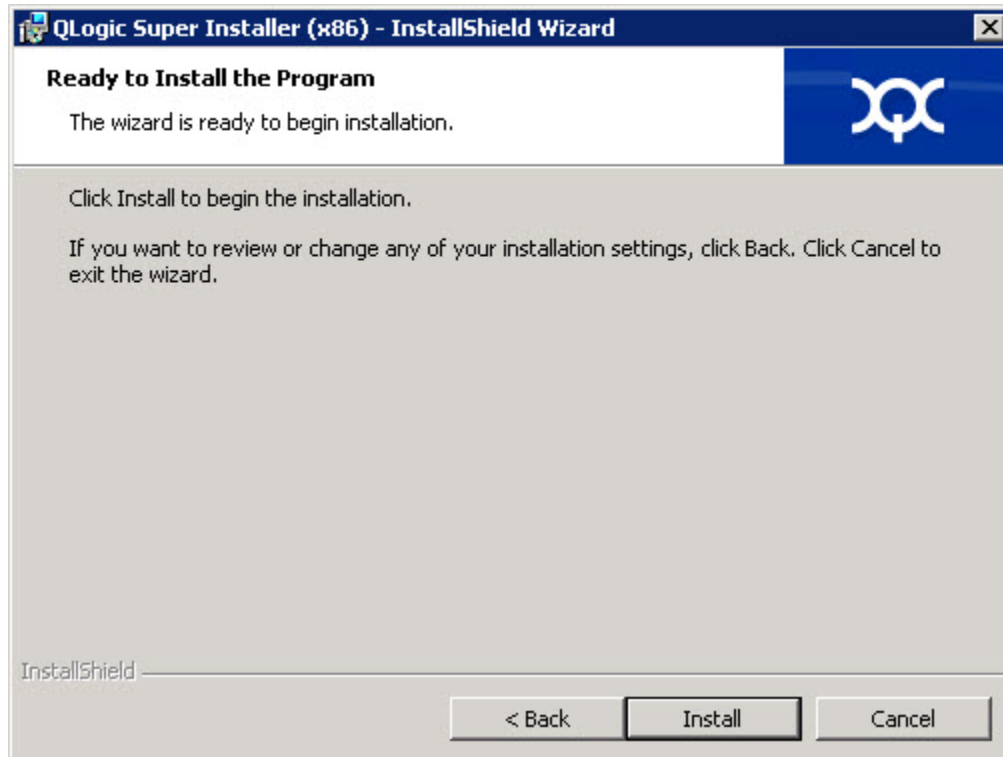
Figure 2-6. Ready to Install the Program Dialog Box

6.   Click **Install** so that the InstallShield Wizard installs the QLogic Adapter drivers and Management Software Installer.

When the installation is complete, the InstallShield Wizard Completed dialog box appears, as shown in Figure 2-7.
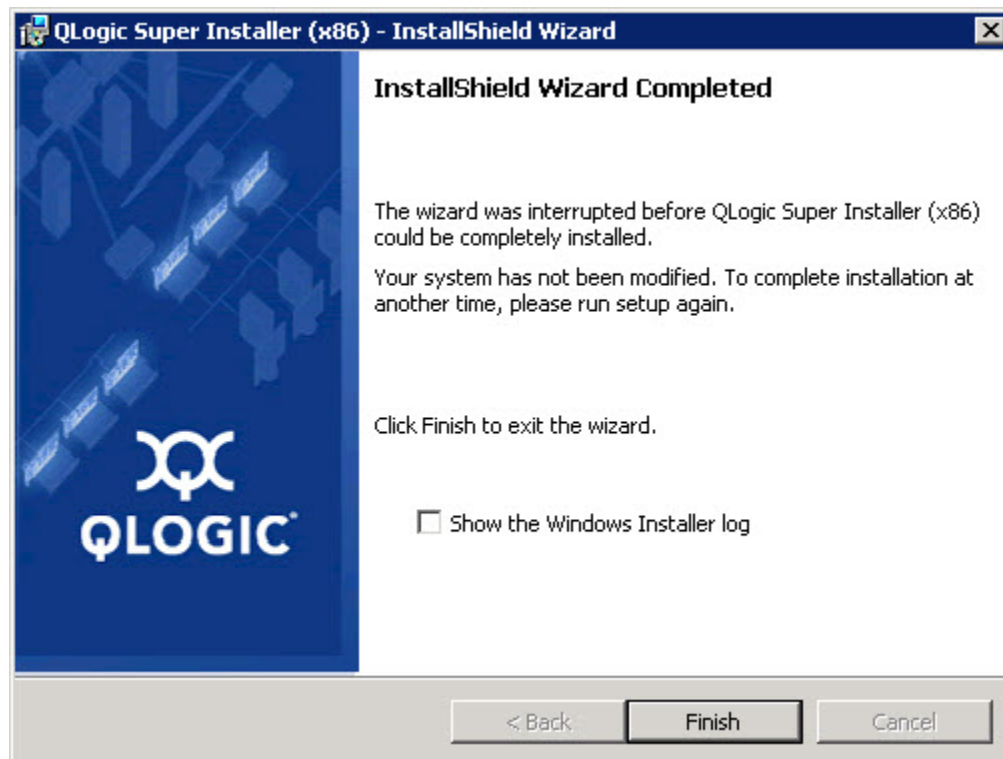
Figure 2-7. InstallShield Wizard Complete Dialog Box

7. Click **Finish** to dismiss the installer.

The Update Package window appears, as shown in Figure 2-8.

Figure 2-8. Update Package Window

8. Click **OK** to close the window.

Options

The following options can be used to customize the DUP installation behavior.

To extract only the driver components to a directory:

`/drivers=<path>`

NOTE

This command requires the `/s` option.

To install/update only the driver components:

`/driveronly`

> **NOTE**
>
> This command requires the `/s` option.

(Advanced) This command sends all text following the `/passthrough` option directly to the QLogic installation software of the DUP. This mode suppresses any provided GUI but not necessarily those of the QLogic software.

`/passthrough`

(Advanced) To return a coded description of this DUP's supported features:

`/capabilities`

> **NOTE**
>
> This command requires the `/s` option.

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 NIC Driver
- Installing the Linux iSCSI Driver
- Installing the Linux FCoE 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>
```

Installing the Linux NIC Driver

To install the Linux NIC driver, refer to the instructions (`INSTALL.qlcnic`) provided with the individual driver package.

Installing the Linux iSCSI Driver

Driver installation makes extensive use of the `build.sh` script located in the driver source (`extras/build.sh`). This section provides installation instructions for the following Linux versions:

-

-

-

Building the iSCSI Adapter Driver SLES 11 SP4

Building and Installing the Adapter Driver

1. Issue the following commands from the directory that contains the driver package file, `qla4xxx-src-x.xx.xx.xx.xx.xx-k.tar.gz`:

    ```
    # tar -xzvf qla4xxx-vx.xx.xx.xx.xx.xx-kx.tar.gz
    # cd qla4xxx-vx-x.xx.xx.xx.xx.xx-kx
    # tar -xvzf qla4xxx-src-vxx.xx.xx.xx.xx-ky.tar.gz
    # cd qla4xxx-vx.xx.xx.xx.xx.xx-kx
    ```

    where `x.xx.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 directory:

    ```
    /lib/modules/2.6.../extra/qlgc-qla4xxx
    ```

    ❑ Adds the appropriate directive in the `modprobe.conf` file (if applicable)

Manually Loading the Adapter Driver

1. Load the driver using one of the following methods:

    ❑ To directly load the driver from the local build directory, issue the following commands:

    ```
    # insmod
    /lib/modules/2.6.../kernel/drivers/scsi/scsi_transport_is
    csi2.ko
    # insmod qla4xxx.ko
    ```

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

```
# modprobe -v qla4xxx
```

2. If the iqlremote agent was previously running, restart the agent by issuing the following command (to ensure that the QConvergeConsole GUI can reconnect to this host):

```
# service iqlremote start
```

Unloading the Adapter Driver

To replace an existing inbox driver with a new out-of-box iSCSI driver, unload the existing driver and load the new driver. To unload the driver, stop all applications using the driver and then unload the driver.

1. If the iqlremote agent is running, stop the agent by issuing the following command:

```
# service iqlremote stop
```

2. To unload the driver using modprobe, issue the following command:

```
# modprobe -r qla4xxx
```

Rebuilding the RAM Disk with the New Driver

1. Edit the `/etc/modprobe.conf` file and add the following entry (if it is not present):

```
alias scsi_hostadapterX qla4xxx
```

Where X is based on the order of the SCSI modules being loaded.

2. To create a backup copy of the RAM disk image, issue the following commands:

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

3. Rebuild the initrd image by issuing the following commands:

```
# mkinitrd -f initrd-[kernel version].img `uname -r`
```

4. Reboot to boot from the new initrd image and new driver.

> **NOTE**
>
> Depending on the server hardware, the RAMDISK file name might be different.

Building the iSCSI Adapter Driver for RHEL 6.5 and SLES 12

Building and Installing the Adapter Driver

1. Issue the following commands from the directory that contains the source driver file, `qla4`*`xxx`*`-src-v`*`x.xx.xx.xx.xx.xx`*`-k.tar.gz`:

    ```
    # tar -xzvf qla4xxx-vx.xx.xx.xx.xx.xx-cx.tar.gz
    # cd qla4xxx-vx.xx.xx.xx.xx.xx-cx
    ```

    where *`x.xx.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 directory:

    ■   For RHEL 6.5:

    ```
    /lib/modules/2.6.../extra/qlgc-qla4xxx/
    ```

    ■   For SLES 12:

    ```
    /lib/modules/2.6.../updates
    ```

    ❑   Adds the appropriate directive in the `modprobe.conf` file (if applicable)

Manually Loading the Adapter Driver

1. To load the driver, use one of the following methods:

   ❑ To load the driver directly from the local build directories, issue the following commands:

   ■ For RHEL 6.5:

   ```
   # insmod /lib/modules/2.6.../kernel/drivers/scsi/
   scsi_transport_iscsi.ko
   insmod
   /lib/modules/2.6.../extra/qlgc-qla4xxx/qla4xxx.ko
   ```

   ■ For SLES 12:

   ```
   # insmod /lib/modules/2.6.../kernel/drivers/scsi/
   scsi_transport_iscsi.ko
   # insmod /lib/modules/2.6.../updates/qla4xxx.ko
   ```

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

   ```
   # modprobe -v qla4xxx
   ```

2. If the iqlremote agent was previously running, restart the agent by issuing the following command (to ensure that the QConvergeConsole GUI can reconnect to this host):

   ```
   # service iqlremote start
   ```

Unloading the Adapter Driver

To replace an existing inbox driver with a new out-of-box iSCSI driver, unload the existing driver and load the new driver. To unload the driver, stop all applications using the driver and then unload the driver.

1. If the iqlremote agent is running, stop the agent by issuing the following command:

   ```
   # service iqlremote stop
   ```

2. To unload the driver using modprobe, issue the following command:

   ```
   # modprobe -r qla4xxx
   ```

Rebuilding the RAM Disk

To automatically load the driver by rebuilding the RAM disk to include the driver, follow these steps:

1. To create a backup copy of the RAM disk image, issue the following command:

   ❑ For RHEL 6.5:

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

   ❑ For SLES 12:

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

2. Rebuild the initrd image with driver by issuing the following command:

   ❑ For RHEL 6.5:

   ```
   # mkinitrd -f /boot/initramfs-[kernel version].img 'uname -r'
   ```

   ❑ For SLES 12:

   ```
   # mkinitrd
   ```

3. Reboot the host to boot from the new initrd image with new driver.

> **NOTE**
>
> Depending on the server hardware, the RAMDISK file name might be different.

Building the iSCSI Adapter Driver for RHEL 6.5 and SLES 11 SP3

Building and Installing the Adapter Driver

1. Issue the following commands from the directory that contains the source driver file, `qla4xxx-src-vx.xx.xx.xx.xx.xx-k.tar.gz`:

   ```
   # tar -xzvf qla4xxx-vx.xx.xx.xx.xx.xx-cx.tar.gz
   # cd qla4xxx-vx.xx.xx.xx.xx.xx-cx
   ```

   where `x.xx.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 directory:

■   For RHEL 6.5:

```
/lib/modules/2.6.../extra/qlgc-qla4xxx/
```

■   For SLES 11 SP3:

```
/lib/modules/2.6.../updates
```

❑   Adds the appropriate directive in the `modprobe.conf` file (if applicable)

Manually Loading the Adapter Driver

1.   To load the driver, use one of the following methods:

❑   To load the driver directly from the local build directories, issue the following commands:

■   For RHEL 6.5:

```
# insmod /lib/modules/2.6.../kernel/drivers/scsi/
scsi_transport_iscsi.ko
insmod
/lib/modules/2.6.../extra/qlgc-qla4xxx/qla4xxx.ko
```

■   For SLES 11 SP3:

```
# insmod /lib/modules/2.6.../kernel/drivers/scsi/
scsi_transport_iscsi.ko
# insmod /lib/modules/2.6.../updates/qla4xxx.ko
```

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

```
# modprobe -v qla4xxx
```

2.   If the iqlremote agent was previously running, restart the agent by issuing the following command (to ensure that the QConvergeConsole GUI can reconnect to this host):

```
# service iqlremote start
```

Unloading the Adapter Driver

To replace an existing inbox driver with a new out-of-box iSCSI driver, unload the existing driver and load the new driver. To unload the driver, stop all applications using the driver and then unload the driver.

1. If the iqlremote agent is running, stop the agent by issuing the following command:

   ```
   # service iqlremote stop
   ```

2. To unload the driver using modprobe, issue the following command:

   ```
   # modprobe -r qla4xxx
   ```

Rebuilding the RAM Disk

To automatically load the driver by rebuilding the RAM disk to include the driver, follow these steps:

1. To create a backup copy of the RAM disk image, issue the following command:

   ❑ For RHEL 6.5:

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

   ❑ For SLES 11 SP3:

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

2. Rebuild the initrd image with driver by issuing the following command:

   ❑ For RHEL 6.5:

   ```
   # mkinitrd -f /boot/initramfs-[kernel version].img 'uname -r'
   ```

   ❑ For SLES 11 SP3:

   ```
   # mkinitrd
   ```

3. Reboot the host to boot from the new initrd image with new driver.

   NOTE

   Depending on the server hardware, the RAMDISK file name might be different.

Installing the Linux FCoE Driver

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

- Building the Driver for RHEL 6.5 Linux
- Building the Driver for SLES 11 SP4 Linux
- Building the Driver for SLES 12 Linux
- Building the Driver for SLES 11 SP3 Linux

Building the Driver for RHEL 6.5 Linux

1. Issue the following commands from the directory that contains the source driver file, `qla2xxx-src-x.xx.xx.xx.xx.xx-k.gz`:

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

   where `x.xx.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/2.6.../extra/qlgc-qla2xxx` directory.

3. Manually load the driver for Linux by issuing the following command:

   ```
   # modprobe -v qla2xxx
   ```

   To unload the driver, issue the following command:

   ```
   # modprobe -r qla2xxx
   ```

4. 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 11 SP4 Linux

1. Issue the following commands from the directory that contains the source driver file, `qla2xxx-src-vx.xx.xx.xx.xx.x-k4.tar.gz`:

```
# tar -xzvf qla2xxx-src-vx.xx.xx.xx.xx.x-k4.tar.gz
# cd qla2xxx-x.xx.xx.xx.xx.x-k4
```

where `x.xx.xx.xx.xx.x` 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/2.6.../updates` directory.

❑ Adds the appropriate directive in the `modprobe.conf` file (if applicable).

3. Manually load the driver for Linux.

❑ 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. To automatically load the driver each time the system boots, rebuild the RAM disk to include the driver as follows:

    a.   Edit the `/etc/sysconfig/kernel` file to modify the INITRD_MODULES directive and append `qla2xxx` to the string. For example:

```
INITRD_MODULES=".... qla2xxx"
```

where `qla2xxx` is appended to the end of the directive.

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

```
# cd /boot
# cp initrd-[kernel version] initrd-[kernel version].bak
# mkinitrd
```

> **NOTE**
>
> Depending on the server hardware, the RAMDISK file name might be different.

    c.   To load the driver, reboot the host.

Building the Driver for SLES 12 Linux

1. In the directory that contains the source driver file, `qla2`*`xxx`*`-src-v`*`x.xx.xx.xx`*`.11.`*`x`*`-k.tgz`, issue the following commands:

   ```
   # tar -xzvf qla2xxx-src-vx.xx.xx.xx.11.x-k.tgz
   # cd qla2xxx-x.xx.xx.xx.xx.xx-k
   ```

   where *`x.xx.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/3.x.../updates` directory.

   ❑ Adds the appropriate directive in the `modprobe.conf` file (if applicable).

3. 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` by `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. To automatically load the driver each time the system boots, rebuild the RAM disk to include the driver.

Create a copy of the current RAMDISK by issuing the following commands:

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

> **NOTE**
>
> Depending on the server hardware, the RAMDISK file name might be different.

5. To load the driver, reboot the host.

Building the Driver for SLES 11 SP3 Linux

1. In the directory that contains the source driver file, `qla2xxx-src-vx.xx.xx.xx.11.x-k.tgz`, issue the following commands:

```
# tar -xzvf qla2xxx-src-vx.xx.xx.xx.11.x-k.tgz
# cd qla2xxx-x.xx.xx.xx.xx.xx-k4
```

where `x.xx.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/3.x.../updates` directory.

❑ Adds the appropriate directive in the `modprobe.conf` file (if applicable).

3. 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` by `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.   To automatically load the driver each time the system boots, rebuild the
RAM disk to include the driver.

Create a copy of the current RAMDISK by issuing the following commands:

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

> **NOTE**
>
> Depending on the server hardware, the RAMDISK file name might be
> different.

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 5.x NIC Driver
- Installing the ESXi 5.x iSCSI Driver
- Installing the ESXi 5.x FCoE Driver
- Installing the ESXi 6.x Fibre Channel Over Ethernet Driver
- Installing the ESXi 6.x iSCSI Driver
- Installing the QConvergeConsole VMware vCenter Server Plug-in
- Installing the vSphere Web Client Plug-in

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 5.*x* NIC Driver

The operating system manages and controls the driver installation process. To install the ESXi 5.*x* 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, refer to the following:

http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2005205

This section provides procedures for the following:

- Updating an Existing Driver or Installing a New Driver for an Existing ESXi Installation with esxcli (ESXi 5.x Only)
- Verifying the Version of the Installed Driver (ESXi 5.x Only)

For other installation procedures, consult the operating system manuals and the driver readme file for more details.

Updating an Existing Driver or Installing a New Driver for an Existing ESXi Installation with esxcli (ESXi 5.*x* Only)

To use the driver bundle (`<offline-bundle>`):

1. Copy the driver bundle (`<offline-bundle>`) to this ESXi host.

2. Install the driver bundle (`<offline-bundle>`) using the following steps:

    a. Type the following command to make a temporary directory:

    ```
    mkdir /install; cd /install
    ```

    b. Unzip the driver bundle in the temporary directory:

    ```
    /install : unzip <offline-bundle>
    ```

    c. Run the following command:

    ```
    esxcli software vib install –d /install/offline-bundle.zip
    ```

To use the driver VIB:

1. Copy the driver VIB
    (`net-<offline-bundle>-<driver-version>.0.0.<esx-build>.x86_64.vib`) to this ESXi host.

2. Install the driver VIB using the following `esxcli` commands:

    a. Type the following command to make a temporary directory:

    ```
    mkdir /install; cd /install
    ```

    b. Run the following command:

    ```
    esxcli software vib install –v /install/<driver-vib>
    ```

Verifying the Version of the Installed Driver (ESXi 5.*x* Only)

Verify the installed package in the system using the following command:

```
esxcli software vib list | grep -i driver version
```

The driver version is embedded in the VIB version.

For example, the output looks like the following:

```
esxcli software vib list | grep qlc
net-qlcnic              5.1.132-1OEM.500.0.0.472560       VMware
VMwareCertified   2012-12-19
```

Installing the ESXi 5.*x* iSCSI Driver

The operating system manages and controls the driver installation process. To install the ESXi 5.*x* 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, refer to the following:
>
> http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2005205

This section provides procedures for the following:

■ Updating an Existing Driver or Installing a New Driver for an Existing ESXi Installation with esxcli (ESXi 5.x Only)

■ Verifying the Version of the Installed Driver (ESXi 5.x Only)

For other installation procedures, consult the operating system manuals and the driver readme file for more details.

Updating an Existing Driver or Installing a New Driver for an Existing ESXi Installation with esxcli (ESXi 5.*x* Only)

To use the driver bundle (`<offline-bundle>`):

1. Copy the driver bundle (`<offline-bundle>`) to this ESXi host.

2. Install the driver bundle (`<offline-bundle>`) using the following steps:

   a. Type the following command to make a temporary directory:

   ```
   mkdir /install; cd /install
   ```

   b. Unzip the driver bundle in the temporary directory:

   ```
   /install : unzip <offline-bundle>
   ```

   c. Run the following command:

   ```
   esxcli software vib install -d /install/offline-bundle.zip
   ```

To use the driver VIB:

1. Copy the driver VIB
    (`scsi-<offline-bundle>-<driver-version>`.0.0.`<esx-build>`.x
    `86_64.vib`) to this ESXi host.

2. Install the driver VIB using the following `esxcli` commands:

    a. Type the following command to make a temporary directory:

        ```
        mkdir /install; cd /install
        ```

    b. Run the following command:

        ```
        esxcli software vib install -v /install/<driver-vib>
        ```

Verifying the Version of the Installed Driver (ESXi 5.*x* Only)

Verify the installed package in the system using the following command:

```
esxcli software vib list | grep -i driver version
```

The driver version is embedded in the VIB version.

For example, the output looks like the following:

```
# esxcli software vib list | grep qla4xxx
scsi_qla4xxx .01.03.2-6vmw.550.0.0.1014658 VMware VMwareCertified
2013-02-2
```

Installing the ESXi 5.*x* FCoE Driver

The operating system manages and controls the driver installation process. To install the ESXi 5.*x* 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, refer to the following:

http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2005205

This section provides procedures for the following:

■ Updating an Existing Driver or Installing a New Driver for an Existing ESXi Installation with esxcli (ESXi 5.x Only)

■ Verifying the Version of the Installed Driver (ESXi 5.x Only)

For other installation procedures, consult the operating system manuals and the driver readme file for more details.

Updating an Existing Driver or Installing a New Driver for an Existing ESXi Installation with esxcli (ESXi 5.*x* Only)

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

1. Copy the driver bundle (`<offline-bundle>.zip`) to this ESXi host.

2. Install using the driver bundle (`<offline-bundle>.zip`):

   a. Type the following command to make a temporary directory:

   ```
   $ mkdir /install
   $ mv <offline-bundle>.zip /install
   $ cd install
   ```

   b. Unzip the driver bundle in the temporary directory:

   ```
   $ unzip <offline-bundle>.zip
   ```

   c. Run one of the following commands.

   For ESX 5.0/5.1:

   ```
   esxcli software vib install -n scsi-qla2xxx -d
   /install/offline-bundle.zip
   ```

   For ESX 5.5:

   ```
   esxcli software vib install -n qlnativefc -d
   /install/offline-bundle.zip
   ```

To use the driver VIB:

1. Copy the driver VIB (for ESX 5.0/5.1:
   `scsi-qla2xxx-<driver-version>.0.0.<esx-build>.x86_64.vib`;
   for ESX 5.5:
   `qlnativefc-<driver-version>.0.0.<esx-build>.x86_64.vib`)
   to this ESXi host.

2. Install the driver VIB using the following `esxcli` commands:

   a. Type the following command to make a temporary directory:

   ```
   mkdir /install; cd /install
   ```

   b. Run the following command:

   ```
   $ esxcli software vib install -v <driver-vib-file>
   ```

   For example:

   ```
   esxcli software vib install -v
   /vmfs/volumes/datastore1/scsi-qla2xxx-934.5.10.0-1OEM.500
   .0.0.472560.x86_64.vib
   ```

Verifying the Version of the Installed Driver (ESXi 5.*x* Only)

Verify the installed package in the system using the following command:

```
esxcli software vib list | grep -i <driver-version / driver name>
```

The driver version is embedded in the VIB version.

For example, the output looks like the following:

```
# esxcli software vib list | grep qla2xxx
scsi-qla2xxx        911.k1.1-16vmw.500.0.0.406165    VMware
VMwareCertified     2011-09-21
```

Installing the ESXi 6.x Fibre Channel Over Ethernet Driver

Updating an Existing Driver or Installing a New Driver for an Existing ESXi Installation with esxcli (for ESXi 6*x* Only)

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) using the following steps:

   a. Type the following command to make a temporary directory:

   ```
   $ mkdir /install $ mv <offline-bundle>.zip /install $ cd
   install
   ```

   b. Unzip the driver bundle in the temporary directory:

   ```
   $ unzip <offline-bundle>.zip
   ```

   c. Run one of the following commands.

   For ESX 6.x:

   ```
   esxcli software vib install -n qlnativefc -d /install
   ```

To use the driver VIB:

1. Copy the driver VIB (for ESX 6.0:
   qlnativefc-<driver-version>.0.0.<esx-build>.x86_64.vib)
   to this ESXi host.

2. Install the driver VIB using the following `esxcli` commands:

   a. Type the following command to make a temporary directory:

   ```
   mkdir /install; cd /install
   ```

   b. Run the following command:

   ```
   esxcli software vib install -v <driver-vib-file>
   ```

   For example, the output looks like the following:

   ```
   esxcli software vib install -v
   /vmfs/volumes/datastore1/qlnativefc-2.1.23.0-1OEM.6
   00.0.0.2159203.x86_64.vib
   ```

Installing the ESXi 6.x iSCSI Driver

Updating an Existing Driver or Installing a New Driver for an Existing ESXi Installation with esxcli (for ESXi 6*x* Only)

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) using the following steps:

    a. Type the following command to make a temporary directory:

    ```
    $ mkdir /install $ mv <offline-bundle>.zip /install $ cd
    install
    ```

    b. Unzip the driver bundle in the temporary directory:

    ```
    $ unzip <offline-bundle>.zip
    ```

    c. Run one of the following commands.

    For ESX 6.x:

    ```
    esxcli software vib install -n scsi-qla4xxx -d /install
    ```

To use the driver VIB:

1. Copy the driver VIB (for ESX 6.0:
    `scsi-qla4xxx_-<driver-version>.<esx-build>.vib)` to this
    ESXi host.

2. Install the driver VIB using the following `esxcli` commands:

    a. Type the following command to make a temporary directory:

    ```
    mkdir /install; cd /install
    ```

    b. Run the following command:

    ```
    esxcli software vib install -v <driver-vib-file>
    ```

    For example, the output looks like the following:

    ```
    esxcli software vib install -v
    /vmfs/volumes/datastore1/scsi-qla4xxx_644.6.04.0-1O
    EM.600.0.0.2159203.vib
    ```

Installing the QConvergeConsole VMware vCenter Server Plug-in

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

1. QConvergeConsole VMware vCenter Server Plug-in—on the vCenter Server

2. QLogic Adapter CIM Provider—on the ESX/ESXi server

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

- Installation Package Contents
- QConvergeConsole VMware vCenter Server Plug-in Installation
- Plug-in Unregistration from a Manual Install
- Enabling and Disabling the Plug-in
- 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 "QConvergeConsole VMware vCenter Server Plug-in Installation" on page 36.

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. The files are as follows (`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_provider_vmware_esx50x-x.x.x`

  This file contains the QLogic Adapter CIM Provider installation package for ESXi 5.0.*x*/5.1.*x*, where `x.x.x` is the version of the CIM Provider.

- `qlogic_adapter_provider_vmware_esx55_60-x.x.x`

  This file contains the QLogic Adapter CIM Provider installation package for ESXi 5.5, where `x.x.x` is the version of the CIM Provider.

- `readme.txt`

  This file is the Read Me document containing 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 contains the release notes that list changes, fixes, known issues, and release details.

For detailed information on installing the QConvergeConsole VMware vCenter Server Plug-in, refer to "QConvergeConsole VMware vCenter Server Plug-in Installation" on page 36.

For detailed information on installing the CIM Provider, refer to "Installing the QLogic Adapter CIM Provider" on page 43.

QConvergeConsole VMware vCenter Server Plug-in Installation

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. Run the installation 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.

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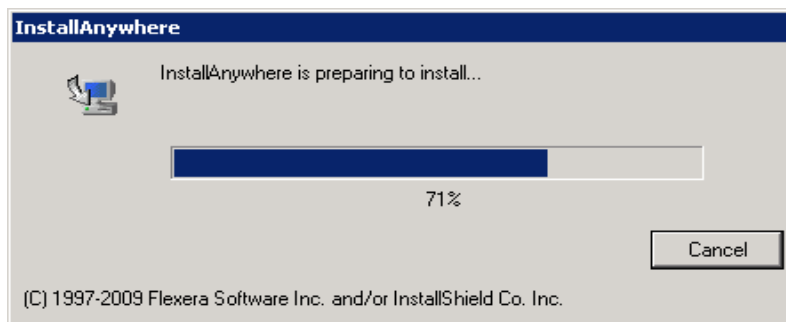
Figure 2-9. InstallAnywhere Initial Window

4.   The Plug-in Registration Wizard opens, as shown in Figure 2-10. Click **Next**.
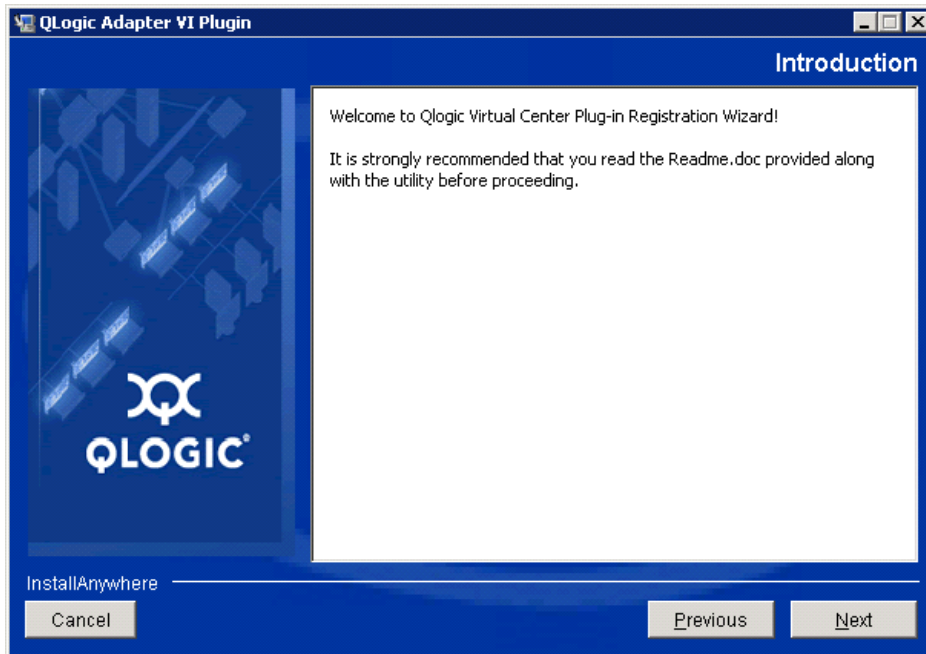
Figure 2-10. QConvergeConsole VMware vCenter Server Plug-in Registration Wizard

5.   Wait while the wizard configures the plug-in (see Figure 2-11).
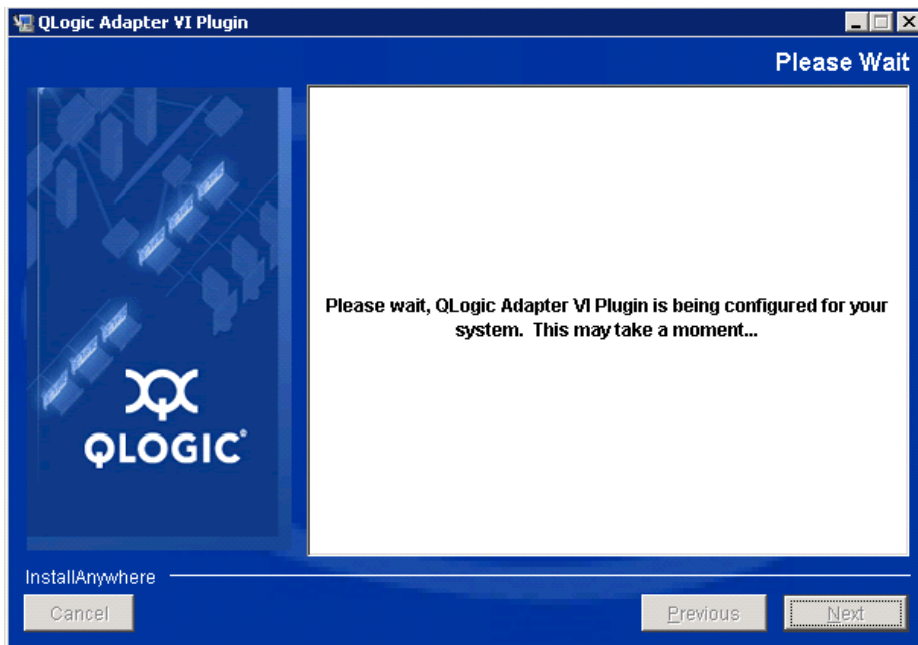
Figure 2-11. QConvergeConsole VMware vCenter Server Plug-in Configuration

6. Select the installation directory and then click **Install** (see Figure 2-12).
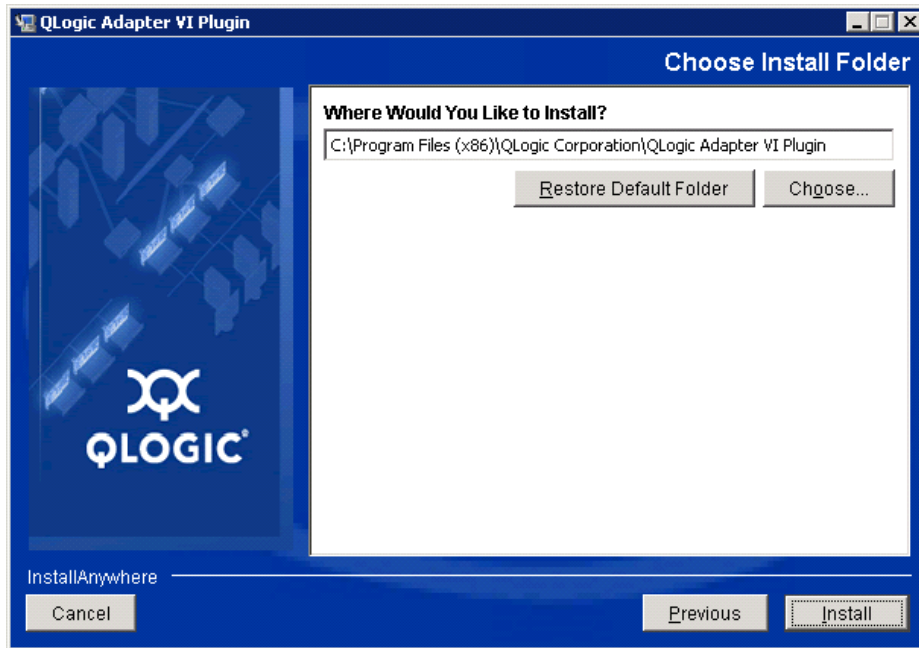
Figure 2-12. Select the Installation Directory

7. Wait while the wizard performs the installation (see Figure 2-13).
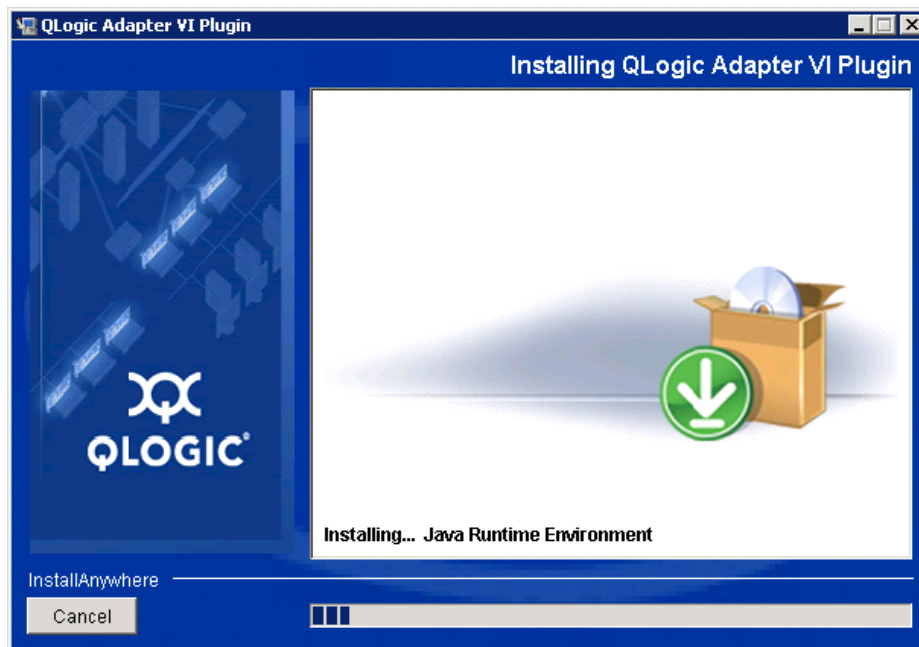
Figure 2-13. Installing the Plug-In

8.   Type in the requested information and then click **Next** (see Figure 2-14).
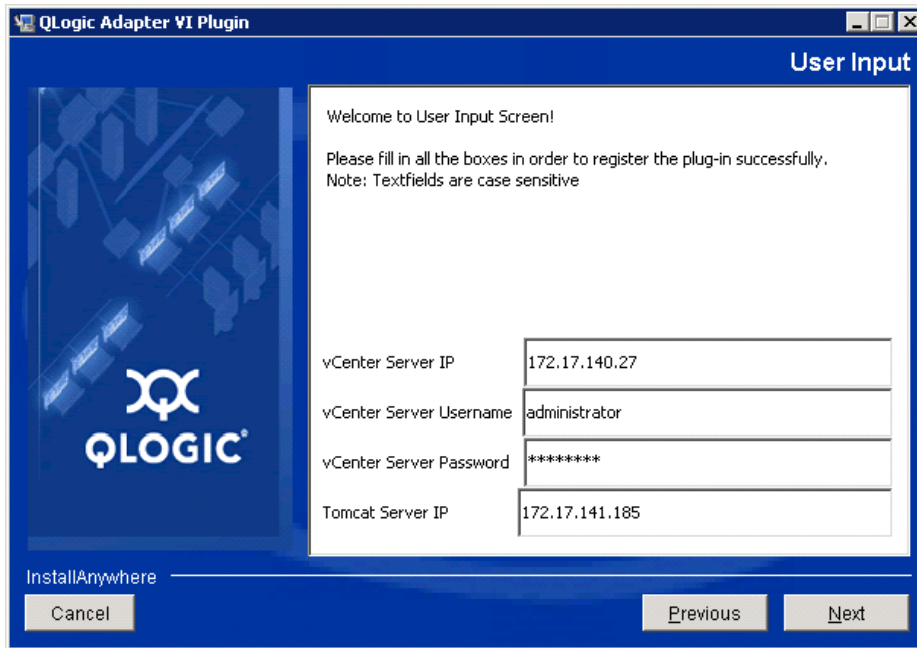
Figure 2-14. User Input Screen

9.   Wait while the wizard finishes configuring the plug-in (see Figure 2-15).
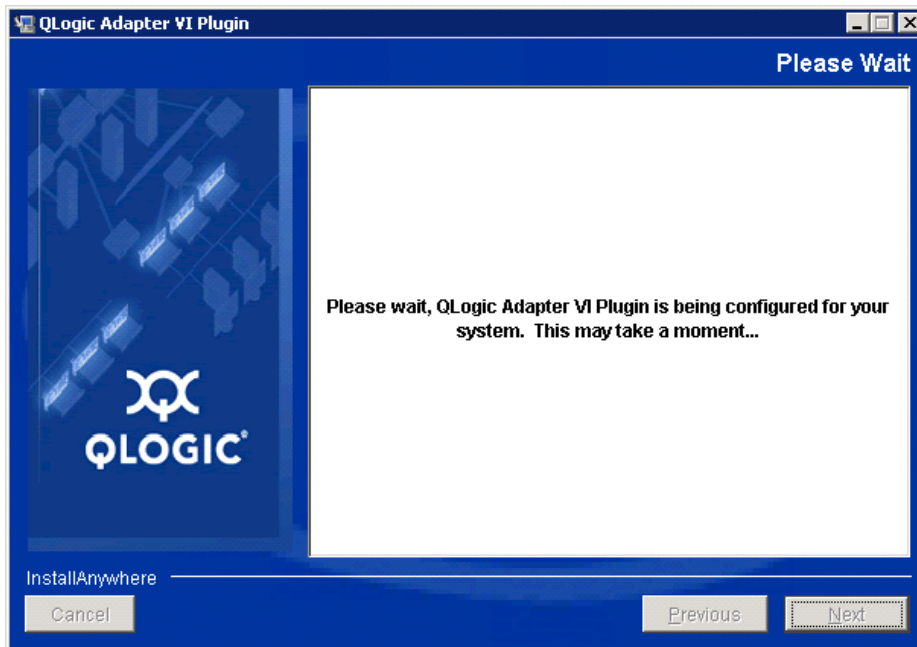
Figure 2-15. QConvergeConsole VMware vCenter Server Plug-in Configuration

10. Figure 2-16 appears when registration is completed. Click **Finish** to exit.
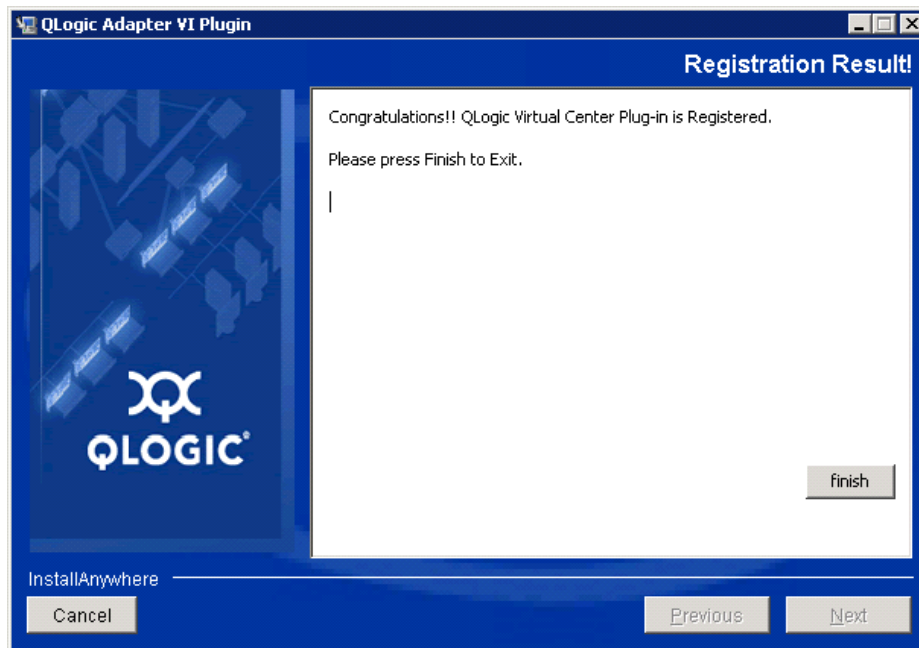
Figure 2-16. Successful Registration

11. 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 type 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 need to download the appropriate VI SDK from VMware:

■ For Perl VMware Infrastructure (VI) software development kit (SDK):
vSphere SDK for Perl
http://www.vmware.com/support/developer/viperltoolkit/

■ For PowerShell: 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"
```

Replace *administrator* and *password* with the correct information to log into the vCenter Server.

Enabling and Disabling the Plug-in

> **NOTE**
>
> If the plug-in installation completed successfully, you do not need to enable the plug-in; it is automatically enabled during installation. You can, however, verify if the plug-in is enabled by using the following procedure.

To enable or disable the QConvergeConsole VMware vCenter Server plug-in, follow these steps:

1. In the vSphere Client window, click **Plug-ins** and then click **Manage Plug-ins**.

    The Plug-in Manager window appears, as shown in Figure 2-17.
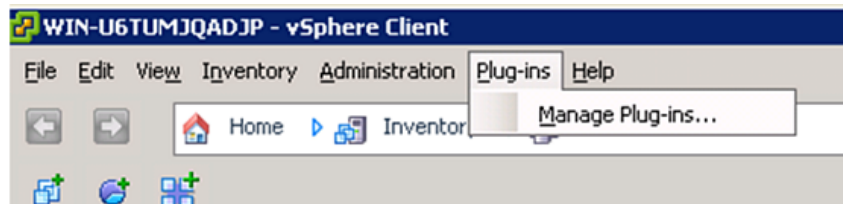
Figure 2-17. Managing Plug-ins in vSphere Client

2. Locate the QConvergeConsole vCenter Server plug-in on the Installed Plug-ins section of the window.

    The plug-in's status (Enabled or Disabled) is displayed in the Status column, as shown in Figure 2-18.
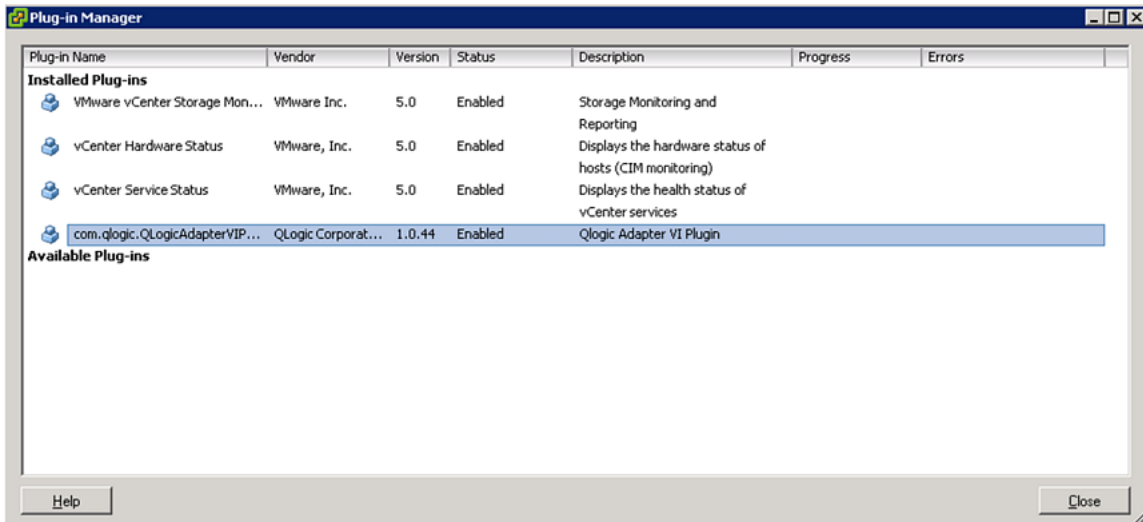
Figure 2-18. QConvergeConsole vCenter Server in Plug-in Manager

3.   If you want to enable or disable the QConvergeConsole plug-in, right-click on the plug-in and select **Enabled** or **Disabled** (the status toggles between the two), as shown in Figure 2-19.

4.   Click **Close** to close the Plug-in Manager window.
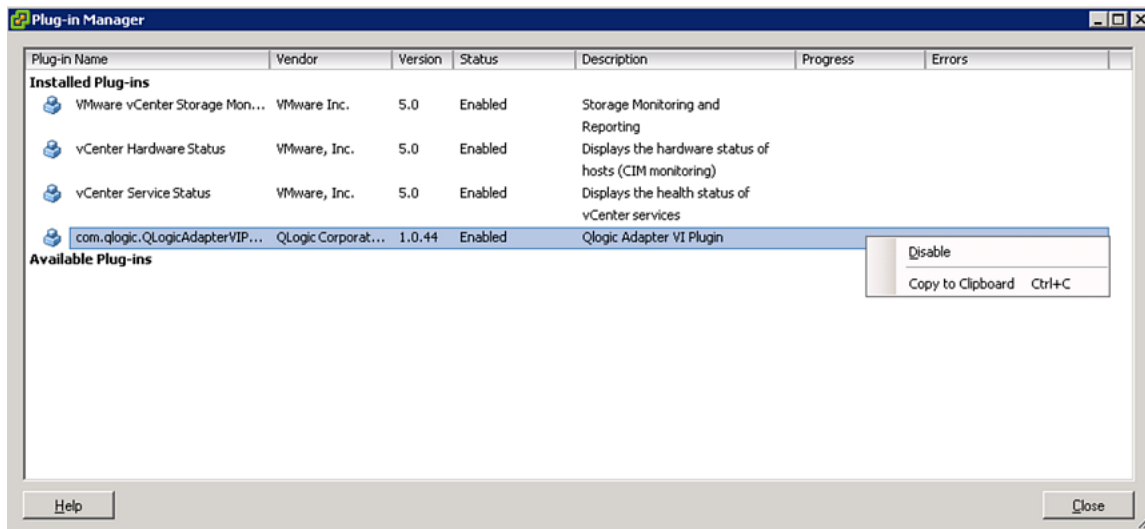
Figure 2-19. Toggling the QConvergeConsole vCenter Server Plug-in Status

Uninstalling the QConvergeConsole VMware vCenter Server Plug-in

To remove the QConvergeConsole VMware vCenter Server Plug-in:

1. In the Windows Control Panel, select **Add or Remove Programs**. (Windows Server 2008 or later only: 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. Follow the instructions in the QConvergeConsole VMware vCenter Server Plug-in installer to remove the plug-in.

Installing the QLogic Adapter CIM Provider

This section describes how to install, start, and remove the QLogic Adapter CIM Provider for VMware ESX and ESXi. There is more than one zip package, so make sure that you pick the zip package that matches your environment—ESXi 5.0, and ESXi 5.1.

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

Initial Installation Methods

Initial installation methods for the CIM Provider include the following:

- Online

  Refer to "Installing the CIM Provider on an ESXi 5.x Host" on page 44 or "Installing the CIM Provider on an ESXi 5.5 Host" on page 44.

- Offline

  Refer to "Existing ESX/ESXi Installation Using VMware Update Manager" on page 44.

- Remote

  Refer to "Remote Installation of the CIM Provider on an ESX/ESXi Host" on page 45.

Installing the CIM Provider on an ESXi 5.*x* Host

1. Copy the `provider-adapter.vib` file to the root directory (/) of the ESXi 5.*x* system.

2. Issue the `esxcli` commands as follows:

```
# cd /
# esxcli software acceptance set --level=CommunitySupported
# esxcli software vib install -v file:/provider-adapter.vib
--maintenance-mode --no-sig-check
```

3. Reboot the system as required.

Installing the CIM Provider on an ESXi 5.5 Host

1. Copy the `qlogic-adapter-provider.zip` file to the root directory (/) of the ESXi 5.5 system.

2. Issue the `esxcli` commands as follows:

```
# cd /
# esxcli software acceptance set --level=CommunitySupported
# esxcli software vib install -d
file:/qlogic-adapter-provider.zip --maintenance-mode
--no-sig-check
```

3. Reboot the system as required.

Existing ESX/ESXi Installation Using VMware Update Manager

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

To install the asynchronous drivers:

1. Extract the contents of the asynchronous driver zip file.

2. Identify the `offline-bundle.zip` file(s).

3. From vCenter Server, select **Home** and then select **Update Manager**.

4. Click the **Patch Repository** tab.

5. Click the **Import Patches** link at the top right of the screen.

6. Click **Finish**.

   The asynchronous driver is now added to the patch repository.

7. Create a baseline and remediate the ESX/ESXi host. For more information, refer to *Installing and Administering VMware vSphere Update Manager* at http://www.vmware.com/support/pubs/vum_pubs.html.

Remote Installation of the CIM Provider on an ESX/ESXi Host

> **NOTE**
>
> Before performing this procedure, ensure that the remote ESX/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 the vMA is hosted.

2. Navigate to the location of the `offline-bundle.zip` file.

3. Run the `vihostupdate` command to install the offline bundle 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 ESX/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 46 to remove the existing VIB. After completing the VIB removal, use the same steps in "Initial Installation Methods" on page 43 to install the new VIB.

Starting the QLogic Adapter CIM Provider

After a system startup, the Small Footprint CIM Broker (SFCB) CIM object manager (CIMOM) in the ESX system starts automatically and loads 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 using 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.

Uninstalling the CIM Provider from an ESXi 5.*x* Host

1. Type the following command to view the VIB list:

    ```
    # esxcli software vib list
    ```

2. Type the following command to remove the QLogic Adapter CIM Provider:

    ```
    # esxcli software vib remove --vibname qlogic-adapter-provider
    --maintenance-mode -f
    ```

Uninstalling the CIM Provider from a Remote Host

> **NOTE**
>
> Before performing this procedure, make sure that the ESX/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 typing 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*:
>
> http://www.vmware.com/pdf/vsphere4/r40/vsp_40_vcli.pdf

Installing the vSphere Web Client Plug-in

1. Gather all information necessary for the installation

   ■ IP address of the vCenter Server

   ■ vCenter Server credentials (user name and password)

   ■ Where to host the QLogic Adapter vSphere Web Client Plug-in (on vCenter Server or other server)

     If you are hosting the vSphere Web Client Plug-in on a non-vCenter Server, make sure the server has Tomcat running as a service and have the IP address of the Tomcat instance ready. Also, make sure the Tomcat `CATALINA_HOME` environment variable is set to the appropriate directory.

2. Run the installer on the server providing the Tomcat service. Provide the information requested by the installer.

   ■ On Windows, double-click on the installer and follow the instructions on the GUI provided.

   ■ On Linux:

     a. Make sure the user is the root user (or has root privileges).

     b. Create the installer executable if one does not already exist. Choose the installer for your system (32-bit or 64-bit), and type the following command:

        ```
        chmod +x <installer>
        ```

        Where `<installer>` is the file name of the installer.

     c. Run the installer by issuing the following command:

        ```
        ./<installer>
        ```

        Where "`<installer>`" is the file name of the installer.

     d. Follow the instructions provided by the installer.

3. Restart the Tomcat service.

   If the vSphere Web Client Plug-in is being hosted on the vCenter Server, you must restart the VMware Virtual Center Management Web services. In Windows, go to the **Administrative Tools** menu, select **Services**, and restart VMware Virtual Center Management Web services. On the vCenter Server Appliance (Linux), issue the following command:

   ```
   /etc/init.d/vmware-vpxd tomcat-restart
   ```

4. Restart any vSphere Web Client sessions.

If you are updating a previous version of the vSphere Web Client Plug-in, restart the vSphere Web Client services. In Windows, go to the **Administrative Tools** menu, select **Services**, and restart VMware vSphere Web Client. On the vCenter Server Appliance (Linux), issue the following command:

```
/etc/init.d/vsphere-client restart
```

Uninstalling the vSphere Web Client Plug-in

- Uninstalling the vSphere Web Client Plug-in on Windows is initiated through the Windows Uninstall Programs control panel. Follow the uninstaller user interface to uninstall the plug-in.

- Uninstalling the vSphere Web Client Plug-in on Linux is initiated by the following command line command:

```
/opt/qlogic/QLogic\ Adapter\ Web\ Client\
Plugin/Uninstall_QLogic\ Adapter\ Web\ Client\
Plugin/Uninstall\ QLogic\ Adapter\ Web\ Client\ Plugin
```

Follow the prompts (user interface or console commands) to uninstall the plug-in by the root user.

3 Adapter Management Applications

Overview

This chapter describes the following adapter management applications:

- General Management with QConvergeConsole
- Switch Independent Partitioning—refer to Chapter 4
- Windows Management Applications
- Linux Management Applications
- VMware Management Applications

General Management with QConvergeConsole

Use the QConvergeConsole GUI and CLI utilities to manage the adapter as follows:

- Configuring the NIC Driver with QConvergeConsole
- Configuring iSCSI with QConvergeConsole
- Configuring FCoE with QConvergeConsole

NOTE

For information on installing and starting the QConvergeConsole GUI, refer to the *QConvergeConsole GUI Installation Guide* (for download instructions, see "Related Materials" on page xii). All procedural information for the QConvergeConsole GUI is covered in the QConvergeConsole GUI's online help system.

Configuring the NIC Driver with QConvergeConsole

For information on configuring the NIC driver using the QConvergeConsole GUI, refer to the *QConvergeConsole GUI Help System* and select **Managing Ethernet (NIC) Ports**.

For information on configuring the NIC driver using the QConvergeConsole CLI, refer to the "NIC Interactive Commands" chapter of the *QConvergeConsole CLI User's Guide*.

Configuring iSCSI with QConvergeConsole

For information on configuring iSCSI using the QConvergeConsole GUI, refer to the *QConvergeConsole GUI Help System* (see "Related Materials" on page xii) and select **Managing iSCSI Ports**.

For information on configuring iSCSI using the QConvergeConsole CLI, refer to the following:

- Configuring FCoE with QConvergeConsole
- Configuring iSCSI Initiators with QConvergeConsole
- Enabling CHAP Authentication with QConvergeConsole
- All other topics: Refer to the *QConvergeConsole CLI User's Guide*

Configuring FCoE with QConvergeConsole

For information on configuring FCoE using the QConvergeConsole GUI, refer to the *QConvergeConsole GUI Help System* and select **Managing Fibre Channel and FCoE Adapters and Ports**.

For information on configuring FCoE using the QConvergeConsole CLI, refer to the "Fibre Channel Interactive Commands" chapter of the *QConvergeConsole CLI User's Guide*.

Configuring iSCSI Offload with QConvergeConsole

The iSCSI offload feature provides full iSCSI offloads that include header and data digest, receive protocol data unit (PDU) parsing, and direct data placement. You can configure iSCSI offload parameters with the following utilities:

■ QConvergeConsole GUI: graphical user interface

■ QConvergeConsole CLI: *interactive mode* (menu driven) and *non-interactive mode* (command-line driven)

> **NOTE**
>
> For information on installing and starting the QConvergeConsole GUI, refer to the *QConvergeConsole GUI Installation Guide* (for download instructions, see "Related Materials" on page xii). All procedural information for the QConvergeConsole GUI is covered in the QConvergeConsole GUI's online help system.

For the interactive mode of the QConvergeConsole CLI, refer to the *QConvergeConsole CLI User's Guide* (for download instructions, see "Related Materials" on page xii). For the non-interactive mode of the QConvergeConsole CLI, refer to the procedures in this section to display and modify the following:

■ Adapter-Level iSCSI Parameters
■ Port-Level iSCSI Parameters
■ Summary of Target Sessions
■ Target Session-Level iSCSI Negotiated Parameters
■ Target Session-Level Persistent iSCSI Parameters

Adapter-Level iSCSI Parameters

This section shows the commands used to display and to modify adapter-level iSCSI parameters.

Displaying Adapter-Level iSCSI Parameters

To view the adapter configured settings, issue the `-ch` command. The positional parameter, `[hba_port_inst]`, is optional. If an `hba_port_inst` is specified, information for only that adapter is shown. If the `hba_port_inst` is not specified, information for all adapters in the system is listed.

Command line options:

```
-ch [hba_port_inst]
```

Example:

```
$qaucli -pr iscsi -ch
```

Or:

```
$qaucli -iscsi -ch
*** hba instance: 0
HBA_Alias                              :  QLogic QLE8262
*** hba instance: 1
HBA_Alias                              :  QLogic QLE8262
```

Modifying Adapter-Level iSCSI Parameters

Use the -nh command to set the adapter-level parameters for single- or multi-port adapters. The positional parameter becomes <hba_port_inst> and a series of one or more parameter name-value pairs. To check the list of parameters, use the -ch option.

Command line options:

```
-nh <hba_port_inst> <config_name|config_alias> <value>
[<config_name|config_alias> <value>]
```

Example:

```
$qaucli -pr iscsi -nh HBA_ALIAS "AccountingHBA"
```

Or:

```
$qaucli -iscsi  -nh HBA_ALIAS  "AccountingHBA"
HBA (adapter) Parameters:
The following table lists the parameters that may be configured using
the -nh option in non-interactive mode.
Full Parameter Name          Alias Name         Allowable Values
------------------           ----------         ----------------
HBA_Alias                    HBAALIAS           Character string
```

Port-Level iSCSI Parameters

This section shows the commands used to display and to modify port-level iSCSI parameters.

Displaying Port-Level iSCSI Parameters

Use the -c command to view port configured settings. The positional parameter, [hba_port_inst], is optional. If the hba_port_inst is specified, only information for that port is shown. If the hba_port_inst is not specified, information on all ports in the system is shown.

Example:

```
$qaucli -pr iscsi -c 0
```

Or:

```
$qaucli -iscsi -c 0
*******************************
*** Displaying Port inst=0 ***
*******************************
*** Displaying HBA (Adapter) Level Information inst=0 ***
HBA_Alias                             :  QLogic QLE8262
HBA_TCP_Max_Window_Size               :  19537
HBA_Default_Fragment_Reass_Timeout    :  0
HBA_Reserved_TCP_Config               :  0x00000000
HBA_Delayed_ACK                       :  off
*** Displaying Port General Summary Information inst=0 ***
 0. HBA: 0 Port: 1 HBA Port Instance: 0 HBA Model: QLE8262
    HBA Serial Number: (000e1e031684)qlutil_GetP3Params1:
    BoardStr=QLogic QLE8262 ; BoardId=0x26; BoardPortNum=1;
PCIFunction=5; MAC_ADDR= 0: E:1E: 3:16:85
 FW Version: 4.2.2 Type: Fibre
    IP Address: 192.168.105.208
    Alias:
    iSCSI Name: iqn.2000-04.com.qlogic:isp8214.000e1e031685.5
User Defined IP Address.
IPv4 Address             : 192.168.105.208
Gateway                  : 192.168.105.178
Subnet Mask              : 255.255.255.0

IPv6 Protocol is currently disabled.
iSNS                  : Disabled.
*** Displaying ISCSI Settings inst=0 ***
Force_Negotiate_Main_iSCSI_Keys    :  off
iSCSI_Send_Markers                 :  off(*)
iSCSI_Header_Digests               :  off
iSCSI_Data_Digests                 :  off
iSCSI_Immediate_Data               :  on
iSCSI_Initial_R2T                  :  off
iSCSI_Data_Seq_In_Order            :  on(*)
iSCSI_Data_PDU_In_Order            :  on(*)
iSCSI_CHAP_Auth                    :  off(*)
iSCSI_Bidi_CHAP_Auth               :  off(*)
iSCSI_Snack                        :  off
iSCSI_Discovery_Logout             :  on
iSCSI_Strict_Login                 :  off
iSCSI_Error_Recovery_Level         :  0(*)
iSCSI_Alias                        :
*** Displaying Firmware Settings inst=0 ***
FW_Marker                          :  on(*)
FW_Stat_Alarm                      :  off(*)
FW_Accept_AEN                      :  off(*)
FW_Access_Control                  :  off(*)
FW_Session_Mode                    :  on(*)
FW_Initiator_Mode                  :  on(*)
FW_Target_Mode                     :  off(*)
```

```
FW_Fast_Posting                    :    off(*)
FW_Sense_Buffer_Desc               :    off(*)
FW_ZIO_Enable_Mode                 :    off
AFW_Device_Timeouts                :    on
AFW_Delayed_Ack                    :    off
AFW_AutoConnect                    :    on
*** Displaying Device Settings inst=0 ***
Large_Frames                       :    off
DevType                            :    0(*)
ExeThrottle                        :    0
FirstBurstLen                      :    32
KeepAliveTO                        :    30
DefaultTime2Retain                 :    20(*)
DefaultTime2Wait                   :    2(*)
MaxBurstLen                        :    512
MaxOutstandingR2T                  :    1
MaxRxDataSegmentLen                :    128(*)
Port                               :    3260(*)
IPv4TOS                            :    0
IPv4TTL                            :    64
*** Displaying Basic Settings inst=0 ***
iSCSI_Discovery_Logout             :    on
iSCSI_Strict_Login                 :    off
TCP_DHCP                           :    off
TCP_Nagle                          :    off
iSCSI_Alias                        :
IP_Address                         :    192.168.105.208
IP_Subnet_Mask                     :    255.255.255.0
IP_Gateway                         :    192.168.105.178
Secondary_DNS                      :    (*)
Secondary_IP_Address               :    (*)
Task_Management_Timeout            :    10
ENABLE_IPV4                        :    on
ENABLE_IPV6                        :    off
LOC_LINK_AUTO                      :    off
ROUTABLE_AUTO                      :    off
LDROUTER_AUTO                      :    off
IPv6_Addr_Local_link               :    fe80::
ENABLE_4022IPV4                    :    on
*** Displaying Advanced Settings inst=0 ***
FW_Marker                          :    on(*)
FW_Stat_Alarm                      :    off(*)
FW_Accept_AEN                      :    off(*)
FW_Access_Control                  :    off(*)
FW_Session_Mode                    :    on(*)
FW_Initiator_Mode                  :    on(*)
FW_Target_Mode                     :    off(*)
FW_Fast_Posting                    :    off(*)
FW_Sense_Buffer_Desc               :    off(*)
FW_ZIO_Enable_Mode                 :    off
AFW_Device_Timeouts                :    on
AFW_Delayed_Ack                    :    off
AFW_AutoConnect                    :    on
DevType                            :    0(*)
ExeThrottle                        :    0
FirstBurstLen                      :    32
```

```
            IP_Fragmentation                  :   on(*)
            IP_ARP_Redirect                   :   off
            VLAN_Enable                       :   off
            VLAN_User_Priority                :   0
            VLAN_ID                           :   0
            IPv4_TOS_ENABLE                   :   off
            Force_Negotiate_Main_iSCSI_Keys   :   off
            iSCSI_Send_Markers                :   off(*)
            iSCSI_Header_Digests              :   off
            iSCSI_Data_Digests                :   off
            iSCSI_Immediate_Data              :   on
            iSCSI_Initial_R2T                 :   off
            iSCSI_Data_Seq_In_Order           :   on(*)
            iSCSI_Data_PDU_In_Order           :   on(*)
            iSCSI_CHAP_Auth                   :   off(*)
            iSCSI_Bidi_CHAP_Auth              :   off(*)
            iSCSI_Error_Recovery_Level        :   0(*)
            KeepAliveTO                       :   30
            DefaultTime2Retain                :   20(*)
            DefaultTime2Wait                  :   2(*)
            MaxBurstLen                       :   512
            MaxOutstandingR2T                 :   1
            MaxRxDataSegmentLen               :   128(*)
            Port                              :   3260(*)
            TCP_Timer_Scale                   :   0(*)
            TCP_Time_Stamp                    :   on
            TCP_Window_Scale                  :   0
            iSCSI_Name                        :
            iqn.2000-04.com.qlogic:isp8214.000e1e031685.5
            ZIO                               :   0
            IPv4TOS                           :   0
            IPv4TTL                           :   64
            IPV6_TCP_Timer_Scale              :   3(*)
            IPv6_TCP_Time_Stamp               :   on
            IPv6_TCP_Window_Scale             :   0
            IPv6_VLAN_ID                      :   0
            IPv6_VLAN_User_Priority           :   0
            IPv6_VLAN_Enable                  :   off
            IPv6_Traffic_Class                :   0
            IPv6_Hop_Limit                    :   64
            IPv6_ND_Reachable_Timer           :   100
            IPv6_ND_Retransmit_Timer          :   100
            IPv6_ND_Stale_Timeout             :   100
            IPv6_DAD_Count                    :   1
            IPv6_Router_Advertised_MTU        :   0(*)
            IPv4_Address_State                :   Valid(*)
            IPv6_Link_Loc_Address_State       :   Invalid(*)
            IPv6_Address0_State               :   Invalid(*)
            IPv6_Address1_State               :   Invalid(*)
            IPv6_Default_Router_State         :   No router(*)
            IPv6_MCast_Listnr_Disco_Enable    :   off
            ACB_Version                       :   2(*)
            AFW_Serlz_Task_Mngmt              :   off
            Large_Frames                      :   off
            *** Displaying IPv6 Settings inst=0 ***
            IPv6_Addr_Local_link              :   fe80::
```

```
IPv6_Addr_Routable0                     : ::
IPv6_Addr_Routable1                     : ::
Default_IPv6_Router                     : ::
IPv6_Port                               : 3260
IPv6_Gratuitious_Neighbor_Ad_Enable :   off
IPv6_Redirect_Enable                    : off
*** Displaying IPv6 TCP Settings inst=0 ***
IPv6_Nagle                              : off
IPV6_TCP_Timer_Scale                    : 3(*)
IPv6_TCP_Time_Stamp                     : on
*** Displaying Remaining parameters inst=0 ***
ACB_Supported                           : on(*)
Values noted with (*) are read only.
```

Modifying Port-Level iSCSI Parameters

Use the `-n` command to modify port-level iSCSI parameters.

Command line options:

```
-n <hba_port_inst> <config_name|config_alias> <value>
<config_name|config_alias> <value>
```

Example:

In the following example, the HBA port instance is 0, and the parameter change is to turn on iSCSI header digests.

```
$qaucli -pr iscsi -n 0 iSCSI_Header_Digests on
```

Or:

```
$qaucli -iscsi -n 0 iSCSI_Header_Digests on
Port Parameters:
The following table lists the parameters that may be configured using
the -n option in non-interactive mode.
Full Parameter Name              Alias Name  Allowable Values
------------------               ----------  ----------------
AFW_Device_Timeouts              AFWDT       on or off
AFW_Delayed_Ack                  AFDACK      on or off
AFW_AutoConnect                  AFWC        on or off
AFW_Serlz_Task_Mngmt             AFWSTM      on or off
ExeThrottle                      ET          0 to 32767
FirstBurstLen                    FB          0 to 32767
Force_Negotiate_Main_iSCSI_Keys  FNMIK       on or off
IP_ARP_Redirect                  IPARP       on or off
IPv6_MCast_Listnr_Disco_Enable   IPV6MLDEN   on or off
iSCSI_Alias                      IALS        Character string
iSCSI_Header_Digests             IHD         on or off
iSCSI_Data_Digests               IDD         on or off
iSCSI_Immediate_Data             IID         on or off
iSCSI_Initial_R2T                IIR2T       on or off
iSCSI_Snack                      ISNACK      on or off
```

```
iSCSI_Discovery_Logout                 ID           on or off
iSCSI_Strict_Login                     IS           on or off
KeepAliveTO                            KATO         0 to ?
Large_Frames                           LRGFRM       on or off
(not for 4010s)
MaxBurstLen                            MBL          0 to ?
MaxOutstandingR2T                      MOR2T        0 to ?
TCP_DHCP                               TCPDHCP      on or off
TCP_Nagle                              TCPN         on or off
TCP_Time_Stamp                         TCPTMS       on or off
TCP_Window_Scale                       WINSCALE     0 to 14
VLAN_Enable                            VLAN         on or off
VLAN_User_Priority                     VLANUPRIOR   0 to 7
VLAN_ID                                VLANID       0 to 4095
IP_Address                             IPAD         IPv4 address format
IP_Subnet_Mask                         IPSM         IPv4 address format
IP_Gateway                             IPGW         IPv4 address format
ZIO                                    ZIO          2 to 16
FW_ZIO_Enable_Mode                     ZIOE         on or off
Task_Management_Timeout                TMTO         0 to 65535
ENABLE_IPV4                            EIPV4        on or off
ENABLE_4022IPV4                        E4022IPV4    on or off
ENABLE_IPV6                            EIPV6        on or off
LOC_LINK_AUTO                          LOCLA        on or off
ROUTABLE_AUTO                          RAUTO        on or off
LDROUTER_AUTO                          LDRA         on or off
IPv6_Addr_Local_link                   IPLL         IPv6 address format
IPv6_Addr_Routable0                    IPR0         IPv6 address format
IPv6_Addr_Routable1                    IPR1         IPv6 address format
Default_IPv6_Router                    IPRR         IPv6 address format
IPv4TOS                                IPV4TOS      0 255
IPv4_TOS_ENABLE                        TOS_ENABLE   on or off
IPv4TTL                                IPV4TTL      0 255
IPv6_Port                              IPV6PORT     0 65535
IPv6_Gratuitious_Neighbor_Ad_Enable IPV6GNAE      on or off
IPv6_Redirect_Enable                   IPV6RDE      on or off
IPv6_Nagle                             TCPV6ND      on or off
IPV6_TCP_Timer_Scale                   TCPV6TS      0 to 7
IPv6_TCP_Time_Stamp                    TCPV6TST     on or off
IPv6_TCP_Window_Scale                  IPV6TCPWS    0 to 14
IPv6_VLAN_ID                           IPV6VLANID   0 to 4095
IPv6_VLAN_User_Priority                IPV6VLANUP   0 to 7
IPv6_VLAN_Enable                       IPV6VLANEN   on or off
IPv6_Traffic_Class                     IPV6TC       0 to 255
IPv6_Hop_Limit                         IPV6HL       0 to 255
(router may override)
IPv6_ND_Retransmit_Timer               IPV6NDRET    0 to 4294967295
(router may override)
IPv6_ND_Stale_Timeout                  IPV6STO      0 to 4294967295
```

```
                                    (router may override)
IPv6_ND_Reachable_Timer             IPV6NDRT    0 to 4294967295
(router may override)
IPv6_DAD_Count                      IPV6DAD     0 to 255
```

Summary of Target Sessions

Use the `-ts` command to display summary information for both persistent and non-persistent targets. Both [`hba_port_inst`] and [`target_id`] are optional parameters. If neither of the parameters is present, the information is displayed for all adapters and all targets. When `hba_port_inst` is entered, target information for all targets on the specified adapter is displayed. If the optional `target_id` keyword is entered, only information on the specified target is displayed.

Command line options:

`-ts [hba_port_inst] [target_id]`

Example:

```
$qaucli -pr iscsi -ts
```

Or:

```
$qaucli -iscsi -ts
Target ID: 2  hba_no: 0   IP: 192.168.105.247    Port: 3260 TGT
Instance #: 2
ISCSI Name:
Alias:
State: No Connection
Target ID: 3  hba_no: 0   IP: 192.168.105.247    Port: 3260 TGT
Instance #: 3
ISCSI Name:
iqn.2003-05.com.stringbeansoftware:apptester-starblazer248-target
Alias:
State: Session Active
Target ID: 2  hba_no: 1   IP: 192.168.105.247    Port: 3260 TGT
Instance #: 2
ISCSI Name:
iqn.2003-05.com.stringbeansoftware:apptester-appstorm245-target
Alias:
State: Session Active
Target ID: 3  hba_no: 1   IP: 192.168.105.247    Port: 3260 TGT
Instance #: 3
ISCSI Name:
iqn.2003-05.com.stringbeansoftware:apptester-starblazer248-target
Alias:
State: Session Active
```

Target Session-Level iSCSI Negotiated Parameters

Use the `-t` command to display information for targets. The positional parameter is *`<hba_port_inst>`*. The optional parameter is [*`target_id`*]. If only the *`hba_port_inst`* is entered, target information for all targets on the specified adapter is displayed. If the optional *`target_id`* is entered, only information on the specified target is displayed.

Command line options:

`-t <hba_port_inst> [target_id]`

Example:

> In the following examples, the HBA port instance is 0, and the target ID is 3.

```
$qaucli -pr iscsi -t 0 3
```

> Or:

```
$qaucli -iscsi -t  0 3
Target ID: 3  hba_no: 0   IP: 192.168.105.247    Port: 3260 TGT
Instance #: 3
ISCSI Name:
iqn.2003-05.com.stringbeansoftware:apptester-starblazer248-target
Alias:
State: Session Active
TGT_iSCSI_Name                         :
iqn.2003-05.com.stringbeansoftware:apptester-starblazer248-target
TGT_Target_ID                          :  3(*)
TGTO_Active                            :  off(*)
TGTO_Access_Granted                    :  off(*)
TGTO_Target_Entry                      :  on(*)
TGTO_Initiator_Entry                   :  off(*)
TGT_RetryCount                         :  0(*)
TGT_RetryDelay                         :  0(*)
TGT_DevType                            :  0(*)
TGT_ExeThrottle                        :  0
TGT_FirstBurstLen                      :  32
TGTIPO_Fragmentation                   :  on(*)
TGTISCSIO_Force_Neg_Main_Keys          :  off
TGTISCSIO_Send_Markers                 :  off(*)
TGTISCSIO_Header_Digests               :  off
TGTISCSIO_Data_Digests                 :  off
TGTISCSIO_Immediate_Data               :  on
TGTISCSIO_Initial_R2T                  :  off
TGTISCSIO_Data_Sequence_In_Order       :  on(*)
TGTISCSIO_Data_PDU_In_Order            :  on(*)
TGTISCSIO_CHAP_Authentication          :  off
TGTISCSIO_Bidi_CHAP_Authentication     :  off
TGTISCSIO_Snack                        :  off
```

```
TGTISCSIO_Discovery_Logout        :  on
TGTISCSIO_Strict_Login            :  off
TGTISCSIO_Error_Recovery_Level    :  0(*)
TGT_KeepAliveTimeout              :  30
TGT_DefaultTimeout                :  2
TGT_DefaultTime2Retain            :  20(*)
TGT_MaxBurstLen                   :  512
TGT_MaxOutstandingR2T             :  1
TGT_MaxRxDataSegmentLen           :  128(*)
TGT_MaxTxDataSegmentLen           :  0(*)
TGT_Port                          :  3260
TGTTCPO_Nagle                     :  off
TGTTCPO_Timer_Scale               :  0(*)
TGTTCPO_Timestamp                 :  on
TGT_TaskManagementTimeout         :  10
TGT_ExeCount                      :  0(*)
TGT_TargetPortalGroupID           :  1(*)
TGT_InitiatorSessID               :  0x000e1e031685
TGT_TargetSessID                  :  9(*)
TGT_TargetIPAddress               :  192.168.105.247
TGT_Window_Scale_Enable           :  on
TGT_Rx_Window_Scale               :  0
TGT_Tx_Window_Scale               :  0(*)
TGT_TimeStamp_Enable              :  64(*)
TGT_DDB_IPv6                      :  off
TGT_IPv6_Address                  :  c0a8:69f7::15:0:0
TGT_IPv6_iSCSIName                :
iqn.2003-05.com.stringbeansoftware:apptester-starblazer248-target
TGT_IPv6_Port                     :  3260
TGT_DIF_Enable                    :  off
TGT_Max_Segment_Size              :  1448
TGT_Local_TCP_Port                :  29912(*)
TGT_Type_of_Service               :  0
TGT_Traffic_Class                 :  0(*)
TGT_Local_IPv6_Address            :  c0a8:69d0::(*)
TGT_Perm_Redirect_Option          :  off(*)
TGT_Temp_Redirect_Option          :  off(*)
TGT_Redirect_IPAddr               :  88.2.60.0(*)
TGT_Redirect_IPAddr_State         :  Not Redirected(*)
TGT_IPv6_Flow_Label               :
TGT_4022_Deleyed_ACK              :  off
TGT_IPv6_Source_Addr_Flg          :  0
TGT_IPv6_Source_Addr              :  c0a8:69d0::(*)
Values noted with (*) are read only.
```

Target Session-Level Persistent iSCSI Parameters

This section shows the commands used to display and to modify target
session-level persistent iSCSI parameters.

Displaying Target Session-Level Persistent iSCSI Parameters

Use the `-tp` command to view target persistent parameter information
(pre-negotiation, from Flash memory). The positional parameter is
<hba_port_inst>. The optional parameter is [*target_id*]. If only the
hba_port_inst is entered, target information for all targets on the specified
adapter is shown. If the optional *target_id* is entered, only information on the
specified target is shown.

Command line options:

```
-tp <hba_port_inst> [target_id]
```

Example:

> In the following examples, the HBA port instance is 0, and the target ID is 3.

```
$qaucli -pr iscsi -tp 0 3
```

> Or:

```
$qaucli -iscsi  -TP 0 3
Target ID: 3  hba_no: 0   IP: 192.168.105.247    Port: 3260 TGT
Instance #: 3
ISCSI Name:
iqn.2003-05.com.stringbeansoftware:apptester-starblazer248-target
Alias:
State: Session Active
TGT_iSCSI_Name                          :
iqn.2003-05.com.stringbeansoftware:apptester-starblazer248-target
TGT_Target_ID                           :  3(*)
TGTO_Active                             :  off(*)
TGTO_Access_Granted                     :  off(*)
TGTO_Target_Entry                       :  on(*)
TGTO_Initiator_Entry                    :  off(*)
TGT_RetryCount                          :  0(*)
TGT_RetryDelay                          :  0(*)
TGT_DevType                             :  0(*)
TGT_ExeThrottle                         :  0
TGT_FirstBurstLen                       :  32
TGTIPO_Fragmentation                    :  on(*)
TGTISCSIO_Force_Neg_Main_Keys           :  off
TGTISCSIO_Send_Markers                  :  off(*)
TGTISCSIO_Header_Digests                :  off
TGTISCSIO_Data_Digests                  :  off
TGTISCSIO_Immediate_Data                :  on
TGTISCSIO_Initial_R2T                   :  off
TGTISCSIO_Data_Sequence_In_Order        :  on(*)
TGTISCSIO_Data_PDU_In_Order             :  on(*)
TGTISCSIO_CHAP_Authentication           :  off
TGTISCSIO_Bidi_CHAP_Authentication      :  off
```

```
TGTISCSIO_Snack                       : off
TGTISCSIO_Discovery_Logout            : on
TGTISCSIO_Strict_Login                : off
TGTISCSIO_Error_Recovery_Level        : 0(*)
TGT_KeepAliveTimeout                  : 30
TGT_DefaultTimeout                    : 2
TGT_DefaultTime2Retain                : 20(*)
TGT_MaxBurstLen                       : 512
TGT_MaxOutstandingR2T                 : 1
TGT_MaxRxDataSegmentLen               : 128(*)
TGT_MaxTxDataSegmentLen               : 0(*)
TGT_Port                              : 3260
TGTTCPO_Nagle                         : off
TGTTCPO_Timer_Scale                   : 0(*)
TGTTCPO_Timestamp                     : on
TGT_TaskManagementTimeout             : 10
TGT_ExeCount                          : 0(*)
TGT_TargetPortalGroupID               : 1(*)
TGT_InitiatorSessID                   : 0x000e1e031685
TGT_TargetSessID                      : 9(*)
TGT_TargetIPAddress                   : 192.168.105.247
TGT_Window_Scale_Enable               : on
TGT_Rx_Window_Scale                   : 0
TGT_Tx_Window_Scale                   : 0(*)
TGT_TimeStamp_Enable                  : 64(*)
TGT_DDB_IPv6                          : off
TGT_IPv6_Address                      : c0a8:69f7::15:0:0
TGT_IPv6_iSCSIName                    :
iqn.2003-05.com.stringbeansoftware:apptester-starblazer248-target
TGT_IPv6_Port                         : 3260
TGT_DIF_Enable                        : off
TGT_Max_Segment_Size                  : 1448
TGT_Local_TCP_Port                    : 29912(*)
TGT_Type_of_Service                   : 0
TGT_Traffic_Class                     : 0(*)
TGT_Local_IPv6_Address                : c0a8:69d0::(*)
TGT_Perm_Redirect_Option              : off(*)
TGT_Temp_Redirect_Option              : off(*)
TGT_Redirect_IPAddr                   : 40.2.45.1(*)
TGT_Redirect_IPAddr_State             : Not Redirected(*)
TGT_IPv6_Flow_Label                   :
TGT_4022_Deleyed_ACK                  : off
TGT_IPv6_Source_Addr_Flg              : 0
TGT_IPv6_Source_Addr                  : c0a8:69d0::(*)
Values noted with (*) are read only.
```

Modifying Target Session-Level iSCSI Parameters

Use the `-tc` command to modify target-session-level iSCSI parameters. The positional parameters are *<hba_port_inst>*, *<target_id>*, and a series of one or more parameter name-value pairs.

Command line options:

`-tc <hba_port_inst> <target_id> <config_name|config_alias> <value> <config_name|config_alias> <value>]`

Example:

> In the following examples, the HBA port instance is 0, the target ID is 3, and the parameter change is to set the keep alive time-out value to 15 seconds.

> `$qaucli -pr iscsi -tc 0 3 TGT_KeepAliveTimeout 15`

Or:

```
$qaucli -iscsi -tc 0 3 TGT_KeepAliveTimeout 15
Target Parameters:
The following table lists the parameters that may be configured using
the -tc option in non-interactive mode.
Full Parameter Name                 Alias Name      Allowable Values
-------------------                 ----------      ----------------
TGT_iSCSI_Name                      TGTINAME        Character string
TGT_ExeThrottle                     TGTET           0 to 32767
TGT_FirstBurstLen                   TGTFB           0 to 32767
TGTISCSIO_Header_Digests            TGTIHD          on or off
TGTISCSIO_Data_Digests              TGTIDD          on or off
TGTISCSIO_Immediate_Data            TGTIID          on or off
TGTISCSIO_Initial_R2T               TGTIIR2T        on or off
TGTISCSIO_Snack                     TGTISNACK       on or off
TGTISCSIO_Discovery_Logout          TGTLDS          on or off
TGTISCSIO_Strict_Login              TGTIS           on or off
TGT_KeepAliveTimeout                TGTKATO         0 to 32767
TGT_DefaultTimeout                  TGTDTO          0 to 32767
TGT_MaxBurstLen                     TGTMB           0 to 32767
TGT_MaxOutstandingR2T               TGTMOR2T        0 to 32767
TGT_Port                            TGTPORT         0 to 65535
TGTTCPO_Nagle                       TGTTCPN         on or off
TGTTCPO_Timestamp                   TGTTMS          on or off
TGT_TaskManagementTimeout           TGTTMT          0 to 65535
TGT_InitiatorSessID                 TGTISID         0x0 to 0xffffffffffff
TGT_TargetIPAddress                 TGTIPADD        IPv4 address format
TGT_Window_Scale_Enable             TGTWINSCALEEN   on or off
TGT_Rx_Window_Scale                 TGTRXWINSCALE   0 to 14
TGT_IPv6_Address                    TGT_DDB_IPv6    IPv6 address format
TGT_IPv6_iSCSIName                  TGTINAME_IPv6   Character string
TGT_IPv6_Port                       TGTPORT_IPv6    0 to 32767
TGT_DIF_Enable                      TGTDIFEN_IPv6   on or off
TGT_Max_Segment_Size                TGTMSS          0 to 65535
TGT_IPv6_Source_Addr_Flg            TGTSRCADDR_IPv6 0 to 3 (0=Don't Care, 1=Link Local,
                                                    2=Address 0, 3=Address 1)
```

Configuring iSCSI Initiators with QConvergeConsole

This section provides procedures on how to configure the following iSCSI initiators using QLogic's QConvergeConsole utility:

■ Configuring the Windows iSCSI Initiator
■ Configuring the Linux iSCSI Initiator
■ Configuring the ESX iSCSI Initiator

> **NOTE**
>
> For information on installing and starting the QConvergeConsole GUI, refer to the *QConvergeConsole GUI Installation Guide* (for download instructions, see "Related Materials" on page xii). All procedural information for the QConvergeConsole GUI is covered in the QConvergeConsole GUI's online help system.

Configuring the Windows iSCSI Initiator

Use the QConvergeConsole CLI to configure the iSCSI initiator for Windows.

To configure a Windows iSCSI initiator:

1. Access the QConvergeConsole CLI either by double-clicking the QConvergeConsole CLI desktop icon or by entering `qaucli` in the CMD window.

2. On the QConvergeConsole CLI Main Menu, select **2, Adapter Configuration**.

3. On the Adapter Type Configuration Selection menu, select **1, CNA Configuration**.

4. On the Converged Network Adapter (CNA) Protocol Type Selection menu, select **CNA iSCSI Configuration** (either **1** or **2**, depending on how many drivers are loaded).

5. On the Converged Network Adapter (CNA) iSCSI Configuration menu, select **3, Port IP Settings**.

6. Select the Converged Network Port you want to configure.

7. Select **2, Configure IP Settings**.

8. Complete the interactive list of settings as follows:

    a. **Enable IPv4 [on]:** Press ENTER to accept the default.

    b. **DHCP to obtain IPv4 Network Information: [off]:** Press ENTER to accept the default.

    c. **IP_Address [0.0.0.0]:** Type the IP address of the initiator system, and then press ENTER.

       d.   **IP_Subnet_Mask [0.0.0.0]:** Type the appropriate subnet mask, and then press ENTER.

       e.   **IP_Gateway [0.0.0.0]:** Press ENTER to accept the default.

       f.   **Enable IPv6 [off]:** Press ENTER to accept the default.

9.   On the options menu that appears, select **3, Save changes and reset HBA (if necessary)**.

10.   At the prompt for both ports, type `Yes`.

11.   To return to the Converged Network Adapter (CNA) iSCSI Configuration menu, type **P** and press ENTER, and then type **P** and press ENTER again.

12.   On the Converged Network Adapter (CNA) iSCSI Configuration menu, select **4, Target Configuration**.

13.   Select the same Converged Network Port you selected in Step 6.

14.   Select **6, Add a Target**.

15.   Complete the interactive list of settings as follows:

       a.   **IPv6 Target? [off]:** Press ENTER to accept the default.

       b.   **TGT_iSCSI_Name [ ]:** Type the iSCSI Qualified Name (IQN) of the iSCSI target to connect to and then press ENTER.

       c.   **TGT_Port [3260]:** Press ENTER to accept the default.

       d.   **TGT_TargetIPAddress [0.0.0.0]:** Type the IP address of the target and then press ENTER.

16.   On the options menu that appears, select **12, Save Target/CHAP Changes**.

The iSCSI initiator is now configured to connect to the iSCSI target.

Configuring the Linux iSCSI Initiator

Use the QConvergeConsole CLI to configure the iSCSI initiator for Linux.

To configure a Linux iSCSI initiator:

1.   Access the QConvergeConsole CLI by typing `qaucli` in a terminal window.

2.   On the QConvergeConsole CLI Main Menu, select **2, Adapter Configuration**.

3.   On the Adapter Type Configuration Selection menu, select **1, CNA Configuration**.

4.   On the Converged Network Adapter (CNA) Protocol Type Selection menu, select **1, CNA iSCSI Configuration**.

5.   On the Converged Network Adapter (CNA) iSCSI Configuration menu, select **3, Port IP Settings**.

6. Select the Converged Network Port you want to configure.

7. Select **2, Configure IP Settings**.

8. Complete the interactive list of settings as follows:

   a. **Enable IPv4 [on]:** Press ENTER to accept the default.

   b. **DHCP to obtain IPv4 Network Information: [off]:** Press ENTER to accept the default.

   c. **IP_Address [ ]:** Type the IP address of the initiator system and then press ENTER.

   d. **IP_Subnet_Mask [255.255.255.0]:** Type the appropriate subnet mask and then press ENTER.

   e. **IP_Gateway [0.0.0.0]:** Press ENTER to accept the default.

   f. **Enable IPv6 [off]:** Press ENTER to accept the default.

9. On the options menu that appears, select **3, Save changes and reset HBA (if necessary)**.

10. At the prompt for both ports, type **Yes**.

11. To return to the Converged Network Adapter (CNA) iSCSI Configuration menu, type **P** and press ENTER and then type **P** and press ENTER again.

12. On the Converged Network Adapter (CNA) iSCSI Configuration menu, select **4, Target Configuration**.

13. Select the same Converged Network Port you selected in Step 6.

14. Select **6, Add a Target**.

15. Complete the interactive list of settings as follows:

   a. **IPv6 Target? [off]:** Press ENTER to accept the default.

   b. **TGT_iSCSI_Name [ ]:** Type the iSCSI Qualified Name (IQN) of the iSCSI target to connect to and then press ENTER.

   c. **TGT_Port [3260]:** Press ENTER to accept the default.

   d. **TGT_TargetIPAddress [0.0.0.0]:** Type the IP address of the target and then press ENTER.

16. On the options menu that appears, select **12, Save Target/CHAP Changes**.

The iSCSI initiator is now configured to connect to the iSCSI target.

Configuring the ESX iSCSI Initiator

The software iSCSI initiator must be enabled for ESX/ESXi to be able to use it for accessing iSCSI storage.

To configure an ESX/ESXi initiator:

1. Log in to the vSphere Client.

2. In the inventory panel, select a server to which to connect.

3. Click the **Configuration** tab.

4. In the Hardware panel, click **Storage Adapters**.

5. From the list of available storage adapters, select the iSCSI initiator you want to configure and then click **Properties**.

6. Click **Configure**.

   The **General Properties** dialog box shows the initiator's status, default name, and alias.

7. To enable the initiator, click **Enabled**.

8. (Optional) To change the default iSCSI name for your initiator, type a new name. The name you enter must be worldwide unique and properly formatted so that all storage devices can recognize the software iSCSI initiator.

9. To save your changes, click **OK**.

> **NOTE**
>
> If you change the iSCSI name, it is used for new iSCSI sessions. Existing sessions do not use new settings until you log out and log in again.

Enabling CHAP Authentication with QConvergeConsole

You can enable CHAP authentication with either the interactive mode or non-interactive mode of QConvergeConsole CLI. For details on the interactive mode, refer to the *QConvergeConsole CLI User's Guide*. For the non-interactive mode of the QConvergeConsole CLI, the following sections describe how to enable CHAP:

- Configuring CHAP with QConvergeConsole CLI
- Linking to a CHAP Target

> **NOTE**
>
> For information on installing and starting the QConvergeConsole GUI, refer to the *QConvergeConsole GUI Installation Guide* (for download instructions, see "Related Materials" on page xii). All procedural information for the QConvergeConsole GUI is covered in the QConvergeConsole GUI's online help system.

Configuring CHAP with QConvergeConsole CLI

To configure CHAP with QConvergeConsole CLI:

1. To add a primary and local CHAP entry (name and secret), issue the `-addchap` command to add a CHAP entry to the persistent CHAP table. The positional parameters are `<hba_port_inst>`, `<CHAP name>`, and `<CHAP secret>`. The optional parameter is `[-BIDI]` indicating the CHAP entry is a bidirectional entry (default is local CHAP). The adapter is reset after this command is issued.

   > **NOTE**
   >
   > The *iSCSI RFC Specification* recommends a minimum CHAP secret length of 12 bytes or characters. The maximum CHAP secret length for QLogic iSCSI cards (the firmware limit) is 100 bytes or characters.

   Command line options:

   ```
   [-BIDI] -addchap <hba_port_inst> <CHAP name> <CHAP secret>
   ```

   In the following examples, the HBA port instance is 0, the CHAP name is `chapdbserver1`, and the CHAP secret is `k9Q038iaZwlqPplq012`.

   ```
   $qaucli -pr iscsi -addchap 0 chapdbserver1 k9Q038iaZwlqPplq012
   ```

   Or:

   ```
   $qaucli -iscsi -addchap 0 chapdbserver1 k9Q038iaZwlqPplq012
   ```

2. To add a peer and BIDI CHAP entry (name and secret), issue the `-addchap` command to add a CHAP entry to the persistent CHAP table. The positional parameters are `<hba_port_inst>`, `<CHAP name>`, and `<CHAP secret>`. The optional parameter is `[-BIDI]` indicating the CHAP entry is a BIDI entry (default is local CHAP). The adapter is reset after this command is issued.

   Command line options:

   ```
   [-BIDI] -addchap <hba_port_inst> <CHAP name> <CHAP secret>
   ```

   Example:

   In the following examples, the HBA port instance is 2, the CHAP name is `chapbidistorage1`, and the CHAP secret is `Z9aujqklaZwlqPplq0827`.

   ```
   $qaucli -pr iscsi -BIDI -addchap 2 chapbidistorage1
   Z9aujqklaZwlqPplq0827
   ```

   Or:

   ```
   $qaucli -iscsi -BIDI -addchap 2 chapbidistorage1
   Z9aujqklaZwlqPplq0827
   ```

3. To view the CHAP map table to determine the CHAP index to use later to link the CHAP entry to a target, issue the `-dspchap` command. The positional parameter for this command is `<hba_port_inst>`.

Command line options:

```
-dspchap <hba_port_inst>
```

In the following examples, the HBA port instance = 0.

```
$qaucli -pr iscsi -dspchap 0
```

Or:

```
$qaucli -iscsi -dspchap 0
CHAP TABLE
Entry: 1
Name: chapdbserver1
Secret: k9Q038iaZwlqPplq012
```

4. Add a persistent Send Target to discover, and log in with dynamic entries to discovered targets by issuing the `-pa` command. The `-pa` command adds a persistent target. The positional parameters are `<hba_port_inst>` and `<ip address>`. The optional parameters are `[-PORT port_num]` and `[-INAME name]`. If the optional port number is not specified, it defaults to 3260. If the optional `INAME` (iSCSI name) is not specified, it defaults to an empty string.

Command line options:

```
-pa <hba_port_inst> <ip address> [-PORT port_num] [-INAME name]
```

Example:

In the following examples, the HBA port instance is 0, and the Send Target IP is 10.14.64.154.

```
$qaucli -pr iscsi -pa 0 10.14.64.154
```

Or:

```
$qaucli -iscsi -pa 0 10.14.64.154
```

5. To display a persistent Send Target entry, issue the `-ps` command (you should initially expect a failed connection because the target is not yet linked to CHAP). The `-ps` command lists persistent (bound) targets. The positional parameter is `[hba_port_inst [target_id]]`. If no `target_id` is specified, all targets for the specified `hba_port_inst` are shown. If neither the `hba_port_inst` nor `target_id` are specified, all `target_id`s for all adapters in the system are shown.

Command line options:

```
-ps <hba_port_inst> [target_id]
```

In the following examples, the HBA port instance is 0, and the Send Target IP is 10.14.64.154.

```
$qaucli -pr iscsi  -ps 0
```

Or:

```
$qaucli -iscsi  -ps 0
Target ID: 2  hba_no: 0  IP: 10.14.64.154  Port: 3260  TGT
Instance #: 2
ISCSI Name:
Alias:
State: Session Failed
```

6. Link the CHAP entry to the target by issuing the `-linkchap` command. The positional parameters are *<hba_port_inst>*, *<chap_no>* and *<target_id>*. The adapter is not reset after this command is issued.

    Command line options:

    ```
    -linkchap <hba_port_inst> <chap_no> <target_id> [<TGTBCA>
    <value>]
    ```

    In the following examples, the HBA port instance is 0, the CHAP number is 1, and the Target ID is 2.

    ```
    $qaucli -pr iscsi -linkchap 0 1 2
    ```

    Or:

    ```
    $qaucli -iscsi -linkchap 0 1 2
    ```

Linking to a CHAP Target

You can link CHAP to a target with active bidirectional (BIDI) CHAP authentication. `TGTBCA` is an optional parameter to turn BIDI CHAP authentication on or off for this target connection.

To link a CHAP target:

1. Link CHAP to a target with BIDI enabled by issuing the `-linkchap` command.

    Command line options:

    ```
    iscli -linkchap <hba_port_inst> <chap_no> <target_id>
    [<TGTBCA> <on|off>]
    ```

    In the following example, the HBA port instance is 2, the CHAP number is 9, and the Target ID is 10.

    ```
    $qaucli -pr iscsi -linkchap 2 9 10 TGTBCA on
    ```

2. View persistent targets by issuing the `-ps` command. You should see only the Send Target.

In the following examples, the HBA port instance is 0.

```
$qaucli -pr iscsi -ps 0
```

Or:

```
$qaucli -iscsi -ps 0
Target ID: 2 hba_no: 0 IP: 10.14.64.154 Port: 3260 TGT
Instance #: 2
ISCSI Name:
Alias:
State: No Connection
```

3. To view all targets linked to the CHAP, issue the `-chapmap` command. This command lists the mapping of targets to CHAP table entries. The positional parameter for this command is *<hba_port_inst>*.

Command line options:

```
-chapmap <hba_port_inst>
```

In the following example, the HBA port instance is 0.

```
$qaucli -pr iscsi -chapmap 0
```

Or:

```
$qaucli -iscsi -chapmap 0
Targets configured for CHAP:
Target ID: 2 IP: 10.14.64.154 Port: 3260
ISCSI Name:
Alias:
Name: chapdbserver1
Secret: k9Q038iaZwlqPplq012
Target ID: 64 IP: 10.14.64.154 Port: 3260
ISCSI Name: iqn.1987-05.com.cisco:00.ba6d7ea87bba.chap1
Alias: chap1
Name: chapdbserver1
Secret: k9Q038iaZwlqPplq012
Target ID: 65 IP: 10.14.64.154 Port: 3260
ISCSI Name: iqn.1987-05.com.cisco:00.00c80ea3857f.chap2
Alias: chap2
Name: chapdbserver1
Secret: k9Q038iaZwlqPplq012
Target ID: 66 IP: 10.14.64.154 Port: 3260
ISCSI Name: iqn.1987-05.com.cisco:00.0b597ef8adf8.chap3
Alias: chap3
Name: chapdbserver1
Secret: k9Q038iaZwlqPplq012
Target ID: 67 IP: 10.14.64.154 Port: 3260
ISCSI Name: iqn.1987-05.com.cisco:00.28182218624e.chap4
Alias: chap4
Name: chapdbserver1
Secret: k9Q038iaZwlqPplq012
```

Windows Management Applications

Windows management applications for the adapter include the following:

- Windows NIC Driver Management Applications
- Windows Teaming
- Windows VLAN Configuration
- User Diagnostics for Windows NIC Driver Management Applications

Windows NIC Driver Management Applications

- Overview
- Viewing and Changing Adapter Properties

Overview

In the QConvergeConsole CLI (qaucli) utility, you can view VLAN and teaming overview information by issuing the `qaucli -nt -zvt` command. (The qaucli is an installation option available when you install the Windows drivers; see "Windows Driver Installation and Configuration" on page 5.)

Viewing and Changing Adapter Properties

This section provides information on using the QConvergeConsole CLI:

- Viewing Adapter Properties
- Changing Adapter Properties

Viewing Adapter Properties

Issue the following commands to view the adapter properties.

To list all detected adapter ports:

```
qaucli -nic -i [cna_port_inst]
```

To view adapter information:

```
qaucli -nic -icna [cna_port_inst]
```

To view port DCBX protocol information:

```
qaucli -nic -idcbx [cna_port_inst]
```

To view configured port settings:

```
qaucli -nic -iset [cna_port_inst]
```

To view physical link status:

```
qaucli -nic -link [cna_port_inst]
```

To view port information:

```
qaucli -nic -pinfo [cna_port_inst]
```

Changing Adapter Properties

> **NOTE**
>
> For an adapter that is teamed or an adapter with VLANs, do not directly modify the adapter properties. To ensure that the properties of all teamed adapters and adapters with VLANs remain synchronized with the team properties, make property changes *only* on the Team Management page (see "Modifying a Team" on page 86).

To configure the adapter port, issue the following command:

```
qaucli -nic -n [cna_port_inst] <config_name|config_alias>
<value> [<config_name|config_alias> <value>]
```

You can set the following properties:

- Port_Alias
- Port_Physical_MAC_Alias
- Port_LAA_MAC_Alias

You can change the variables listed in Table 3-1; however, some variables cannot be changed on specific OSs or configuration states. To determine which ones can be changed, use the -c keyword.

Table 3-1. Port Adapter Variables and Values

Variable	Values
Checksum_Offload_Enable	on, off
IPv4_Checksum_Offload_Enable	off, Rx, Tx, RxTx
IPv4_TCP_Checksum_Offload_Enable	off, Rx, Tx, RxTx
IPv6_TCP_Checksum_Offload_Enable	off, Rx, Tx, RxTx
IPv4_UDP_Checksum_Offload_Enable	off, Rx, Tx, RxTx
IPv6_UDP_Checksum_Offload_Enable	off, Rx, Tx, RxTx
Large_Send_Offload_Enable	on, off
IPv4_Large_Send_Offload_v1_Enable	on, off
IPv4_Large_Send_Offload_v2_Enable	on, off
IPv6_Large_Send_Offload_v2_Enable	on, off
Receive_Side_Scaling_Enable	on, off

Table 3-1. Port Adapter Variables and Values (Continued)

Variable	Values
Header_Data_Split_Enable	on, off
Jumbo_Frames_MTU_9000_Enable	on, off
Jumbo_Frames_MTU_9000_Enable_Rx	on, off
Jumbo_Frames_MTU_9000_Enable_Tx	on, off
LOCAL_Administered_Address_MAC	xx:xx:xx:xx:xx:xx
Port_Wake_On_LAN_Option	0=Disabled, 1=Wake on Magic Frame
VLAN_ID	1.4094

To set the adapter configuration alias, issue the following command:

```
qaucli -nic -nh [cna_port_inst] <config_name|config_alias>
<value> [<config_name|config_alias> <value>]
```

Windows Teaming

- Overview
- Teaming Modes
- Using the CLI for Teaming
- Using the Team Management GUI
- Teaming Configuration
- Viewing Teaming Statistics

Overview

You can group together multiple network adapters in a server to make a team. Individual adapters that are part of a team operate as a team rather than standalone adapters. A team provides traffic load balancing across the member adapters and fault tolerance when some, but not all, of the members lose connectivity.

To enable teaming functionality, install the teaming driver in addition to the basic NIC.

Team MAC Address

At initialization, the teaming driver selects the team's MAC address to be the MAC of one of the teamed adapters. In general, the first adapter to come up is chosen to be the *preferred primary adapter*. The preferred primary's MAC address is assigned to the MAC address of the team. Alternately, you can choose any valid MAC address as the team's static MAC address, also called the LAA. Make sure any provided LAA is unique for the local Ethernet network. This provision gives the system administrator more flexibility in configuring the MAC address for a team when necessary.

Teaming Modes

Teaming is designed to improve reliability and fault tolerance of networks and to enhance performance by efficient load balancing.

The following NIC teaming modes are provided:

- Failsafe Mode ensures that an alternate standby or redundant adapter becomes active if the primary network connection fails.

- Switch-Independent Load Balancing Mode ensures distribution of transmit loads across the teamed adapters.

- Link Aggregation Mode (802.3ad static, 802.3ad dynamic [active and passive LACP) enables the use of multiple adapters together as a single, virtual adapter with the aggregated capacity of its individual adapters.

All team types—failsafe, switch-independent load balancing, and link aggregation—can be heterogeneous as well as homogeneous. Every team must have at least one QLogic adapter.

Table 3-2 shows that Failsafe and Tx load-balancing modes are switch independent, which means they do not require switch configuration. LACP or 802.3ad requires switch ports configured for LACP.

Table 3-2. Windows Teaming Modes

Mode	Failover Capability	Switch Dependency	SFT (System Fault Tolerance)	Load Balancing	Quantity of Ports per Team (Range [a])
Failsafe	Yes: Layer 2	No	Yes	No	1–16
Tx load balancing	Yes	No	Yes	Yes: Layers 3 or 4	1–16
Static 802.3ad	Yes	Yes	Yes	Yes	1–16
Dynamic 802.3ad	Yes	Yes	Yes	Yes	1–16

[a] 16×16 ports can be aggregated per system: 16 ports per team and 16 teams per system.

Failsafe Mode

The failsafe mode provides Layer 2 fault tolerance. The failsafe mode provides high reliability through redundancy in the event of port failure. When the primary network connection is down, data traffic is automatically transferred to a secondary, standby connection. The preferred primary adapter can be specified either by the system administrator or by the teaming driver (if the administrator does not select the preferred adapter). When the teaming driver needs to make the selection, it selects the best adapter in terms of bandwidth, health, and capability. The preferred primary must always be a QLogic adapter.

The administrator can also choose one of the following failback types to specify the behavior when connection to preferred primary is restored after a period of failure:

■ *None*

   When the preferred primary becomes operational again, the driver does not automatically switch back the primary to the active adapter.

■ *Preferred Primary*

   When the preferred primary becomes operational again, the driver automatically switches back the primary as the active adapter. The network traffic resumes to the primary adapter from the standby adapter. The traffic stays with the secondary adapter only as long as the primary adapter is down.

■ *Auto Select*

   Use this option to enable the teaming driver to automatically select the best adapter based on parameters such as bandwidth, link state, health.

In failsafe mode, the standby adapter could be dissimilar in the individual features supported and capacity and might come from a different vendor.

All the adapters in the team share a common team MAC address. This is a locally administered MAC address or can be a default MAC address specified by the driver. Only one adapter at a time in the team is active for network traffic. No two same MAC addresses are exposed to the switch at the same time.

Failsafe mode is inherent in all other teaming modes and is switch agnostic.

Switch-Independent Load Balancing Mode

Switch-independent load balancing mode provides a failsafe feature and supports transmit load balancing. For receive load balancing, use the 802.3ad modes.

In this mode, the outbound traffic is efficiently distributed across the member adapters to increase the transmit bandwidth. Traffic load balancing is connection-based to avoid out-of-order packet delivery. The administrator can select one of the following load distribution types:

■ *Auto Select* indicates that the load is distributed based on the target IP address (IPv4 or IPv6) and port number. This option ensures a one-to-one correspondence between a traffic flow and a team adapter.

■ *MAC address based* indicates that the load is distributed based on the target MAC address.

In switch-independent load balancing, a team receives the traffic on the preferred primary adapter. If the preferred primary adapter fails, the receive load switches to a secondary adapter (failover operation). If the preferred primary adapter becomes operational again, the receive load fails back to the preferred primary adapter (failback operation). Thus, a switch-independent load balancing team also behaves like a failsafe team. Each time the preferred primary changes due to failover or failback, other network elements are notified of the change in the primary adapter through team gratuitous address resolution protocols (ARPs).

Link Aggregation Mode

Link aggregation provides increased bandwidth and high reliability by combining several NICs into a single, logical, network interface called a LAG. The link aggregation is scalable, meaning an adapter can be added or deleted either statically or dynamically from a team.

Traffic from all the team ports that form a LAG have the same MAC address, which is the MAC address of the team. If a new adapter joins the LAG, or an adapter forming the LAG fails, the LAG becomes operational again after a brief exchange of protocols between the switch and the server. QLogic adapters are rapidly aggregated, with a latency of 1 to 2 seconds.

Two options are available in the link aggregation mode:

- Static Link Aggregation
- Dynamic Link Aggregation

> **NOTE**
>
> The switch must support the IEEE 802.3ad standard for the preceding two link aggregation modes to work.

Static Link Aggregation

Static link aggregation (SLA, 802.3ad static protocols with generic trunking) is a switch-assisted teaming mode, where the switch must be 802.3ad compliant. The switch ports must be configured so that the switch perceives adapters from a LAG as a single, virtual adapter.

In SLA, the ports on the switch are active by default. There is no negotiation between the switch and the intermediate driver to decide on adapters participating in a LAG.

In SLA mode, the protocol stack responds to ARP requests with a single, advertised MAC address and an IP address corresponding to the LAG. Each physical adapter in the team uses the same team MAC address during transmission. As the switch (at the other end of link) is aware of the trunking teaming mode, it appropriately modifies the forwarding table to indicate the trunk as a single virtual port. This modification ensures correct traffic routing on the receive side as well. In this mode, the switch also distributes receive traffic across the member adapters.

Dynamic Link Aggregation

Dynamic link aggregation (DLA) with LACP is similar to SLA except that LACP allows self configuration of LAG through handshaking between the switch and the intermediate driver. For the team to function, LACP must be enabled at both ends of the link: the server and the switch. LACP (802.3ad dynamic) allows switch ports to dynamically communicate with the intermediate driver, allowing controlled addition and removal of ports from the team.

Link aggregation mode has transmit load balancing and fail safety support. If a link connected through a participant port of a link-aggregated team goes down, LACP provides failover and load balancing across the remaining members of the team. In addition, if a new member port is added to the team or is removed from the team, the switch performs load re-balancing for the receive operation and the driver performs load balancing for the transmit operation, to accommodate the change in configuration.

Transmit load distribution in LACP provides the following options:

- *None* indicates no traffic distribution. Only a single "active" adapter is used for transmit. The driver selects the active adapter based on LACP state information.

- *Auto Select* indicates that the load is distributed based on the target IP address and port number. This option ensures a one-to-one correspondence between a traffic flow and a team adapter.

- *MAC address based* indicates that the load is distributed based on the target MAC address.

Using the CLI for Teaming

You can view, create, configure, and delete teams using QConvergeConsole.

To view a list of teams, issue the following command:

```
qaucli -nic -teamlist
```

To view team information, issue the following command:

```
qaucli -nic -teaminfo <team_inst|ALL>
```

To preview available ports before configuring a new team, issue this command:

```
qaucli -nic -teamnew_portspreview
```

To configure a new team, issue the following command:

```
qaucli -nic -teamnew <team_type> <port_insts|ALL>
```

where *port_insts* are the ports indices separated by commas (for example, 1,2) and *team_type* is either 1=Fail Over or 2=Load Balanced.

To delete a team, issue the following command:

```
qaucli -nic -teamdel <team_inst|ALL>
```

> **NOTE**
>
> The following applies to configuring teaming and VLAN using the QConvergeConsole CLI:
>
> **Windows Server 2012 and later:**
> QConvergeConsole CLI does not support teaming and VLAN configuration. Use the native Windows teaming interface instead of QConvergeConsole CLI.

Using the Team Management GUI

Use the Team Management property page to manage the following teaming-related activities:

- Viewing network topology
- Creating, modifying, and deleting teams
- Viewing and changing team properties
- Adding and deleting virtual adapters

To launch the Team Management property page:

1. In Windows, access the **Computer Management** dialog box and then click **Device Manager** in the left pane.

2. Under **Network adapters**, right-click the **QLogic 10 Gigabit Ethernet adapter** and then select **Properties**.

3. Click the **Team Management** tab to bring that page to the front (see Figure 3-1) and perform teaming-related management.
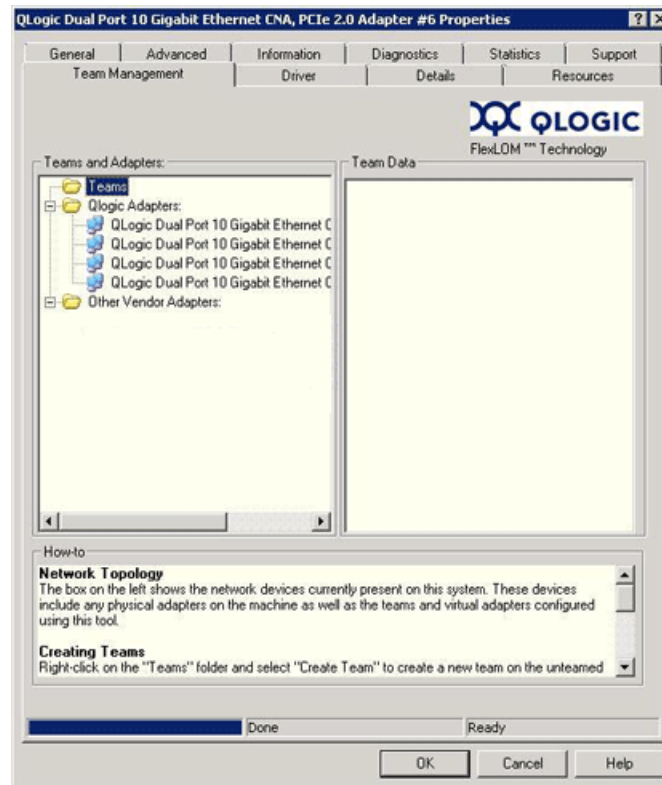
Figure 3-1. Team Management Property Page

On the Team Management page, the Teams and Adapters pane on the left lists the network devices currently present on this system, including:

■ Teams and virtual adapters, as well as their member physical adapters
■ QLogic and other vendor adapters.

Procedures for creating a team, adding virtual adapters, and more are provided in the How-to box on the bottom of the Team Management page.

Teaming Configuration

Teaming configuration includes creating, modifying, and deleting teams, and viewing team statistics on the Team Management property page. To launch the Team Management property page, see "Using the Team Management GUI" on page 79.

Information on teaming configuration includes the following:

■ Creating a Team
■ Modifying a Team
■ Deleting a Team
■ Saving and Restoring Teaming Configuration

Creating a Team

To create a team use the following procedure:

1. Right-click the **Teams** folder icons and then click **Create Team** (see Figure 3-2).
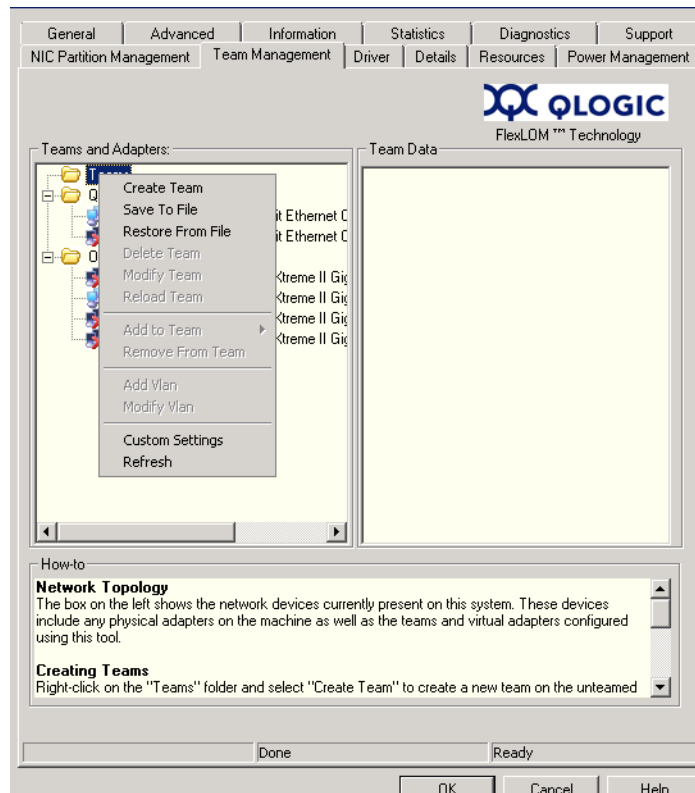
Figure 3-2. Creating a Team

2. The software automatically picks a unique team name, or you can enter your own team name. Team names must be unique on a system.

3. On the **Create Team** dialog box, specify the following (see the message pane at the bottom of the dialog box for more details) and then click **OK** to return to the adapter properties:

   ❑ **Name**—Type a name for the new team.

❑ **Type**—Select the teaming mode by clicking either **Failsafe Team**, **802.3ad Static Team**, **802.3ad Dynamic Team**, or **Switch Independent Load Balancing**. If you select the 802.3ad dynamic option, you must also select one of the following options:

  ■ **Active LACP**: LACP is a Layer 2 protocol that is used control the teaming of physical ports into an aggregated set. LACP discovers if a host's ports are connected to a switch that supports aggregation on the connected ports and configures those ports into an aggregation bundle. For LACP to operate, one side has to be Active LACP. The active LACP side of the protocol initiates the protocol.

  ■ **Passive LACP**: The passive LACP side just responds to the active LACP requests.

❑ **Adapters to Add**—Select the check box next to each adapter that should form the team.

❑ **Use default MAC Address**—Select this check box to have the driver assign a MAC address, or clear the check box to select a locally administered MAC address from the list.

❑ **Failback Delay**—Type the failback delay in seconds.

❑ **Select Preferred Primary Adapter**—Choose a preferred primary adapter for the team from the list of teamed adapters or **None** to allow the driver to assign the preferred primary adapter.

❑ **Failback Type**—If this is a Failsafe Team, select a failback type of either **None**, **Auto Select**, or **Preferred primary**.

❑ **Load Balancing Type**—If this is an 802.3ad Static Team or 802.3ad Dynamic Team, select the type of load balancing: **Auto**, **MAC Address Based**, or **None**.

❑ **Distribution Type**—If this is a Switch Independent Load Balancing team type, select a distribution type of either **Auto Select** or **MAC Address Based**.

❑ **Advanced**—Click this button to configure QLogic-specific team capabilities such as RSS, MTU, or various offloads. These properties are used to configure the member adapters to avoid any conflict after a team has been created. Figure 3-7 shows the **Advanced Team Properties** dialog box.

The following figures show the configuration of various teaming modes.
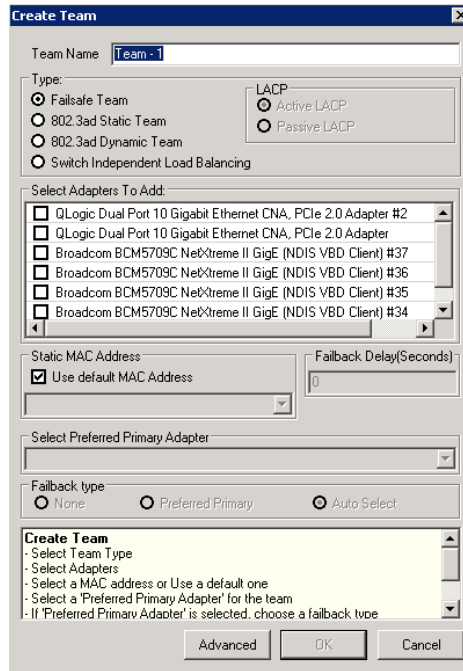
Figure 3-3. Creating a Failsafe Team
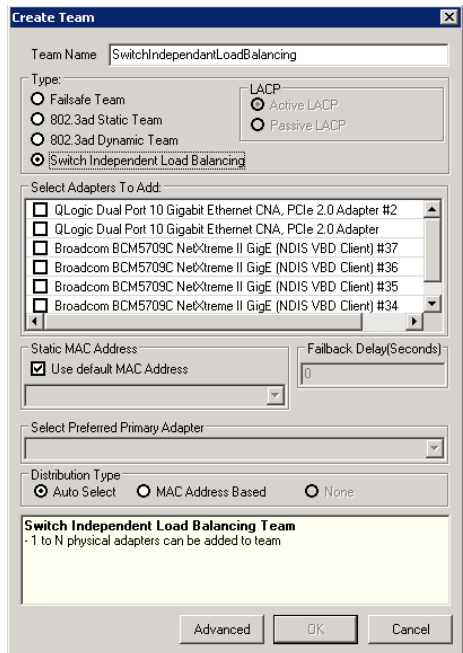
Figure 3-4. Creating a Switch-Independent Load Balancing Team
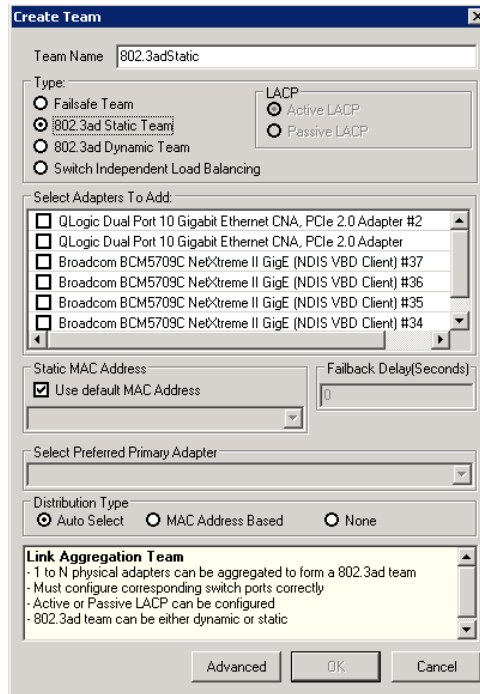
Figure 3-5. Creating an 802.3ad Static Team
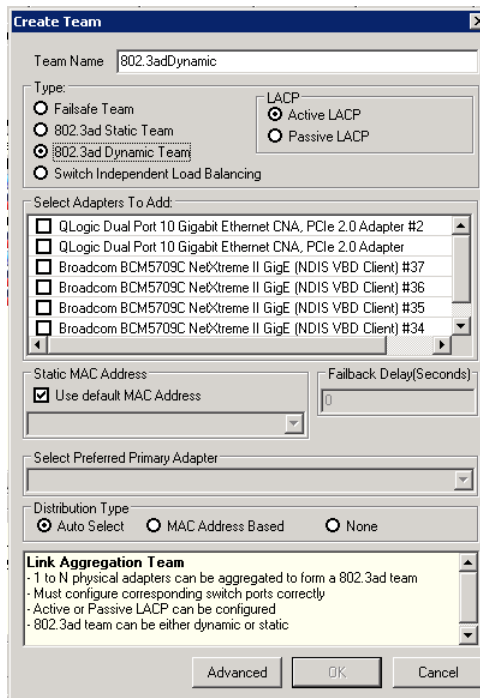
Figure 3-6. Creating an 802.3ad Dynamic Team
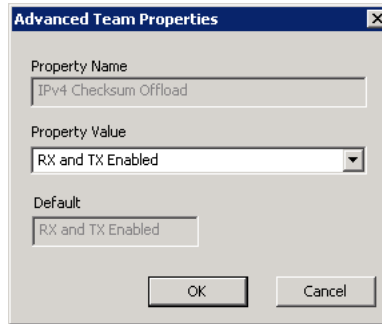
Figure 3-7. Setting Advanced Team Properties

4.   To confirm if a team has been successfully created, view the **Team and Adapters** pane on the **Team Management** page.

Figure 3-8 shows an example of a newly formed team. The **Team Data** pane on the right shows the properties, information, and status of the team or adapter that is currently selected in the **Teams and Adapters** pane on the left.
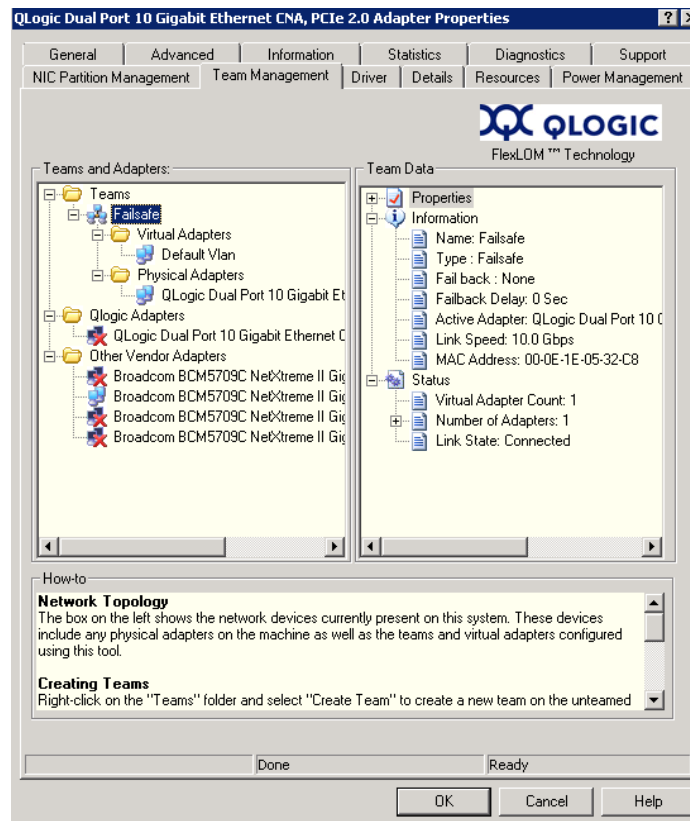
Figure 3-8. Confirming New Team Creation

Modifying a Team

A team can be modified by doing the following:

- Adding or removing one or more team members to a team
- Modifying the team properties

To add team members:

1. On the **Team Management** property page, right-click the unteamed adapter to add to a team.

2. On the shortcut menu, point to **Add to Team** and then click the team to which you want to add the adapter (see Figure 3-9).

> **NOTE**
>
> You cannot add an adapter to a team that is already a member of another team. Teaming of teams (nested teams) is not supported.
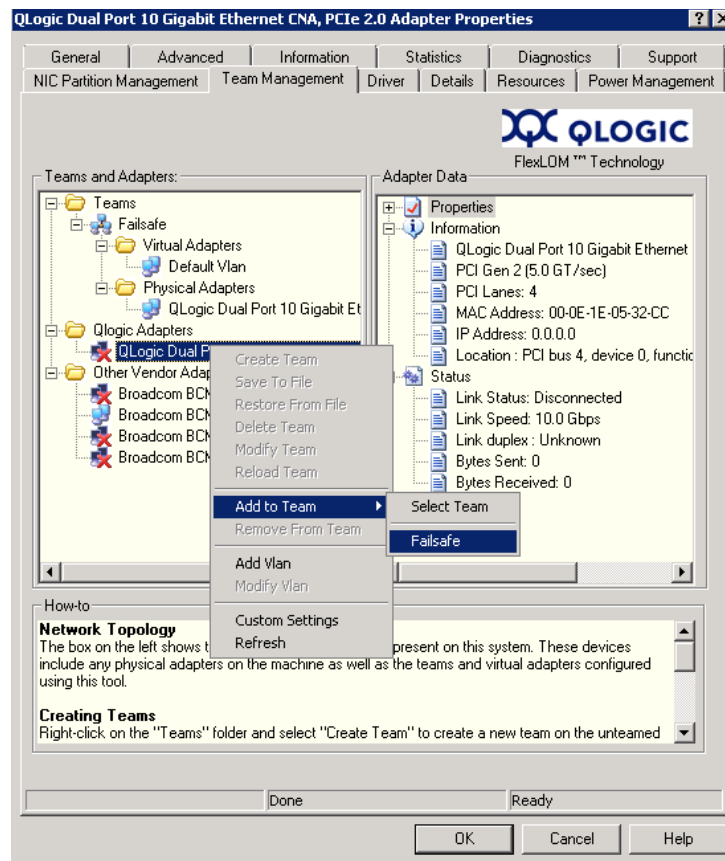
Figure 3-9. Adding a Team

To remove an adapter from a team:

> **NOTE**
>
> A team must include at least one QLogic adapter. A QLogic adapter is allowed to be deleted from a team only if it is not the last QLogic teamed adapter.

1. On the **Team Management** property page, right-click the adapter to be removed from the team.

2. On the shortcut menu, click **Remove from Team**.

At least two adapters must be present in a team. If an adapter is no longer required to be a member of a team, it can be removed from the team.

To change a team property:

For the VLAN and teaming solution to work correctly, the properties of all teamed adapters and adapters with multiple VLANs must remain synchronized with the team properties. Ensure that you change the properties of a team and an adapter with VLANs *only* on the **Team Management** page.

1. On the **Team Management** page, in the right pane under **Team Data**, expand the **Properties** list.

2. Double-click the team property you need to change.

3. In the **Advanced Team Properties** dialog box (see Figure 3-10), specify a new **Property Value** and then click **OK**.
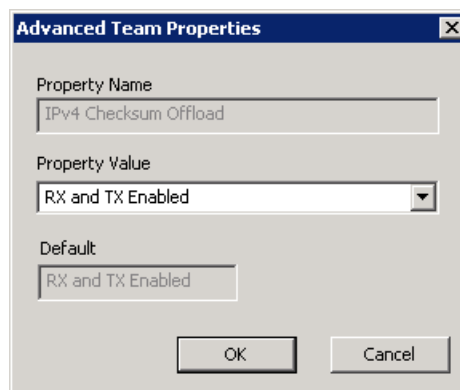
Figure 3-10. Modifying Advanced Team Properties

The team properties change takes effect immediately. Changing team properties causes the driver to reload, which might result in a momentary loss of connectivity.

> **NOTE**
>
> To ensure that the properties of all teamed adapters and adapters with VLANs remain synchronized with the team properties, do not directly modify the adapter properties on the Advanced page. If an adapter property becomes unsynchronized with its team properties, change either the team or adapter property so that they are the same on each and then reload the team. To reload a team: On the **Team Management** page, in the left pane under **Teams and Adapters**, right-click the team name and then click **Reload Team**.

To modify team composition:

1. On the **Team Management** page, in the left pane under **Teams and Adapters**, right-click the team name whose properties are to be changed.

2. On the shortcut menu, click **Modify Team** (see Figure 3-11).
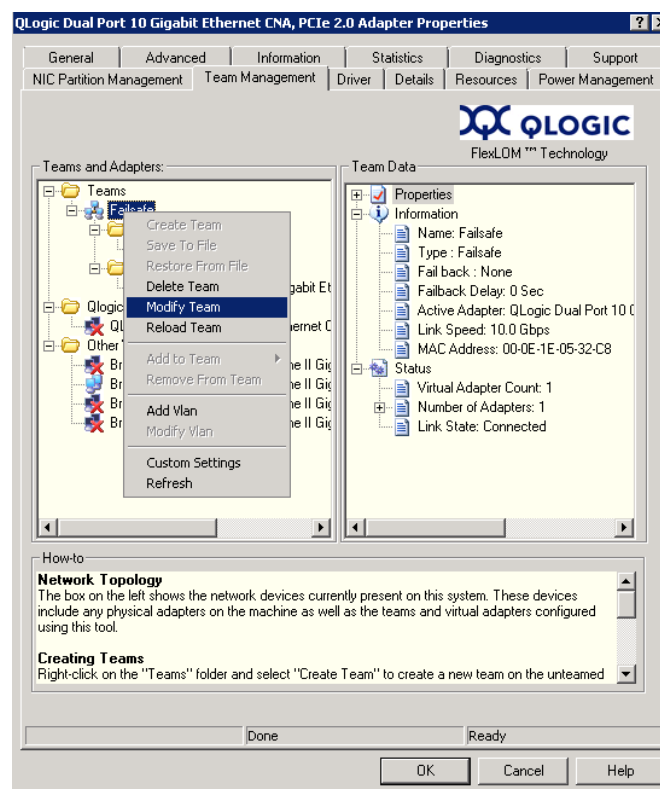
Figure 3-11. Modifying Team Properties

3. On the **Modify Team Properties** dialog box, change the team parameters as needed and then click **OK**.

Example 1: For a failsafe team, you can change the team name, assigned team static MAC address, preferred primary adapter, and failback type, as shown in Figure 3-12.
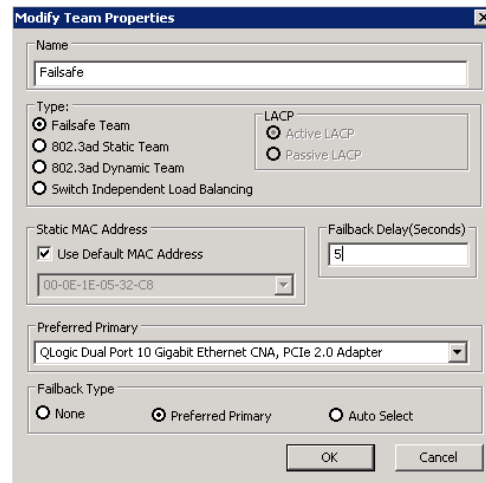
Figure 3-12. Modifying Failsafe Team Properties

Example 2: You can change the team type and the corresponding team attributes. For example, you can change from failsafe to switch-independent load balancing or from 802.3ad static team to 802.3ad dynamic team. Figure 3-13 shows a failsafe team modification, which shows the new team type and default values for team attributes. You can manually change attribute values.
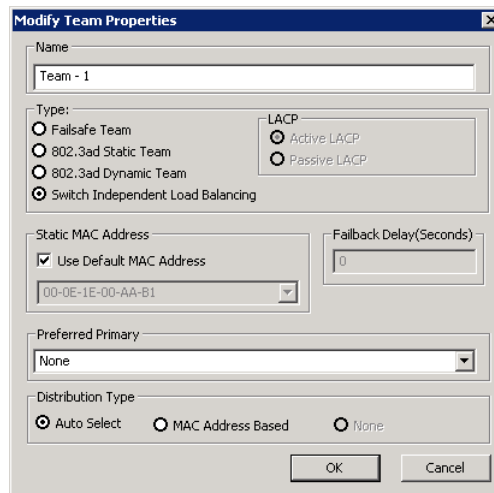
Figure 3-13. Modifying the Team Type

Deleting a Team

To delete a team:

1. On the **Team Management** property page, in the left pane under **Teams and Adapters**, right-click the team name to be deleted.

2. On the shortcut menu, click **Delete team**.

Saving and Restoring Teaming Configuration

It is recommended that you periodically save the configuration to prevent any accidental loss of network topology and settings. Current configuration including the teams, VLANs, and properties can be saved to a file. Restoring an earlier configuration results in the destruction of the current configuration.

To save a configuration:

1. On the **Team Management** page under **Teams and Adapters**, right-click the **Teams** folder.

2. On the shortcut menu, click **Save to File**.

3. Enter a location to save the configuration.

To restore a configuration:

1. On the **Team Management** page under **Teams and Adapters**, right-click the **Teams** folder.

2. On the shortcut menu, click **Restore From File**.

3. Select a previously saved configuration file.

Viewing Teaming Statistics

You can view teaming and Ethernet statistics using the QConvergeConsole utility.

To view teaming statistics in the CLI, issue the following command:

```
qaucli -nic -statport
```

To reset the Ethernet statistics counter, issue the following command:

```
qaucli -nic -sreset [cna_port_inst]
```

To display Ethernet port statistics, issue the following command:

```
qaucli -nic -statport [cna_port_inst]
```

To undo the reset of Ethernet statistics counters, issue the following command:

```
qaucli -nic -sunreset [cna_port_inst]
```

Windows VLAN Configuration

The term *VLAN* refers to a collection of devices that communicate as if they were on the same physical LAN. VLAN information covered in this section includes the following:

- VLAN Properties
- Using the CLI for VLANs
- Using the GUI for VLANs

VLAN Properties

The VLAN protocol permits insertion of a tag into an Ethernet frame to identify the VLAN to which a frame belongs. If present, the four-byte VLAN tag is inserted into the Ethernet frame between the source MAC address and the length and type field. The first two bytes of the VLAN tag consist of the 802.1q tag type; the second two bytes include a user priority field and the VLAN identifier (VID).

For QLogic adapters, VLAN tagging is done according to the IEEE 802.1q protocol and the process is the same whether it is a single port or a team.

Multiple VLANs can be configured over a single port or a team of ports up to a maximum of 64 VLANs. Each VLAN is represented by a virtual network interface that is bound to the protocols.

Using the CLI for VLANs

Use the QConvergeConsole CLI utility to view, add, and delete VLANs.

To display a list of VLANs, issue the following command:

```
qaucli -nic -vlanlist
```

To view VLAN information, issue the following command:

```
qaucli -nic -vlaninfo [vlan_inst|ALL]
```

To preview a VLAN before adding it to a port or team, issue the following command:

```
qaucli -nic -vlanadd_preview
```

To add a VLAN to a port or team, issue the following command:

```
qaucli -nic -vlanadd <list_insts> <vlan_id>
```

where *list_insts* are comma-separated port indices (for example, 1,2) and *vlan_id* is the comma-separated numeric value (for example, 1...4095 or 100,555). Use the -vlanadd_preview command to preview list indices.

To preview a VLAN before removing it from a port or team, issue the following command to list the indices to use in the `-vlandel` command:

```
qaucli -nic -vlandel_preview
```

To remove a VLAN from a port or team, issue the following command:

```
qaucli -nic -vlandel <list_insts|ALL> <vlan_id|ALL>
```

where *list_insts* specifies the comma-separated port indices (for example, 1,2) and *vlan_id* specifies a comma-separated numeric value (for example, 1...4095 or 100,555). Use the `-vlandel_preview` command to preview list indices.

To view the content of VLAN and teaming state configuration, issue the following command:

```
qaucli -nic -vtcfgview [state_cfg_file]
```

To save the VLAN and teaming state configuration, issue the following command:

```
qaucli -nic -vtsave [state_cfg_file]
```

To restore VLAN and teaming state configuration, issue the following command:

```
qaucli -nic -vtrestore [state_cfg_file]
```

> **NOTE**
>
> In the preceding commands, if the *state_cfg_file* option is not specified, the default file is `vtstate.cfg` in the installation directory.

Using the GUI for VLANs

You can use the Team Management GUI to perform the following VLAN activities:

- Adding and Configuring a VLAN
- Deleting a VLAN
- Viewing VLAN Statistics

Adding and Configuring a VLAN

Use the Team Management page to add VLAN on a team, either directly on an unteamed adapter or on an unteamed adapter while configuring parameters for the adapter's NIC driver. Whenever a team is created, a virtual adapter is implicitly created that becomes the default VLAN for the team. There can be only one untagged VLAN on a team.

Multiple virtual adapters can be created on top of the teamed adapter by creating different instances with different VLAN IDs assigned to each.

To add and configure a VLAN:

1. On the **Team Management** page under **Teams and Adapters**, right-click either a team or an unteamed adapter.

2. On the shortcut menu, click **Add VLAN** (see Figure 3-14).
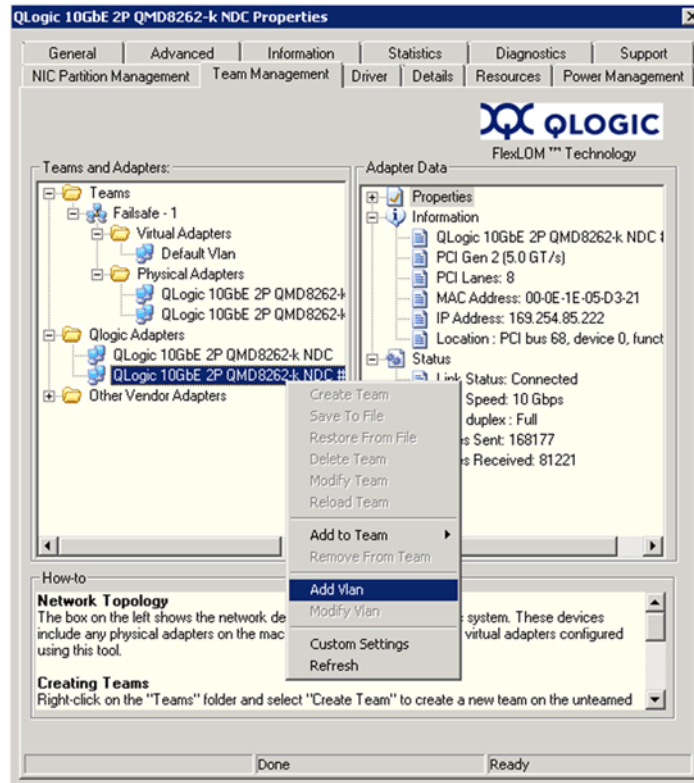
Figure 3-14. Adding a VLAN

3. On the **Configure VLAN** dialog box (see Figure 3-15), type values in the **VLAN Name** and **VLAN ID** boxes, click an appropriate **VLAN Type**, and then click **OK**.
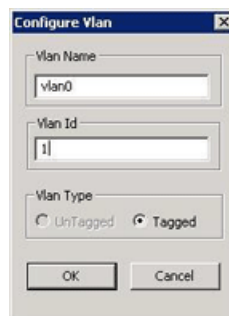
Figure 3-15. Configuring a VLAN

When the VLAN addition is complete, the added VLAN is visible as a **Virtual Adapter** on the **Team Management** page under **Teams and Adapters**.

4.   Click the added virtual adapter to view all the properties, information, and status of the virtual adapter in the **VLAN Data** pane (see Figure 3-16).
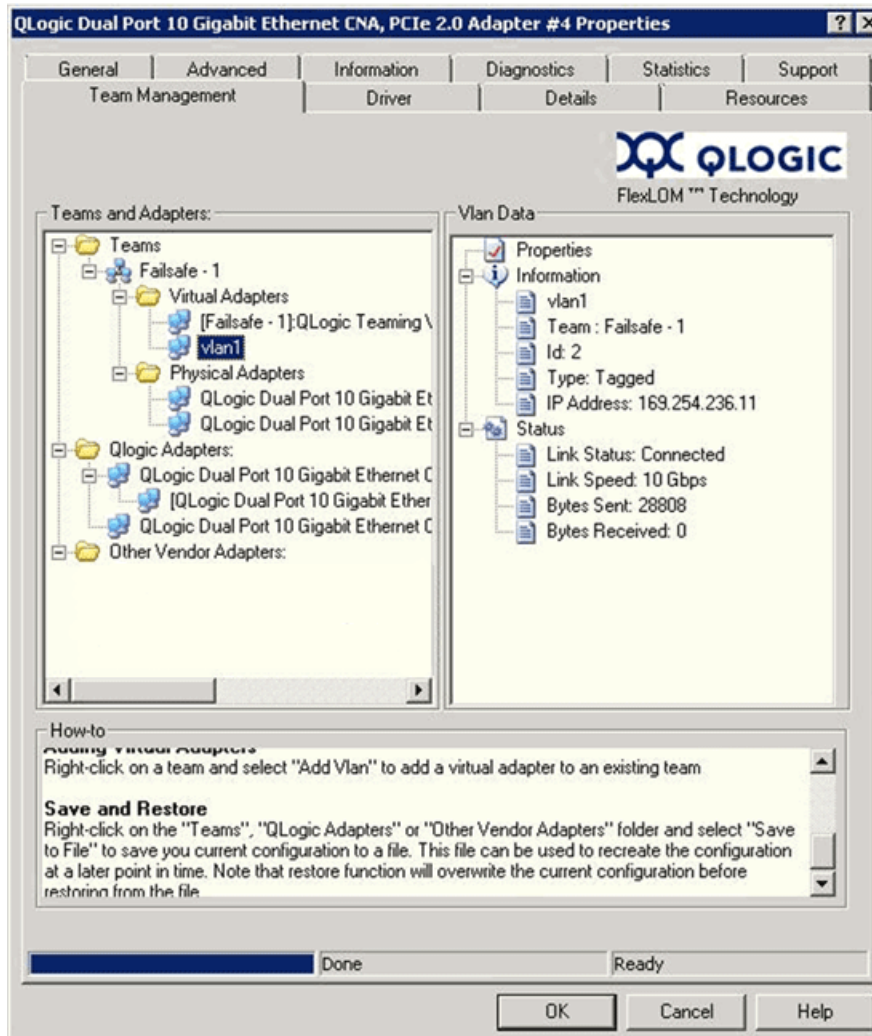
Figure 3-16. Viewing VLAN Data Properties

Deleting a VLAN

If VLAN is not needed on a team, you can delete it.

To delete a VLAN:

1.   On the **Team Management** page, right-click the VLAN that you want to remove.

2.   On the shortcut menu, click **Remove from Team**.

> **NOTE**
>
> To allow VLAN deletion, there must be at least one VLAN on the team. Deleting the last VLAN on the team results in deletion of the entire team.

Viewing VLAN Statistics

Follow these steps to view statistics for a selected VLAN.

To view VLAN statistics:

1. On the **Team Management** page, click a team name in the left pane under the **Teams** folder.

2. Check the statistical data listed in the **Team Data** pane on the right, including the VLAN type, VLAN ID, team IP, link status, link speed, and exchanged bytes.

User Diagnostics for Windows NIC Driver Management Applications

This section covers the following information for user diagnostics for Windows NIC driver management applications:

- Running Windows User Diagnostics
- Windows Diagnostic Test Descriptions
- Windows Diagnostic Test Messages

Running Windows User Diagnostics

You can run user diagnostics using either the GUI or the CLI.

> **NOTE**
>
> For information on installing and starting the QConvergeConsole GUI, refer to the *QConvergeConsole GUI Installation Guide* (for download instructions, see "Related Materials" on page xii). All procedural information for the QConvergeConsole GUI is covered in the QConvergeConsole GUI's online help system.

To run user diagnostics in the GUI:

1. Access the Windows Control Panel and then open the Device Manager.

2. In the Device Manager, right-click the QLogic 10Gb Ethernet adapter, and then, on the shortcut menu, click **Properties**.

3. On the adapter properties page, click the **Diagnostics** tab.

Figure 3-17 shows the Diagnostics page.
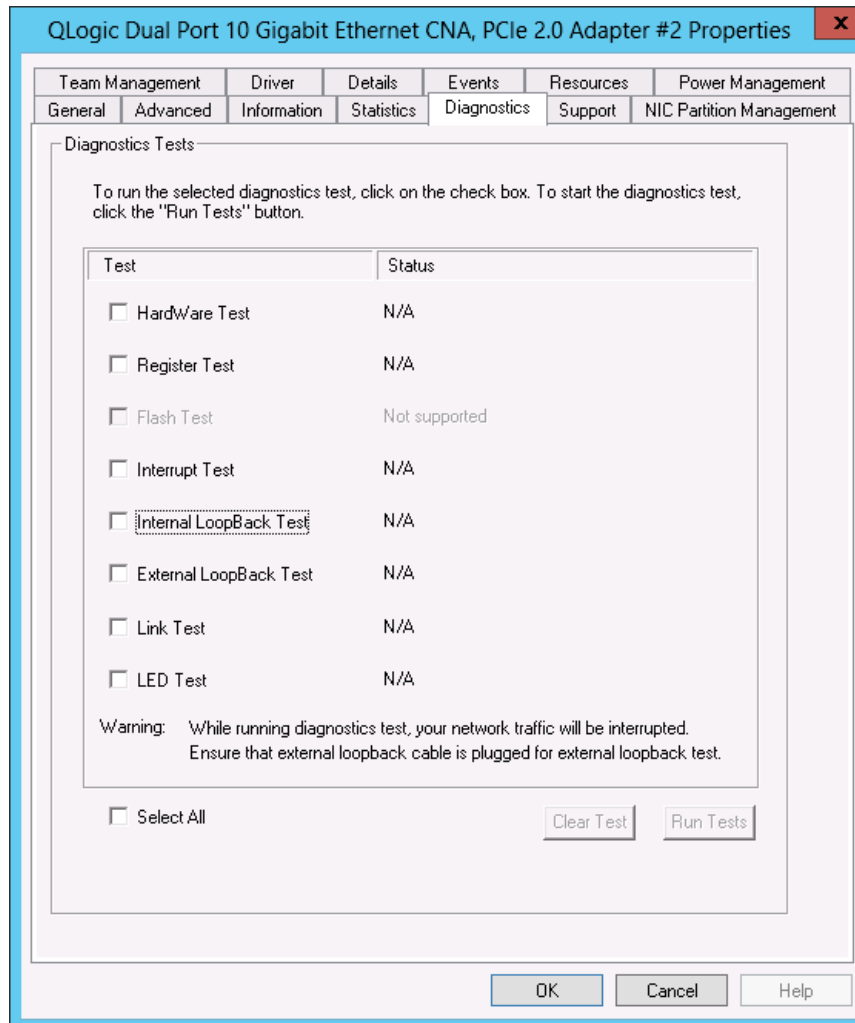
Figure 3-17. Diagnostics Tests on Windows

4. Under **Diagnostic Tests**, select one or more check boxes indicating the tests you want to run: **Hardware Test**, **Register Test**, **Interrupt Test**, **Internal Loopback Test**, **External Loopback Test**, and **Link Test**. ("Windows Diagnostic Test Descriptions" on page 100 describes each test type.)

5. Click **Run Tests**.

> **NOTE**
>
> Only one test can run at a time. Multiple tests can run sequentially.

To run user diagnostics in the CLI:

Use QConvergeConsole CLI (qaucli), a unified command line utility, to manage all QLogic adapter models, including running user diagnostics.

The overall option (`-pr <protocol>`) allows you to start the utility with a specific protocol type, either NIC, iSCSI, or Fibre Channel. If you do not specify a protocol, all protocols are enabled by default. Table 3-3 and Table 3-4 list the QConvergeConsole commands for selecting a protocol.

Table 3-3. Selecting a Protocol in Menu Mode

Command	Description
`qaucli`	Start QConvergeConsole CLI in interactive mode
`qaucli -pr nic [options]`	Issue NIC command line options
`qaucli -pr iscsi [options]`	Issue iSCSI command line options
`qaucli -pr fc [options]`	Issue Fibre Channel and FCoE command line options
`qaucli -npar [options]`	Issue NIC partitioning (NPAR) (Switch Independent Partitioning) command line options

Table 3-4. Selecting a Protocol in Legacy Mode

Command	Description
`qaucli -nic [options]`	Use NIC legacy command line
`netscli [options]`	Use NIC legacy command line
`qaucli -iscsi [options]`	Use iSCSI legacy command line
`iscli [options]`	Use iSCSI legacy command line
`qaucli -fc [options]`	Use Fibre Channel legacy command line
`scli [options]`	Use Fibre Channel legacy command line

Diagnostic help commands, and command options available for each specific protocol, are available by specifying `-h` to the protocol, as shown in Table 3-5.

Table 3-5. Getting Help

Command	Description
`-h`	Print usage of a specific adapter type and then exit
`qaucli -pr nic -h`	Print NIC protocol usage and then exit

Table 3-5. Getting Help (Continued)

Command	Description
`qaucli -pr fc -h`	Print Fibre Channel and FCoE protocol usage and then exit
`qaucli -pr iscsi -h`	Print iSCSI protocol usage and then exit
`qaucli -npar -h`	Print NPAR (Switch Independent Partitioning) commands usage and then exit

Table 3-6 lists miscellaneous Windows diagnostics commands.

Table 3-6. Miscellaneous Commands

Command	Description
`qaucli -v`	Print version number and then exit
`qaucli -h`	Print usage and then exit

Table 3-7 lists the Windows CLI diagnostic test commands. Note that while running these tests, network traffic is interrupted.

Table 3-7. Diagnostic Test Commands

Command [a]	Description
`-i` `--interface`	Specifies the interface type (NX_NIC, NX_NIC1, and so on)
`-a` `--all`	Perform all test, regardless of default
`-D` `--default`	Perform only default test
`-R` `--CRegs`	Test all control registers (default)
`-I` `--IRQS`	Test interrupt mechanism (default)
`-L` `---IntLB`	Internal loopback test (default)
`-H` `--Hw`	Hardware test (default)
`-S` `--LinkST`	Link status test (default)
`-nR` `--noCRegs`	No control registers test (combine `-D` or `-a`)
`-nI` `--noIRQS`	No interrupt test (combine `-D` or `-a`)

Table 3-7. Diagnostic Test Commands (Continued)

Command [a]	Description
`-nL` `--noIntLP`	No internal loopback test (combine `-D` or `-a`)
`-nH` `--noHw`	No hardware test (combine `-D` or `-a`)
`-nS` `--noLinkSt`	No link status test (combine `-D` or **`-a`**)
`-h` `--help`	View help text

[a] All commands must be prefaced by `qaucli -pr nic -qldiag`. For example, to view the help, issue the following commands:
```
qaucli -pr nic -qldiag --help
```

For every test performed, the diagnostic utility shows the following data:

- Total number of attempts
- Number of successful completions
- Number of failures

You can either perform all tests in succession or perform only specific tests specified by the preceding command-line parameters.

You can run additional diagnostics in the CLI as listed in the following table. To determine the *cna_port_inst*, issue the `qaucli -nic -i` command as shown in Table 3-8.

Table 3-8. Running Windows Diagnostic Tests in the CLI

Test Type	Command
External Loopback	`qaucli -nic -extloopback <cna_port_inst> <tests_num> <on_error>` where `<tests_num>` specifies the number of tests, 1–65535, and `<on_error>` is either 0=Ignore or 1=Abort. **NOTE:** This test requires a pass-through module to be configured for both ports. Test runs between two ports. Single port loopback is not supported.
Flash	`qaucli -nic -testflash [cna_port_inst]`
Hardware	`qaucli -nic -testhw [cna_port_inst]`
Internal Loopback	`qaucli -nic -intloopback <cna_port_inst> <tests_num> <on_error>` where `<tests_num>` specifies the number of tests, 1–65535, and `<on_error>` is either 0=Ignore or 1=Abort.
Interrupt	`qaucli -nic -testinterrupt [cna_port_inst]`

Table 3-8. Running Windows Diagnostic Tests in the CLI (Continued)

Test Type	Command
Link	`qaucli -nic -testlink [`*`cna_port_inst`*`]`
Ping (IPv4)	`qaucli -nic -ping <`*`cna_port_inst`*`> <`*`hostname_or_IPv4`*`>` `[<`*`count`*`> <`*`pocket_size`*`> <`*`timeout_ms`*`> <`*`TTL`*`>]` where the default values are <*count*>=5, <*pocket_size*>=525, <*time-out_ms*>=1000, and <*TTL*>=30.
Register	`qaucli -nic -testregister [`*`cna_port_inst`*`]`
Transceiver DMI Data	`qaucli -nic -trans [`*`cna_port_inst`*`]`

Windows Diagnostic Test Descriptions

This section provides descriptions of the following Windows diagnostic tests.

■ Hardware Test

The hardware test checks the status of various hardware blocks, including DMA engines, receive engine, and on-board processor meta cores.

■ Register Test

The register test performs device register read/write accesses.

■ Interrupt Test

The interrupt test checks the ability of the hardware to create an interrupt and the ability of the driver to process the interrupt by forcing the generation of a predetermined number of interrupts. The test succeeds if the device generates the interrupts and the driver processes all interrupts expected.

■ External Loopback Test

The external loopback test verifies the transmit and receive path by sending a packet out onto an external loopback cable.

■ Internal Loopback Test

The internal loopback test verifies the capability of the ProductLine to internally loop back packets.

NOTE

Loopback tests are enabled only when the 8200 and 3200 Series Adapters are running firmware version 4.09.24 or later. When the loopback tests are running at the same time Fibre Channel or iSCSI protocols are running, refresh messages may appear. To avoid these messages, either click Cancel to ignore the messages or stop the `qlremote` and `iqlremote` agents while running the loopback tests on a NIC port.

■　Link Test

The link test inspects the link status (up or down) by checking the physical communication channel between the host and the firmware.

Windows Diagnostic Test Messages

If a test fails, an appropriate error code is generated and displayed, as shown in Table 3-9. Note that this table does not list error messages for the Interrupt and Link tests.

Table 3-9. Windows Diagnostic Test Messages

Test	Error Message	Description
Loopback	LB_TEST_OK	Loopback test has passed
Loopback	LB_SEND_WAIT_QUEUE_ERR	Send queue blocked
Loopback	LB_NORCV_ERR	Receive packet not received
Loopback	LB_NOMEM_ERR	No memory error
Loopback	LB_TX_QUEUE_ERR	Transmit queue error
Loopback	LB_SHORT_DATA_ERR	Looped data short error
Loopback	LB_SEQUENCE_ERR	Looped data out of sequence
Loopback	LB_DATA_ERR	Looped data corrupted
Loopback	LB_ERR_CNT	Looped error count
Register	CR_TEST_OK	Control register test passed
Register	CR_NIU_MODE	Network interface unit (NIU) error
Register	CR_PHY	Physical layer (PHY) error
Register	CR_ERRCNT	Control register error count
Hardware	HW_TEST_OK	Hardware test has passed
Hardware	HW_DMA_BZ_0	DMA channel 0 is busy
Hardware	HW_DMA_BZ_1	DMA channel 1 is busy
Hardware	HW_DMA_BZ_2	DMA channel 2 is busy
Hardware	HW_DMA_BZ_3	DMA channel 3 is busy
Hardware	HW_SRE_PBI_HALT	Segmentation and reassembly engine currently halted
Hardware	HW_SRE_L1IPQ	Segmentation and reassembly engine currently paused due to L1 IPQ discard failure
Hardware	HW_SRE_L2IPQ	Segmentation and reassembly engine currently paused due to L2 IPQ discard failure
Hardware	HW_SRE_FREEBUF	Segmentation and reassembly engine free buffer list is currently empty.
Hardware	HW_IPQ	IPQ is currently not empty

Table 3-9. Windows Diagnostic Test Messages (Continued)

Test	Error Message	Description
Hardware	HW_PQ_W_PAUSE	PQ write pause previously detected
Hardware	HW_PQ_W_FULL	PQ write full previously detected
Hardware	HW_IFQ_W_PAUSE	IFQ write pause previously detected
Hardware	HW_IFQ_W_FULL	IFQ write full previously detected
Hardware	HW_MEN_BP_TOUT	Memory backpressure time-out previously detected
Hardware	HW_DOWN_BP_TOUT	Downstream backpressure time-out previously detected
Hardware	HW_FBUFF_POOL_WM	Free buffer pool low watermark previously detected
Hardware	HW_PBUF_ERR	Packet buffer error previously detected
Hardware	HW_PBUF_ERR	Packet buffer error previously detected
Hardware	HW_FM_MSG_HDR	FM message header error previously detected
Hardware	HW_FM_MSG	FM message error previously detected
Hardware	HW_EPG_CTRL_Q	Egress packet generator (EPG) control queue is backed up
Hardware	HW_EPG_MSG_BUF	EPG message buffer error
Hardware	HW_EPG_QREAD_TOUT	EPG read queue time-out
Hardware	HW_EPG_QWRITE_TOUT	EPG write queue time-out
Hardware	HW_EPG_CQ_W_FULL	EPG completion queue write full
Hardware	HW_EPG_MSG_CHKSM	EPG message checksum error
Hardware	HW_EPG_MTLQ_TOUT	EPG MTL queue fetch time-out
Hardware	HW_PEG0	Peg 0 is not used
Hardware	HW_PEG1	Peg 1 is not used
Hardware	HW_PEG2	Peg 2 is not used
Hardware	HW_PEG3	Peg 3 is not used
Hardware	HW_ERRCNT	Hardware error count
Interrupt	INT_TEST_OK	Test okay
Interrupt	INT_TEST_ERR	Test error
Link	LINK_TEST_UP	Link up
Link	LINK_TEST_DOWN	Link down

For example:

```
qaucli  -nic -testlink
=== Link Test for 1. CNA Port Index ===
Function is not supported by this hardware/driver/api stack
=== Link Test for 2. CNA Port Index ===
Function is not supported by this hardware/driver/api stack
=== Link Test for 3. CNA Port Index ===
Function is not supported by this hardware/driver/api stack
=== Link Test for 4. CNA Port Index ===
Function is not supported by this hardware/driver/api stack
=== Link Test for 5. CNA Port Index ===
Link Test Starts ...
Test Status:           Passed  (Passed=1, Failed=0, ErrorCode=0)
Register Test Results:
  Status=Passed
  Passed=1,  Failed=0, ErrorCode=0

=== Link Test for 6. CNA Port Index ===
Link Test Starts ...
Test Status:           Passed  (Passed=1, Failed=0, ErrorCode=0)
Register Test Results:
  Status=Passed
  Passed=1,  Failed=0, ErrorCode=0
```

Linux Management Applications

Linux management applications for the adapter include the following:

- Linux NIC Driver Management Applications
- User Diagnostics for Linux NIC Driver Management Applications

Linux NIC Driver Management Applications

The following sections describe how to configure and manage the driver and adapter using Linux management utilities:

- Overview
- Viewing and Changing Adapter Properties on Linux

Overview

The following sections describe how to configure and manage the driver and adapter using Linux management utilities.

Viewing and Changing Adapter Properties on Linux

Use the following Linux utilities to view and change adapter properties:

- QConvergeConsole CLI (qaucli) Utility
- ethtool Utility

QConvergeConsole CLI (qaucli) Utility

Install QConvergeConsole CLI (qaucli) from the following packages supplied by QLogic:

- Package file:
  `QConvergeConsoleCLI-<version>_linux_<arch>.install.tar.gz`

- RPM installer package file:
  `QConvergeConsoleCLI-<version>_<arch>.rpm`

Example RPM package installation:

1. To determine if QConvergeConsole is installed and to find the full name of the installed QConvergeConsole RPM package, issue the following command using the partial name "QConvergeConsoleCLI" as an argument to grep:

   `rpm –qa | grep QConvergeConsoleCLI`

2. To check for an older version of the RPM package, issue the following command:

   `rpm –qa QConvergeConsoleCLI`

3. If an older version is found, erase it by issuing the following command:

```
rpm –e QConvergeConsoleCLI
```

4. To install the new version, issue the following command:

```
rpm –ihv QConvergeConsoleCLI-<version>.i386.rpm
```

The utility is installed in the `/opt/QLogic_Corporation/QConvergeConsoleCLI` directory.

Some software releases require firmware to be updated in the NIC's Flash memory. See the release notes for your software package to see if this is necessary. If so, use the QConvergeConsole CLI tool to write to the card's Flash memory.

The qaucli utility forces the QConvergeConsole CLI to use the NIC protocol interface to update the Flash memory. For example:

```
.\qaucli -pr nic .updimages [cna_instance] image_file
```

To have the QConvergeConsole CLI automatically select a protocol interface to update the Flash memory, issue the following command:

```
.\qaucli -nic .updimages [cna_instance] image_file
```

ethtool Utility

Use the ethtool utility to view adapter statistics and configure interface options. For additional details, refer to the qlcnic driver man page and ethtool man page.

Examples:

To disable transmit segmentation offload, issue the following command, where [*n*] represents a numerical value for a specific instance:

```
ethtool -K eth[n] tso off
```

To list interface statistics, issue the following command, where [*n*] represents a numerical value for a specific instance:

```
ethtool –S eth[n]
```

Sample Output 1:
```
ethtool –S eth8
NIC statistics:
xmit_called: 6
xmit_finished: 6
rx_dropped: 0
tx_dropped: 0
csummed: 0
rx_pkts: 0
```

```
        lro_pkts: 0
        rx_bytes: 0
        tx_bytes: 468
        lrobytes: 0
        lso_frames: 0
        xmit_on: 0
        xmit_off: 0
        skb_alloc_failure: 0
        null_skb: 0
        null_rxbuf: 0
        rx_dma_map_error: 0
```

In the following example, ethtool eth[*n*] lists interface settings.

Sample Output 2:

Ethtool eth8
```
Settings for eth8:
Supported ports: [ TP FIBRE ]
Supported link modes:
Supports auto-negotiation: No
Advertised link modes:  10000baseT/Full
Advertised auto-negotiation: No
Speed: 10000Mb/s
Duplex: Full
Port: FIBRE
PHYAD: 1
Transceiver: external
Auto-negotiation: off
Supports Wake-on: g
Wake-on: g
Current message level: 0x00000000 (0)
Link detected: yes
```

User Diagnostics for Linux NIC Driver Management Applications

This section covers the following information for user diagnostics for Linux NIC driver management applications:

- Running Linux User Diagnostics
- Linux Diagnostic Test Descriptions
- Linux Diagnostic Test Messages

Running Linux User Diagnostics

Linux user diagnostics include QConvergeConsole diagnostics and ethtool diagnostics.

QConvergeConsole Diagnostics

> **NOTE**
>
> For information on installing and starting the QConvergeConsole GUI, refer to the *QConvergeConsole GUI Installation Guide* (for download instructions, see "Related Materials" on page xii). All procedural information for the QConvergeConsole GUI is covered in the QConvergeConsole GUI's online help system.

QConvergeConsole CLI-based diagnostics include the following commands:

■ To enable or disable the port beacon, issue the following command:

```
qaucli -pr nic -beacon [cna_port_inst] <on|off>
```

■ To run an internal loopback test, issue the following command:

```
qaucli -pr nic -intloopback <cna_port_inst> <tests_num> <on_error>
```

where *tests_num* is the number of tests (1–65535) and *on_error* is either 0=Ignore or 1=Abort

■ To perform a Flash test, issue the following command:

```
qaucli -pr nic testflash [cna_port_inst]
```

■ To perform a hardware test, issue the following command:

```
qaucli -pr nic -testhw [cna_port_inst]
```

■ To perform an interrupt test, issue the following command:

```
qaucli -pr nic -testinterrupt [cna_port_inst]
```

■ To perform a link test, issue the following command:

```
qaucli -pr nic -testlink [cna_port_inst]
```

■ To perform a register test, issue the following command:

```
qaucli -pr nic -testregister [cna_port_inst]
```

■ To display transceiver DMI data, issue the following command:

```
qaucli -pr nic -trans [cna_port_inst]
```

Ethtool Diagnostics

To perform an adapter self-test using ethtool-based diagnostics, issue the following command:

```
# ethtool -t eth<x> offline
```

The self-test includes the following:

■ Loopback test
■ Interrupt test
■ Link test
■ Register test

Examples:

```
# ethtool -t eth8 offline
   The test result is PASS
   The test extra info:
   Register_Test_on_offline        0
   Link_Test_on_offline            0
   Interrupt_Test_offline          0
   Loopback_Test_offline           0
```

```
# ethtool -t eth4
   The test result is PASS
   The test extra info:
   Register_Test_on_offline          0
   Link_Test_on_offline              0
   Interrupt_Test_offline            0
   Loopback_Test_offline             0
```

Linux Diagnostic Test Descriptions

- The *internal loopback test* performs internal packet loopback.

- The *Flash test* verifies the Flash read and write.

- The *hardware test* verifies that the hardware is running.

- The *interrupt test* enables and disables the interrupt and functional verification tests.

- The *link test* verifies that the port is linked, meaning that the port has a good cable attached to the port and that other end of the cable is connected to an operational Ethernet port, either another NIC port or a network device, such as a switch.

- The *register test* verifies the NIC register read and write.

Linux Diagnostic Test Messages

Test information and PASS or FAIL messages are displayed for each of the tests listed in "Linux Diagnostic Test Descriptions" on page 110.

VMware Management Applications

VMware management applications for the adapter include the following:

- VMware NIC Driver Management Applications
- User Diagnostics for VMware NIC Driver Management Applications

VMware NIC Driver Management Applications

This section provides the following information about the VMware NIC driver management applications:

- Overview
- Using Switch Independent Partitioning Under ESX

Overview

VMware vSphere provides a CIM monitoring framework for both classic ESX and ESXi. Use this framework to configure and manage Ethernet interfaces. On an ESX Server, you can configure and manage Ethernet interfaces using a command line interface. On an ESXi Server, use the remote CLIs (RCLIs).

Using Switch Independent Partitioning Under ESX

All Switch Independent Partitioning Ethernet functions are enumerated by the hypervisor, controlled by the driver running in the hypervisor, and configured similar to other Ethernet interfaces. For more details, see "Switch Independent Partitioning" on page 122.

You would typically create a virtual switch (vSwitch) for each Switch Independent Partitioning interface. You can configure virtual machines (VMs) to use the standard virtual network devices, such as VMXNET 3 adapters. On each interface, you can configure features such as NetQueue. Although it is possible to configure all standard aggregation and failover configurations supported by ESX using Switch Independent Partitioning interfaces as uplinks, it is not useful to have multiple uplinks with NIC partitions belonging to the same physical port.

The vSwitch switches the packets when VMs are assigned to share the same NIC partition. The embedded switch (eSwitch) on the adapter switches packets when packets are sent by a VM on a NIC partition destined to another VM that is connected to a vSwitch on a different NIC partition belonging to the same physical port. If the two NIC partitions are associated with different physical ports, the external switch forwards packets between them.

Perform vSwitch configuration using either standard service console commands or RCLI commands. For specific command options, refer to the VMware documentation.

> **NOTE**
>
> The adapter has eSwitches. Because all of the interfaces are assigned to the hypervisor, eSwitch operation is transparent, and the system administrator is not required to perform any additional configuration.

User Diagnostics for VMware NIC Driver Management Applications

This section covers user diagnostics for VMware NIC driver management applications.

Ethtool

Use the ethtool utility to view adapter statistics and configure interface options.

Supported options include the display of device information, protocol offload options, driver information, adapter statistics, and more. For details, refer to the ethtool man page.

To view driver and firmware versions, issue the following command:

```
ethtool -i vmnic[n]
```

Example:

```
# ethtool -i vmnic7
driver: qlcnic
version: <version>
```

```
firmware-version: <version>
bus-info: 0000:10:00.3
```

To view interface statistics, issue the following command:

```
ethtool -S vmnic[n]
```

Example:

```
# ethtool -S vmnic7
NIC statistics:
    bad_skblen: 0
    rcvd_badskb: 0
    xmit_called: 617618
    xmit_finished: 617618
    rx_dropped: 0
    tx_dropped: 0
    csummed: 777792
    rx_pkts: 897558
    lro_pkts: 0
    rx_bytes: 78433279
    tx_bytes: 60526284
```

Unified Extensible Firmware Interface

This section provides user notes on the UEFI driver contained in the 10Gb multi-boot image. It includes following sections:

- UEFI Package Contents
- Supported Features
- Fibre Channel Adapter Configuration
- Updating the UEFI (EfiUtilx64)

UEFI Package Contents

Files in the UEFI driver package include the following:

- `p3pxxxxx.bin` is a combined binary file, which includes the binaries for BIOS, FCode, UEFI driver, and reduced instruction set computer (RISC) firmware.

- `EfiUtil.EFI` or `EfiUtilx64.EFI` is the UEFI utility to update the UEFI driver, firmware, and NVRAM.

- `QL8XXX.DRV` is the auxiliary driver file used for updating the adapter.

- `UPDATE.NSH` is the UEFI script file used to update the adapter on a UEFI system.

- `EFICFG.PDF` is the adapter configuration and diagnostic PDF file.

- `EFIUTIL.PDF` is the adapter EfiUtil utility PDF file.

The preceding PDF files are included in the boot code release package in the `EFI` directory.

> **NOTE**
>
> All bin, uefi, and nsh files are required to update the adapter on a UEFI system.

Supported Features

The UEFI driver supports the following features:

- UEFI specification 1.10, 2.1
- UEFI protocols: Extended SCSI Pass Thru, Block IO, SCSI IO, Driver Diagnostics, Driver Configuration, and Component Name
- 382 targets
- 2,047 LUNs per target
- Fabric topology
- Boot from SAN
- Compatible with RHEL 6.1, SLES® 11 SP4, SLES 11 SP 1 x64, and Windows Server 2008 x64 on UEFI Intel® x64 based systems

Fibre Channel Adapter Configuration

To configure the adapter parameters, refer to the `EFICFG.PDF` file included in the boot code release package in the `EFI` directory.

Updating the UEFI (EfiUtilx64)

For systems that do not include a built-in UEFI Shell, obtain the UEFI Shell program from an open source site. For example: https://www.tianocore.org/.

EfiUtilx64 is a QLogic Flash programming utility for UEFI. To run the utility, boot the UEFI Shell. Do not run this utility from a drive connected to a Converged Network Adapter.

Make sure `update.nsh`, `efiutilx64.efi`, `ql8xxx.drv`, and `p3pxxxxx.bin` source files are in the same directory.

To update the UEFI driver and RISC firmware on all adapters:

1. Unzip and copy the update kit files to the root level of a USB storage device.

> **NOTE**
>
> Do not change or rearrange the kit directories or files that were unzipped and copied to the USB storage device.

2. Connect the USB device to the UEFI-based system with the 8*xxx* adapters.

3. At the system's UEFI shell prompt, issue the `map -r` command to map the USB device file system. You can check the mapping as follows:

```
map -b
```

4. Locate the USB device and change to that device. For example, if the USB device is mapped to fs9 after the `map -r`:

```
fs9: <Enter>
```

The UEFI Shell prompt changes as follows: `fs9:\>`

5. To update the UEFI driver and RISC firmware, run the `update.nsh` script. For example:

```
fs9:> update.nsh
```

The `update.nsh` calls `efiutilx64.efi` to update all of the adapters.

6. For the changes to take effect, power the system off and then back on.

To update NVRAM or perform other manual EfiUtilx64 Flash functions, refer to the information in the `efiutil.pdf` file.

To write, read, or verify the multi-boot image on the adapter, use EfiUtilx64.

NOTE

To perform other manual EfiUtilx64 Flash functions, refer to the `efiutil.pdf`.

Configuring iSCSI over DCBX

This section describes the configuration of iSCSI over DCBX on a Converged Enhanced Ethernet (CEE) switch port on a Brocade 8000 Series FCoE switch. This configuration enables an iSCSI VLAN on a QLogic iSCSI Host Bus Adapter. It also enables iSCSI over DCBX on the Brocade 8000 switch to set priority flow control (PFC) for iSCSI traffic. With these configurations, the iSCSI devices attached to the switch performs iSCSI over DCBX. The configuration is achieved using the switch's CLI to configure the switch port to run iSCSI traffic with the user-specified bandwidth and PFC and the QConvergeConsole CLI utility to enable and configure the VLAN for the QLogic adapter.

NOTE

- iSCSI over DCBX applies only to the iSCSI Host Bus Adapter. It does not apply to iBFT/SW or an iSCSI function type on a NIC port configured with Switch Independent Partitioning.
- When bandwidth settings exist for both Switch Independent Partitioning and DCBX, *DCBX takes precedence over Switch Independent Partitioning*. DCBX sets the bandwidth for iSCSI and NIC traffic, and then Switch Independent Partitioning sets the bandwidth for the NIC partitions by dividing the NIC bandwidth allocated by DCBX. For more information, see "Interoperation of Bandwidth Settings for DCBX and Switch Independent Partitioning" on page 120.
- For information related to the CEE switch, refer to the manufacturer's documentation.
- This section provides an example in which the port number is 0 and the VLAN ID is 99. Replace these values with whatever is appropriate for your situation.

The procedure is detailed in the following sections:

- Configuring the iSCSI VLAN on the QLogic Adapter
- Configuring the Switch for iSCSI over DCBX
- Interoperation of Bandwidth Settings for DCBX and Switch Independent Partitioning
- Choosing DCBX or Switch Independent Partitioning

Configuring the iSCSI VLAN on the QLogic Adapter

Follow these steps to configure a VLAN to run iSCSI traffic on the adapter:

NOTE

Before you perform the following procedure:

- The QConvergeConsole utility (qaucli) must be installed.
- The iSCSI IP settings must be configured.

---

1. Open a command prompt.

2. Issue the following commands to configure a VLAN on port 0 with VLAN ID 99:

   ```
   - qaucli –pr iscsi –n 0 VLAN on
   - qaucli –pr iscsi –n 0 VLANID 99
   ```

3. Issue the following command to confirm the VLAN settings for port 0:

```
- qaucli –pr iscsi –c 0
```

Configuring the Switch for iSCSI over DCBX

Configuring the Brocade 8000 CEE switch involves the following steps:

1. Verify the Version of the Switch Firmware
2. Create and Configure the iSCSI VLAN on the Switch
3. Create and Configure the CEE Map for iSCSI Traffic Bandwidth and PFC
4. Configure LLDP/DCBX for the iSCSI TLV
5. Configure the CEE Port's iSCSI Traffic Class
6. Verifying Adapter/Switch Status for iSCSI Login, Traffic, and PFC

Verify the Version of the Switch Firmware

The switch firmware must be version 6.4.1 or later to support the iSCSI TLV. To check the switch firmware version, issue the following command:

```
admin> version
```

Update the switch firmware if the version is earlier than 6.4.1. For instructions on updating the firmware, refer to the switch's documentation.

Create and Configure the iSCSI VLAN on the Switch

Issue the following commands to create the iSCSI VLAN on the switch.

```
admin> cmsh
# configure terminal
swd77(config)#interface vlan 99
swd77(conf-if-vl-99)#exit
```

The preceding commands create a VLAN with VLAN ID 99 on the switch.

Create and Configure the CEE Map for iSCSI Traffic Bandwidth and PFC

In this example, the following commands are issued to set iSCSI traffic for 90 percent bandwidth and PFC with group ID 7 and to set all other traffic for 10 percent bandwidth and no PFC with group ID 6.

1. Create a CEE map with name `iscsi-dcbx`:

   ```
   >swd77(config)#cee-map iscsi-dcbx
   ```

2. Allocate 90 percent bandwidth for the priority traffic in group 7:

   ```
   swd77(conf-ceemap)#priority-group-table 7 weight 90 pfc
   ```

3. Allocate 10 percent bandwidth for the nonpriority traffic in group 6:

   ```
   swd77(conf-ceemap)#priority-group-table 6 weight 10
   ```

4. Map COS to group for the allocated traffic:

```
swd77(conf-ceemap)#priority-table 6 6 6 6 6 6 6 7
swd77(conf-ceemap)#exit
```

Configure LLDP/DCBX for the iSCSI TLV

The following commands configure link layer discovery protocol (LLDP) for the iSCSI type-length-value (TLV).

1. Configure the LLDP:

   ```
   swd77(config)#protocol lldp
   ```

2. Enable the LLDP:

   ```
   swd77(conf-lldp)#no disable
   ```

3. Advertise DCBX TLV in the LLDP:

   ```
   swd77(conf-lldp)#advertise dcbx-tlv
   ```

4. Advertise iSCSI over DCBX in the LLDP:

   ```
   swd77(conf-lldp)#advertise dcbx-iscsi-app-tlv
   swd77(conf-lldp)#exit
   ```

When you use the Linux lldptool to enable the link layer discovery protocol (LLDP) management type-length-value (TLV) from the host, data center bridging (DCB) becomes disabled.

Cause: Even when the Linux driver indicates that LLDP is managed by another entity (for example, if LLDP is managed by the adapter firmware, and the driver capability has the DCBX_CAP_LLDP_MANAGED flag), the LLDP Agent Daemon (lldpad) reads local values from its `config` file rather than from the driver. This misread causes lldpad to send different values than the adapter firmware sends. The incorrect values lead to a multiple LLDP neighbors condition, and according to the *DCB Capability Exchange Protocol Base Specification*, data center bridging exchange (DCBX) is disabled when this condition is true.

Fix: If the driver capability has the DCBX_CAP_LLDP_MANAGED flag, lldpad should either read all values from the driver through dcbnl_ops, or it should send non-DCB TLVs. The lldpad has been fixed to address the issue using the non-DCB TLV option. The fix is available in the latest lldpad and lldptool, which you can download from the following Web site:

http://open-lldp.org/git/?p=open-lldp.git;a=shortlog

To obtain the latest code from the Web site's tree, click **snapshot** to save the source code to a setup. Use the code to build, install, and test lldpad and lldptool on the setup.

Configure the CEE Port's iSCSI Traffic Class

The following commands configure the switch port to which the QLogic adapter is connected. In this example, the adapter is connected to port 0/16 of the switch.

1. Set the switching characteristics:

```
swd77(config)#interface tengigabitethernet 0/16
swd77(conf-if-te-0/16)#switchport
```

2. Set the interface as converged:

```
swd77(conf-if-te-0/16)#switchport mode converged
```

3. Add VLAN 99 on this switch port (iSCSI traffic host is configured for VLAN 99, so this port allows iSCSI traffic):

```
swd77(conf-if-te-0/16)#switchport converged allowed vlan add 99
```

4. Enable TX and RX pause on this switch port:

```
swd77(conf-if-te-0/16)#qos flowcontrol tx on rx on
```

5. Configure cos 7 for iSCSI on this switch port:

```
swd77(conf-if-te-0/16)#qos cos 7
```

6. Configure previously created CEE map on this switch port:

```
swd77(conf-if-te-0/16)#cee iscsi-dcbx
```

7. Apply LLDP iSCSI priority bits to interface:

```
swd77(conf-if-te-0/16)#lldp iscsi-priority-bits 0x80
```

8. Bring up the link for the port:

```
swd77(conf-if-te-0/16)#no shutdown
swd77(conf-if-te-0/16)#exit
```

Verifying Adapter/Switch Status for iSCSI Login, Traffic, and PFC

The best way to verify the iSCSI login, traffic, and PFC is to use a network analyzer to analyze the network traffic. The switch and adapter counters also can be used to verify iSCSI traffic and PFCs.

The following commands (in bold text) can be used to verify the port settings:

```
swd77(config)#do show running-config interface tengigabitethernet 0/16
switchport
switchport mode converged
switchport converged allowed vlan add 99
no shutdown
lldp iscsi-priority-bits 0x80
qos cos 7
cee iscsi-dcbx
```

```
swd77(config)#do show lldp interface tengigabitethernet 0/16
LLDP information for Te 0/16
State: Enabled
Mode: Receive/Transmit
Advertise Transmitted: 30 seconds
Hold time for advertise: 120 seconds
Re-init Delay Timer: 2 seconds
Tx Delay Timer: 1 seconds
DCBX Version : CEE
Auto-Sense : Yes
Transmit TLVs: Chassis ID Port ID
TTL IEEE DCBx
DCBx FCoE App DCBx FCoE Logical Link
Link Prim Brocade Link DCB
x iSCSI App
DCBx FCoE Priority Bits: 0x8
DCBx iSCSI Priority Bits: 0x80


swd77(config)#do show cee maps iscsi-dcbx
CEE Map iscsi-dcbx
Precedence 1
Priority Group Table
6: Weight 10, PFC Disabled, TrafficClass 6, BW% 10
7: Weight 90, PFC Enabled, TrafficClass 7, BW% 90
15.0: PFC Disabled
15.1: PFC Disabled
15.2: PFC Disabled
15.3: PFC Disabled
15.4: PFC Disabled
15.5: PFC Disabled
15.6: PFC Disabled
15.7: PFC Disabled
Priority Table
 CoS: 0 1 2 3 4 5 6 7
-------------------------------------------
PGID: 6 6 6 6 6 6 6 7
FCoE CoS: None
Enabled on the following interfaces
Te 0/16
```

Interoperation of Bandwidth Settings for DCBX and Switch Independent Partitioning

If you want to run iSCSI and NIC traffic together, DCBX can be used to set the bandwidth percentage to be shared among the iSCSI and NIC. If you want to run partitioned NIC traffic, Switch Independent Partitioning should be used to set the percentage of bandwidth that is shared among the multiple NIC partitions. The NIC partition uses the share of NIC bandwidth that was allocated by DCBX; therefore, DCBX takes precedence over Switch Independent Partitioning for NIC traffic.

If NIC traffic is partitioned, the total bandwidth of any NIC partition is the NIC bandwidth allocated by DCBX multiplied by the bandwidth that Switch Independent Partitioning allocates to the NIC partition.

For example, suppose that DCBX allocates 60 percent of the total bandwidth to iSCSI traffic, and 40 percent to NIC traffic. That 40 percent is then re-allocated by Switch Independent Partitioning between the NIC partitions. If Switch Independent Partitioning allocates 75 percent to NIC partition 1 and 25 percent to NIC partition 2, then NIC partition 1 is allocated 30 percent of the total bandwidth (40 percent x 75 percent), and NIC partition 2 is allocated 10 percent of the total bandwidth (40 percent x 25 percent).

Choosing DCBX or Switch Independent Partitioning

The choice of DCBX or Switch Independent Partitioning is determined by the types of traffic that need to be run (iSCSI or NIC).

The following guidelines can be used when choosing and using DCBX and Switch Independent Partitioning for bandwidth settings.

If running:

■ Both iSCSI and NIC traffic, *with* partitioning of the NIC traffic:

Use DCBX to allocate the bandwidth for the iSCSI and NIC traffic; then, use Switch Independent Partitioning to allocate the NIC bandwidth (allocated by DCBX) among the NIC partitions.

■ Both iSCSI and NIC traffic, *without* partitioning of the NIC traffic:

Use DCBX to split the total bandwidth among iSCSI and NIC traffic. There is no need to use Switch Independent Partitioning.

■ NIC traffic only (no iSCSI traffic) *with* partitioning of the NIC traffic:

Use Switch Independent Partitioning to allocate the total bandwidth among the NIC partitions. DCBX is not needed.

- NIC traffic only (no iSCSI traffic) *without* partitioning of the NIC traffic:
  Neither Switch Independent Partitioning nor DCBX needs to be used.

- iSCSI traffic only (no NIC traffic):
  Neither Switch Independent Partitioning nor DCBX needs to be used.

Table 3-10 summarizes these guidelines.

Table 3-10. Choosing DCBX and Switch Independent Partitioning for Bandwidth Settings

To run these types of traffic			Use these tools	
iSCSI	NIC		DCBX	Switch Independent Partitioning
	Unpartitioned	Partitioned		
X	–	–	–	–
–	X	–	–	–
–	–	X	–	X
X	X	–	X	–
X	–	X	X	X

4 Switch Independent Partitioning

Overview

This chapter provides the following information about the QLogic Switch Independent Partitioning feature:

- Switch Independent Partitioning Setup Requirements
- Switch Independent Partitioning Configuration
- Switch Independent Partitioning Setup and Management Options
- Switch Independent Partitioning Setup

Switch Independent Partitioning Setup Requirements

This section provides hardware and software requirements for applying Switch Independent Partitioning functionality to QLogic adapters installed in host servers within SANs.

Hardware Requirements

Table 4-1. Hardware Requirements

QLogic Adapters	Platforms
QMD8262-k [a] Blade Network 10GbE Daughter Card	M620, M820
QLE8262[a] Monolithic Server 10GbE Standup Card	T420, T620 R320, R420, R520, R620, R715, R720, R815, R820, R910
QME8262-k 10GbE Blade Mezzanine Card	M420, M520, M610, M620, M710HD, M720, M820, M910, M915

[a] Link speed limitation: QMD8262-k, and QME8262-k support only 10GbE operation and do *not* link at 1Gbps with Dell M6220, M6348, and Cisco® 3x3x switches.

Software Requirements

Table 4-2. Operating System Requirements

Operating Systems	Platforms
Linux	■ RHEL 6.5—x64 only ■ SLES 11 SP3 ■ SLES 12
VMware	■ ESX/ESXi 5.0/5.1/5.5
Windows	■ Windows Server 2008 SP2, Windows Server 2008 R2 SP1 ❏ Hyper-V® [a] ■ Windows 2012 ■ Windows 2012 R2

[a] If a partitioned NIC is configured for use in a Hyper-V network virtualization stack, virtual message queue (VMQ) must be enabled. To enable VMQ on a management OS with physical network adapters less than 10Gbps, issue the following command in a command prompt window:

```
reg add HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\services\VMSMP\Parameters
/v BelowTenGigVmqEnabled /t REG_DWORD /d 1 /f
```

Table 4-3. Management Tool and Driver Requirements

SW Components [a]	File Names and Download Locations
Management Tools	
Dell System Setup, Lifecycle Controller, or other human interface infrastructure (HII) browser	http://support.dell.com
QLogic OptionROM	■ Pre-installed, written on the adapter's Flash memory at Dell factory
QLogic QConvergeConsole GUI/CLI	http://support.dell.com
QLogic QConvergeConsole Plug-ins for vSphere	http://support.dell.com
Drivers	
Adapter drivers	http://support.dell.com ■ VMware: Drivers are available on VMware Web site

[a] Ecosystem Requirements: Switch Independent Partitioning is a new feature from Dell and it requires the necessary ecosystem software/firmware to operate correctly. So in addition to the hardware/software requirements, it also requires user to upgrade to the latest System Setup, BIOS, Chassis Management Controller (CMC), and iDrac available from support.dell.com, especially if users are using the FlexAddress feature.

Switch Independent Partitioning Configuration

This section defines Switch Independent Partitioning configuration and describes the configuration options and the management tools you can use to set up Switch Independent Partitioning on QLogic adapters installed in 11th and 12th generation Dell PowerEdge blade servers.

In addition to defining what Switch Independent Partitioning is, this section describes:

- What Is Switch Independent Partitioning?
- Switch Independent Partitioning Options
- Personality Changes
- Quality of Service
- eSwitch
- Configuration Management Tools

What Is Switch Independent Partitioning?

Switch Independent Partitioning provides the ability to create multiple physical functions on the PCIe bus that share a single physical port. Each physical function is a PCI endpoint (PCIe) that can have a device driver attached to it.

The Switch Independent Partitioning feature in QLogic adapters allows you to partition a single 10GbE NIC port into up to four individual partitions with user-configurable bandwidth and interface type (personality). The partitioning options are not limited to NIC; they extend to converged fabric partitioning by enabling you to assign iSCSI or FCoE protocols to certain partitions.

For example, each partition can be either native Ethernet NIC or configured to support iSCSI or FCoE storage devices with different PCIe endpoint device class code (subject to the restrictions listed in Table 4-4 on page 129). Both iSCSI and FCoE operate in full hardware offload mode.

The Dell/QLogic Switch Independent Partitioning solution is OS and switch agnostic, which means Switch Independent Partitioning does not require a proprietary switch to operate; however, the adapter does require the OS-specific QLogic adapter driver for each supported protocol (NIC, iSCSI, and FCoE). The solution also means Switch Independent Partitioning bandwidth allocation can only regulate transmit traffic but not receive traffic.

After you have configured the NIC partitions as needed on the adapter ports, you must reboot the Dell PowerEdge server to make the personality changes take effect.

You can modify the minimum and maximum bandwidth for each switch-independent partition. The changes take effect immediately without rebooting the server. The minimum and maximum bandwidths are specified as percentages of the link bandwidth, where:

- The minimum bandwidth is the minimum bandwidth guaranteed to a partition.
- The maximum bandwidth is the maximum value that a partition is permitted to use.

> **NOTE**
>
> When bandwidth settings exist for both Switch Independent Partitioning and DCBX, *DCBX takes precedence over Switch Independent Partitioning*. DCBX sets the bandwidth for iSCSI and NIC traffic, and then Switch Independent Partitioning sets the bandwidth for the NIC partitions by dividing the NIC bandwidth allocated by DCBX. For more information, see "Interoperation of Bandwidth Settings for DCBX and Switch Independent Partitioning" on page 120.

Switch Independent Partitioning Options

The Switch Independent Partitioning feature in QLogic adapters provides the ability to create multiple PCIe physical functions for each physical 10GbE port on the adapter. Each PCIe function appears as an independent interface to the host operating system or hypervisor.

When the adapter is configured as an Ethernet-only adapter, it contains eight Ethernet functions.

By default, Switch Independent Partitioning functionality is disabled on the adapters, having only two Ethernet functions enabled. Depending on the feature personality mapping supported on the adapter, you can enable additional Ethernet or storage functions.

The PCI function number assignment is as follows:

- Functions 0 and 1 are always NIC, function 0 for port 1 and function 1 for port 2; any of the other functions can be individually enabled or disabled.
- Functions 2 and 3 can only be NIC personalities.
- Functions 4 and 5 can be configured with either iSCSI or NIC personality.
- Functions 6 and 7 can be configured with either FCoE or NIC personality.
- You can configure at most one iSCSI and one FCoE personality for each physical port.

Virtual machine (VM)-to-VM Ethernet traffic between VMs on different vSwitches is routed by the eSwitch if the communicating VMs are attached to NIC partitions derived from the same physical port. The eSwitch handles VM-to-VM communication by learning MAC addresses of the virtual NICs (VNICs) of the VMs. This capability enables the eSwitch to switch packets destined to another VM on the same host.

The adapter supports a maximum of 64 Layer-2 MAC address filters across all partitions. The NIC driver evenly distributes the number of filters across all NIC partitions. For example, if the NIC adapter has four NIC partitions, two NIC partitions per physical port, each NIC partition gets 16 filters (64/4 = 16).

The available number of MAC address filters per partition determines the number of VMs that can communicate to other VMs through the eSwitch without the eSwitch having to forward the traffic to an external switch.

> **NOTE**
>
> For most deployments, this limit of 64 Layer-2 MAC address filters is not reached and is only applicable when a large number of VMs need to communicate with other VMs on the same host with the eSwitch.

The VLAN and teaming solutions on partitioned NIC functions have the following restrictions:

- A fail-safe team cannot be created using NIC functions that belong to the same physical port. For example, you cannot choose PF2 as a backup for PF0 because both functions are partitions of the same physical port.
- 802.3ad link aggregation teams are not allowed on partitioned NIC functions.

Figure 4-1 shows the default Switch Independent Partitioning function settings.

NOTE

In NPAR configurations with teaming on ESXi 5.1 and ESXi 5.5, QLogic recommends setting the driver module parameter `defq_filters` to 0 by issuing the following command, and then rebooting the system for the setting to take effect.

`esxcfg-module –s defq_filters=0`

This setting enables the driver-provided MAC learning to be used rather than the vmkernel-provided callback to notify vMAC addresses.
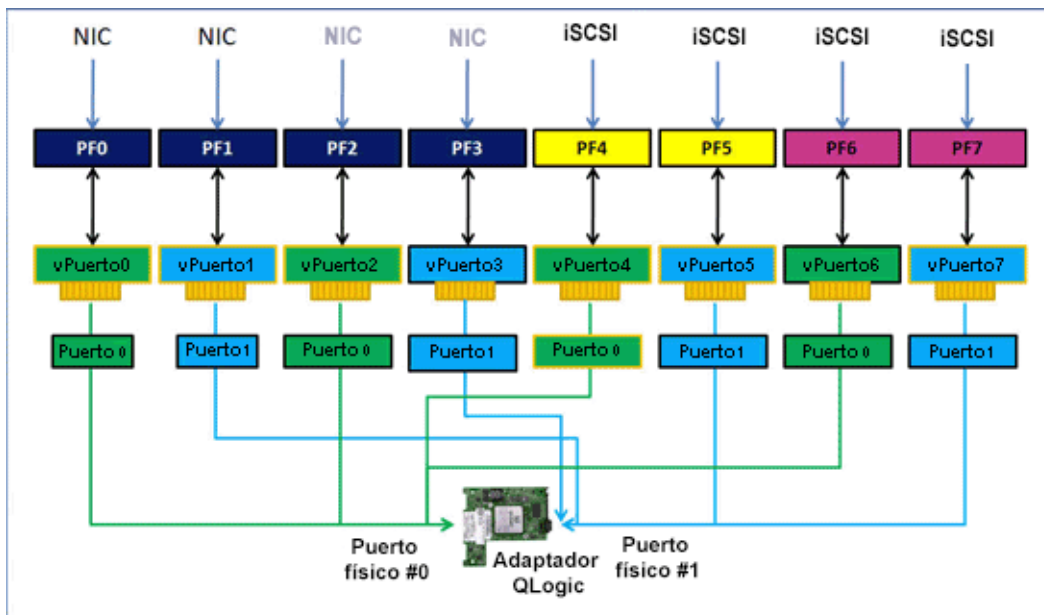
Figure 4-1. Switch Independent Partitioning Default Configuration (NIC, FCoE, and iSCSI)

Figure 4-2 shows the possible configurations.
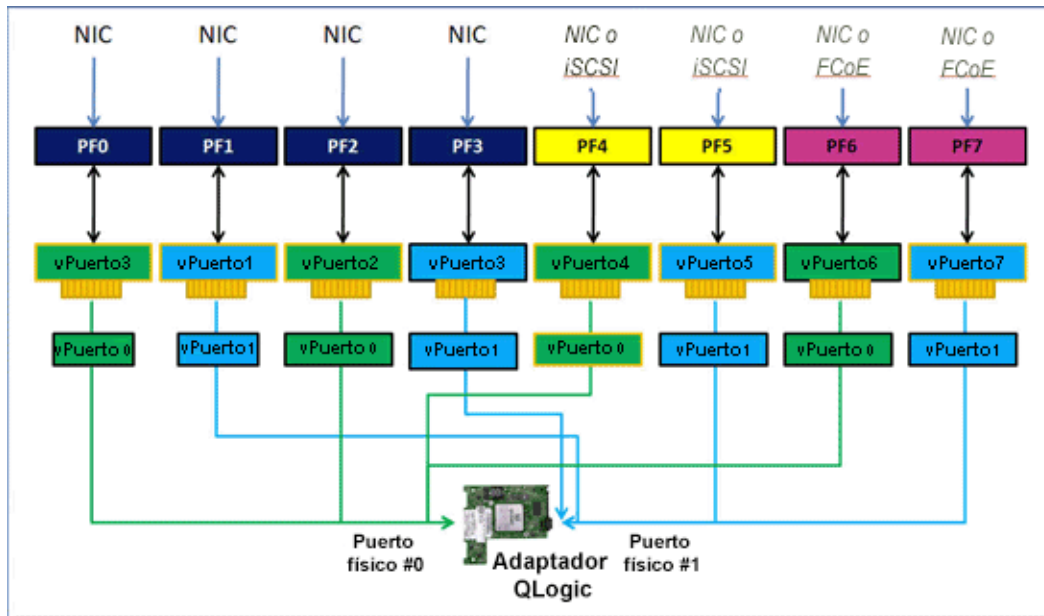
Figure 4-2. Switch Independent Partitioning Configuration Options (Personalities)

Personality Changes

Based on your operating environment, you can use your preferred management tool to change or disable PCI functions on either physical port. Using this feature lets you divide each physical port into up to four partitions, configured to support one of the following PCI function types: NIC, FCoE, or iSCSI.

> **NOTE**
>
> This document uses the terms *personality* and *function type* interchangeably.

Table 4-4 shows the port identifications and the possible Switch Independent Partitioning configurations.

Table 4-4. Configuration Options

Function Number	Function Type	Physical Port Number	
		User Label[a]	**System Number**[b]
0	NIC	1	0
1	NIC	2	1

Table 4-4. Configuration Options (Continued)

Function Number	Function Type	Physical Port Number	
		User Label [a]	System Number [b]
2	Disabled/NIC	1	0
3	Disabled/NIC	2	1
4	iSCSI/NIC/Disabled	1	0
5	iSCSI/NIC/Disabled	2	1
6	FCoE/NIC/Disabled	1	0
7	FCoE/NIC/Disabled	2	1

[a] The physical port number is displayed as Port 1 or Port 2 on the adapter's port's label.

[b] The physical port number is displayed as **Phy Port 0** or **Phy Port 1** on the Switch Independent Partitioning configuration screens for most of the management tools, except where noted otherwise.

Quality of Service

Quality of Service (QoS) refers to the *bandwidth allocation* assigned to each partition used to send and receive data between the adapter port and connected devices.

> **NOTE**
>
> When bandwidth settings exist for both Switch Independent Partitioning and DCBX, *DCBX takes precedence over Switch Independent Partitioning*. DCBX sets the bandwidth for iSCSI and NIC traffic, and then Switch Independent Partitioning sets the bandwidth for the NIC partitions by dividing the NIC bandwidth allocated by DCBX. For more information, see "Interoperation of Bandwidth Settings for DCBX and Switch Independent Partitioning" on page 120.

Each physical port on a QLogic adapter can send and receive data at up to 10Gbps in both directions at the same time. When the physical port is partitioned into four partitions, the port bandwidth is divided between each port partition according to traffic demands.

You can set QoS for each port partition by setting minimum and maximum percentages of the physical port's bandwidth for each partition. This feature helps guarantee a transmission rate for each partition that requires a particular bandwidth to run critical applications using port partitions. The setting for a given QoS can resolve bottlenecks that exist when VMs contend for port bandwidth.

Enhanced transition services (ETS) control the actual bandwidth allocation at the network port. The bandwidth allocation under ETS is typically 50 percent for FCoE traffic and 50 percent for non-FCoE traffic (NIC and iSCSI). This means that Switch Independent Partitioning QoS allocations among the NIC partitions for a given port, allocate a percentage of the non-FCoE portion of the bandwidth.

Switch Independent Partitioning QoS allows NIC and iSCSI partitions to each allocate a minimum guaranteed portion of the available bandwidth. However, the user interface tools do not support setting the bandwidth limits for iSCSI partition. This means that the iSCSI partition might not get the preferred bandwidth. For example, if the total minimum allocated bandwidth across the NIC partitions equals 100 percent, the iSCSI partition is limited to 1 percent of the NIC bandwidth portion in high-usage conditions.

To ensure that iSCSI has more than one percent of bandwidth available in high-usage conditions, set the total Switch Independent Partitioning QoS minimum bandwidth settings so that they equal less than 100 percent.

For example:

- A Switch-Independent-Partitioning-enabled port has two NIC partitions, one iSCSI partition and one FCoE partition.

- ETS allocates 50 percent of the network bandwidth to FCoE traffic and 50 percent to non-FCoE traffic.

- The Switch Independent Partitioning QoS minimum bandwidth setting for each NIC partition is 50 percent.

  ❑ This means that each NIC partition is guaranteed 50 percent of 50 percent of 10Gb or 2.5Gb each.

- If at any time the FCoE partition is using 5Gb of bandwidth and each NIC partition is using 2.5Gb, the iSCSI partition is left with only 50Mb of bandwidth.

- If, however, the NIC partitions each allocated 45 percent of the non-FCoE traffic, the total allocated bandwidth would be 90 percent.

  ❑ The remaining 10 percent (or 500Mb) is then effectively reserved for the iSCSi partition.

eSwitch

The adapters provide embedded switch functionality. This provides a basic Layer-2 switch for Ethernet frames. Each physical port has one instance of an eSwitch, which supports all NIC partitions on that physical port.

The eSwitch operation is transparent, and the administrator does not need to perform any specific configuration. The ability to view eSwitch statistics depends on your operating environment and management tool.

The QLogic drivers download the VM MAC addresses to the firmware. This enables the firmware and hardware to switch the packets destined for VMs on the host.

For traffic to flow from one eSwitch to another it must first pass through an external switch or have been forwarded by a VM that has a path through both eSwitches.

Configuration Management Tools

Depending on your operating environment and preferred system management techniques, you can use any of the following tools to set up NIC partitions on adapter ports to meet your system's networking requirements:

- Dell System Setup
- QLogic OptionROM at POST
- QConvergeConsole GUI
- QConvergeConsole CLI
- Windows Device Manager
- CIM Provider and QConvergeConsole VMware vCenter Server Plug-in for VMware ESX/ESXi

Dell System Setup

Administrators of 11th and 12th generation Dell PowerEdge servers can use the server's embedded Dell System Setup to set up switch-independent partitions and view eSwitch statistics on installed adapters. The System Setup function provides OS-independent management functionality. It includes the following:

- Updates
- Hardware configuration
- OS deployment
- System diagnostics

This tool is especially helpful for configuring adapters without relying on an operating system because it runs from the server controller's firmware by initiating it during the power-on self test (POST).

For details about the System Setup, visit the Dell Web site (http://support.dell.com).

For procedures on using the Dell System Setup to set up Switch Independent Partitioning on QLogic adapters, see "Dell System Setup" on page 137.

QLogic OptionROM at POST

The QLogic OptionROM comes factory installed, written on the adapter's Flash memory at the Dell factory. When you first start the Dell PowerEdge server that contains QLogic adapters, the POST starts. Running the POST test gives you access to the OptionROM utility.

For procedures on setting up Switch Independent Partitioning and eSwitch parameters using the OptionROM while powering up the host server, see "QLogic OptionROM at POST" on page 144.

QConvergeConsole GUI

The QConvergeConsole Unified Adapter Web Management Interface is a Web-based client/server application that allows for centralized management and configuration of QLogic adapters within the entire network (LAN and SAN).

On the server side, the QConvergeConsole runs as an Apache Tomcat server Web application. After the application is launched on the Web server, you can connect to the QConvergeConsole's GUI through a browser, either locally on the server or remotely from another computer. Your browser window becomes the client used to connect to Dell PowerEdge blade servers that host the adapters and connected storage devices within the network.

In addition to the configuration and management tools available through the QConvergeConsole GUI, QConvergeConsole enables you to partition and configure NIC ports and eSwitch parameters on the adapters.

Follow the procedures in "Installing the QConvergeConsole GUI" on page 237 to install the QConvergeConsole GUI application on a Windows or Linux server. Before using the QConvergeConsole GUI to configure Switch Independent Partitioning on your adapters, you must also install the drivers on the Dell PowerEdge blade server where the adapter resides.

Before configuring Switch Independent Partitioning, do the following:

- Use the QLogic SuperInstaller for your host server's operating system (Windows or Linux) to install the Fibre Channel/FCoE, NIC, and iSCSI drivers on the server where the adapters reside. To download the installers and drivers, go to http://support.dell.com.

- Make sure the remote agents are running on the Tomcat server where the QConvergeConsole GUI application resides:

  - ❑ Fibre Channel/FCoE (`qlremote`)
  - ❑ NIC (`netqlremote`)
  - ❑ iSCSI (`iqlremote`)

QConvergeConsole CLI

QConvergeConsole CLI is a management utility that centralizes management and configuration of QLogic adapters within the entire network (LAN and SAN).

The QConvergeConsole CLI manages iSCSI, Ethernet, and FCoE functions on QLogic adapters installed on a Dell PowerEdge blade server on either a Linux or Windows environment.

In addition to the configuration and management capabilities available through the QConvergeConsole CLI, the QConvergeConsole CLI enables you to partition and configure NIC ports and eSwitch parameters on QLogic adapters.

Follow the procedures for your operating system in the *QConvergeConsole CLI User's Guide* to install the application on the host server.

For command references needed while using the QConvergeConsole CLI, refer to the *QConvergeConsole CLI User's Guide.*

Before using the QConvergeConsole CLI to configure Switch Independent Partitioning on your adapters, you must install the OS-specific drivers on the Dell PowerEdge server where the adapter resides. You can use the QLogic SuperInstaller for your host server's operating system (Windows or Linux) to install the Fibre Channel/FCoE, NIC, and iSCSI drivers. To download the installers and drivers, go to http://support.dell.com.

For procedures on setting up Switch Independent Partitioning and eSwitch parameters using the QConvergeConsole CLI, see "QConvergeConsole CLI" on page 153.

Windows Device Manager

Dell PowerEdge servers that run on supported Windows operating systems have the Windows-based tools available for configuring QLogic adapters. This enables you to use the Windows Device Manager to set up and manage NIC partitions.

For system requirements, see "Software Requirements" on page 123.

For procedures on using this native server management tool on a Windows server to configure Switch Independent Partitioning and on QLogic adapter NIC ports, see "Windows Device Manager" on page 158.

CIM Provider and QConvergeConsole VMware vCenter Server Plug-in for VMware ESX/ESXi

VMware vSphere provides a CIM monitoring framework for both classic ESX and ESXi. Use this framework to configure and manage Ethernet interfaces. On an ESX Server, you can configure and manage Ethernet interfaces using a command line interface. On an ESXi Server, use the remote CLIs.

All Switch Independent Partitioning Ethernet functions are enumerated by the hypervisor, controlled by the driver running in the hypervisor, and configured similar to other Ethernet interfaces.

You would typically create a vSwitch for each Switch Independent Partitioning interface. You can configure VMs to use the standard virtual network devices, such as VMXNET 3 adapters. On each interface, you can configure features such as NetQueue. Although it is possible to configure all standard aggregation and failover configurations supported by ESX using Switch Independent Partitioning interfaces as uplinks, it is not useful to have multiple uplinks with NIC partitions belonging to the same physical port.

The vSwitch switches the packets when VMs are assigned to share the same NIC partition. The eSwitch on the adapter switches packets when packets are sent by a VM on a NIC partition destined for another VM that is connected to a vSwitch on a different NIC partition belonging to the same physical port. If the two NIC partitions are associated with different physical ports, the external switch forwards packets between them.

Perform vSwitch configuration using either standard service console commands or RCLI commands. For specific command options, refer to the VMware documentation.

NOTE

Because all of the interfaces are assigned to the hypervisor, eSwitch operation on the adapter is transparent, and the system administrator is not required to perform any additional configuration.

To download the QLogic VMware ESX/ESXi vCenter Plug-in and vSphere Web Client Plug-in packages, go to http://driverdownloads.qlogic.com and select your 3200 Series Adapter or 8200 Series Adapter model for the VMware ESX/ESXi OS. For system requirements and installation procedures, see the associated Read Me document and user's guide provided with the package.

For procedures on using the vCenter Server Plug-in to configure Switch Independent Partitioning on QLogic adapter NIC ports, see "CIM Provider and vCenter Server Plug-in for VMware ESX/ESXi" on page 167.

Switch Independent Partitioning Setup and Management Options

This section describes how to configure NIC partitions on QLogic adapters installed in a Dell PowerEdge server (*host server*) within a SAN. Procedures for establishing QoS for each partition and viewing the eSwitch parameters and statistics are included.

This section provides setup procedures using the following management tools:

- Dell System Setup
- QLogic OptionROM at POST
- QConvergeConsole GUI
- QConvergeConsole CLI
- Windows Device Manager
- CIM Provider and vCenter Server Plug-in for VMware ESX/ESXi

> **NOTE**
>
> These procedures assume you have either local or remote access to a host server with at least one installed QLogic adapter, as well as the necessary drivers and management tools.

Overview

Depending on your operating environment and preferred system management techniques, you can use any of the tools described in this section to set up NIC partitions on QLogic adapter ports to meet your system's networking requirements.

When you first start the Dell PowerEdge server that contains QLogic adapters, the POST starts. Running POST gives you access to two of the configuration tools you can use to set up NIC partitions on QLogic adapter ports: Dell System Setup and the QLogic OptionROM.

The QConvergeConsole GUI and CLI tools work on both Linux and Windows servers. If you prefer using a browser-based GUI interface, you can use the QConvergeConsole GUI to partition Ethernet ports into NIC, FCoE, or iSCSI partitions and establish QoS by adjusting the bandwidth settings. As an alternative, you can use the QConvergeConsole CLI to set up partitions using a command line interface in either interactive or non-interactive mode.

> **NOTE**
>
> When bandwidth settings exist for both Switch Independent Partitioning and DCBX, *DCBX takes precedence over Switch Independent Partitioning*. DCBX sets the bandwidth for iSCSI and NIC traffic, and then Switch Independent Partitioning sets the bandwidth for the NIC partitions by dividing the NIC bandwidth allocated by DCBX. For more information, see "Interoperation of Bandwidth Settings for DCBX and Switch Independent Partitioning" on page 120.

On Windows host servers, you can use the Windows Device Manager to set up and manage NIC partitions.

For CIM Provider VMware host servers, you can use CIM Provider and vCenter Server Plug-in for VMware ESX/ESXi.

Dell System Setup

When you first start the host server that contains QLogic adapters, the POST starts. Running the POST gives you access to the Dell System Setup.

To set up Switch Independent Partitioning using the Dell System Setup:

> **NOTE**
>
> Depending on your server model and System Setup version, the screens you see might differ from those shown.

1. While running POST, press F2.

   The Main menu for the Dell System Setup opens.

2.   Select **Device Settings**, as shown in Figure 4-3.
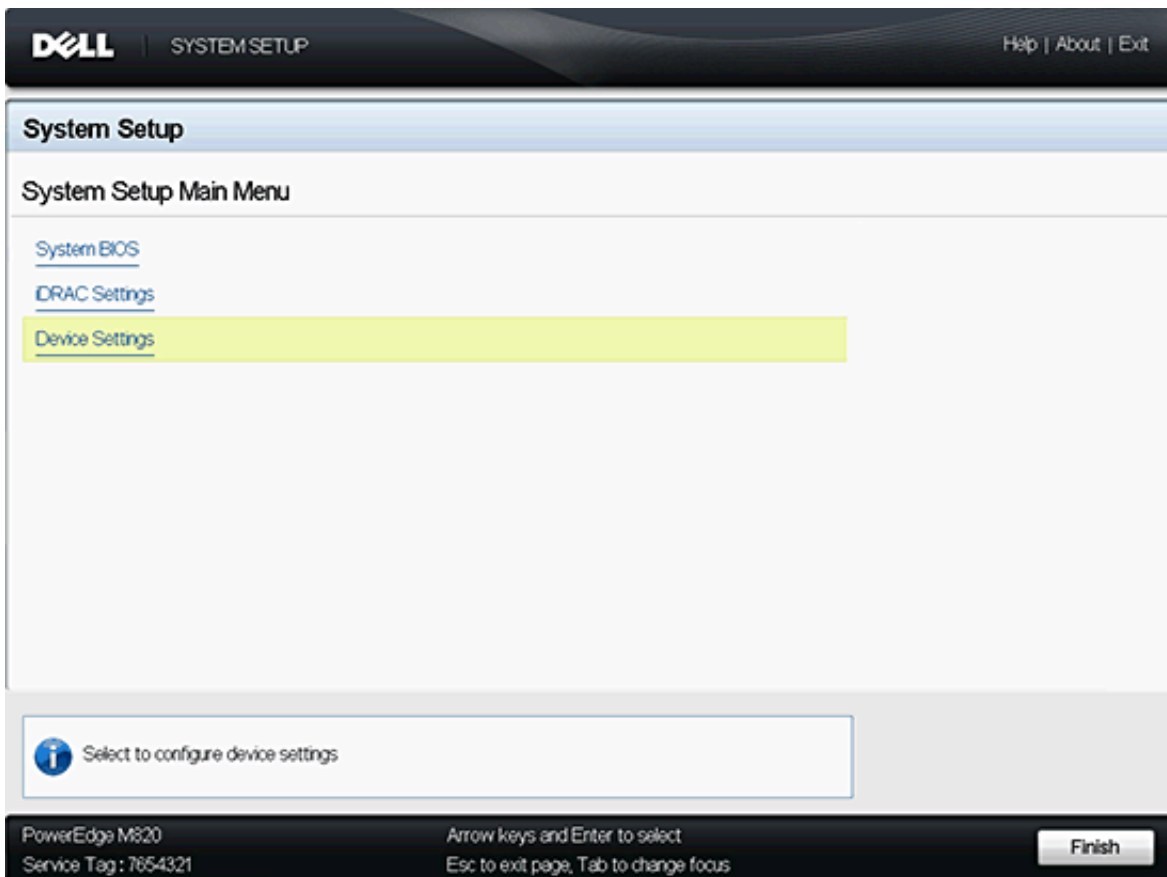
Figure 4-3. Main Menu Page

3. In the Device Settings screen, select the adapter that you want to configure (see Figure 4-4).
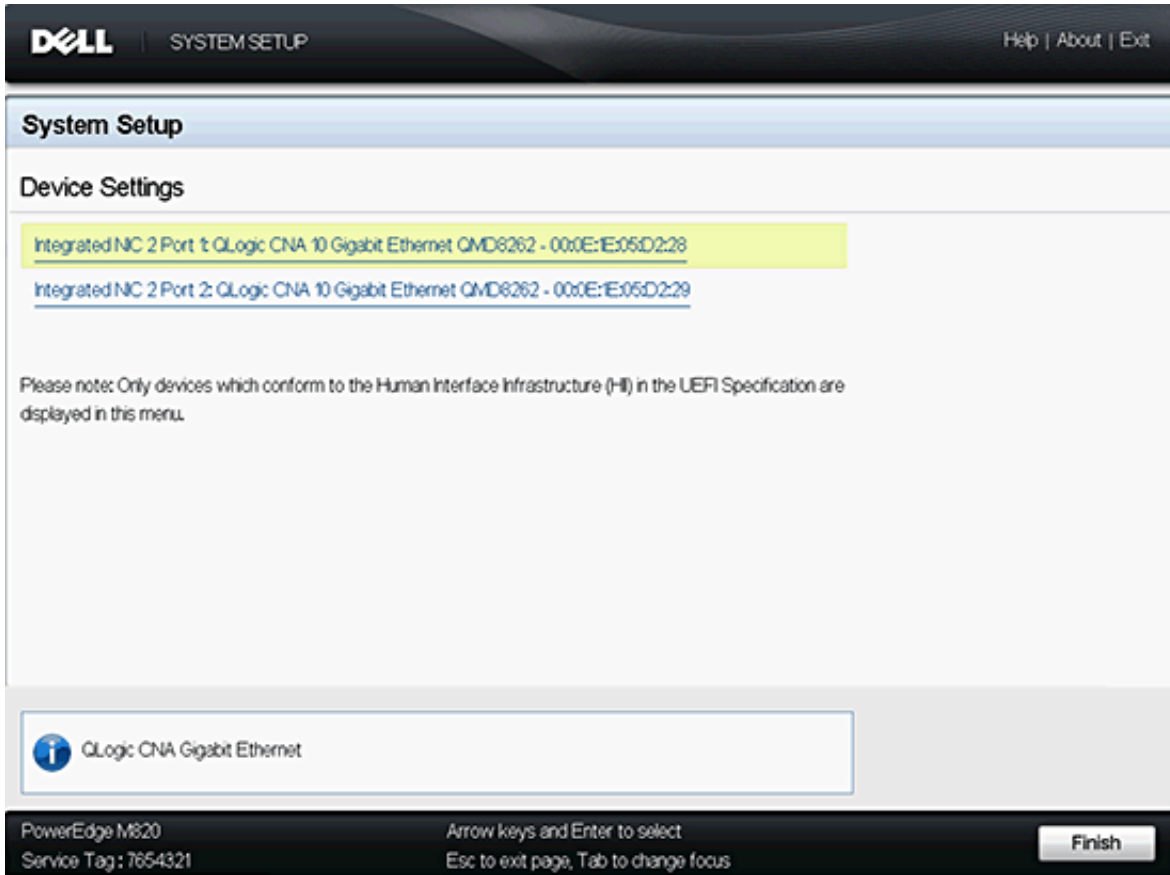
Figure 4-4. Selecting the Device to Configure

The next screen that appears (see Figure 4-5) is the Main Configuration page, which lists information about the selected adapter and the available setup options for the adapter.
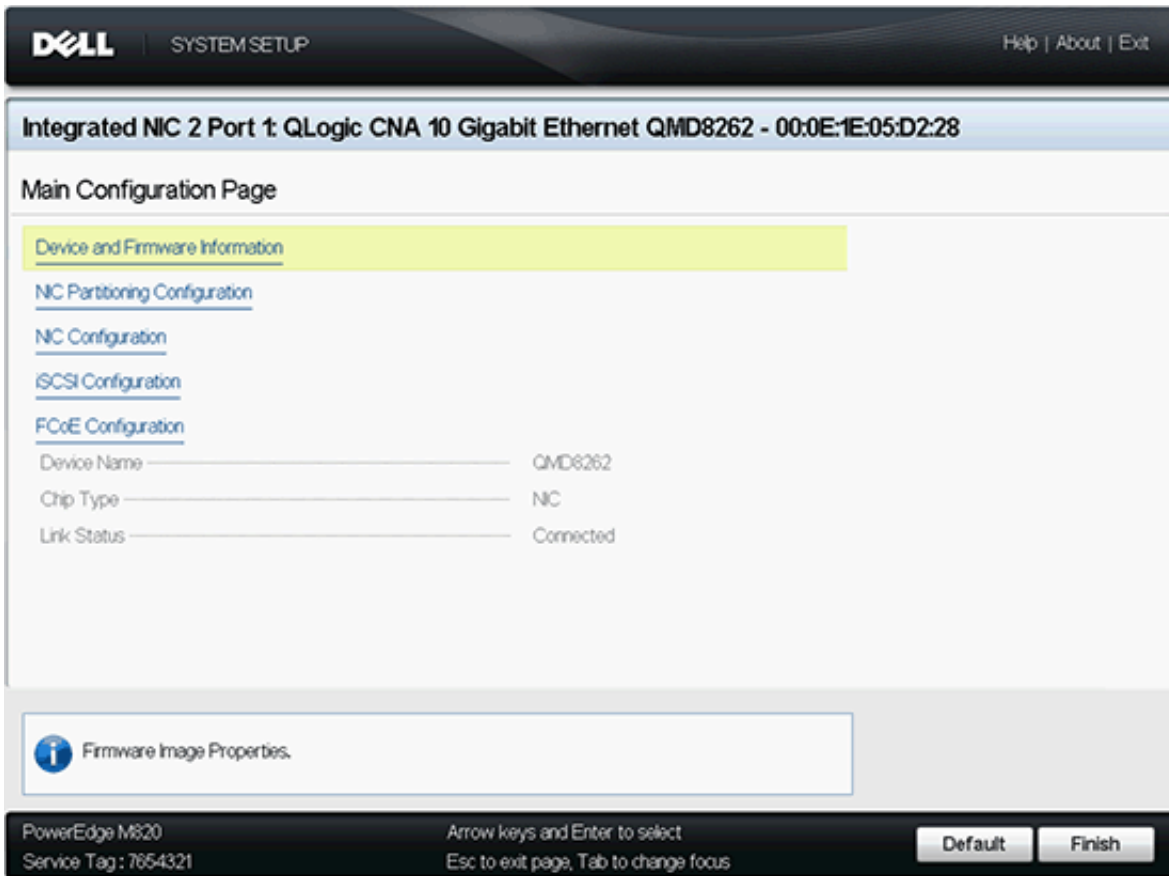
Figure 4-5. Main Configuration

4. Select **NIC Partitioning** (Switch Independent Partitioning) **Configuration** from the Main Configuration page.
   The NIC Partitioning Configuration page opens (see).
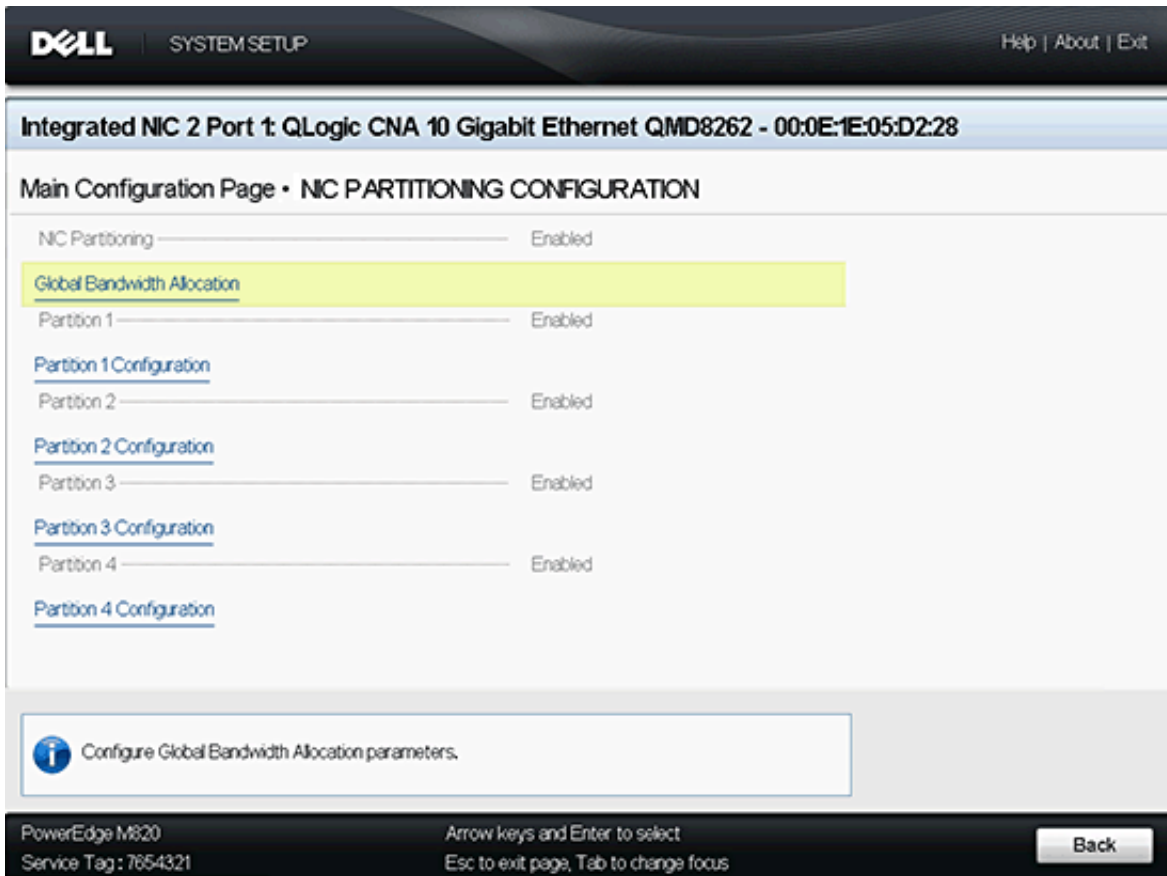
Figure 4-6. NIC Partitioning (Switch Independent Partitioning) Configuration Page

NOTE

For a list of Switch Independent Partitioning configuration options, see .

5. Select **Global Bandwidth Allocation** to open the Global Bandwidth Allocation page (see Figure 4-7).

Figure 4-7. Global Bandwidth Allocation Page

6. Set the relative and maximum bandwidth (between `0-100` percent) as needed for each partition.

❑ The relative bandwidth setting guarantees that at least this much bandwidth is available to the partition.

■ The cumulative relative bandwidth settings per port (partitions 1–4) must not exceed 100 percent.

■ This is software enforced.

❑ The maximum bandwidth setting enforces a bandwidth ceiling over which the partition cannot subscribe.

■ The adapter supports oversubscription.

■ Oversubscription allows for the sum total of the maximum bandwidth settings to equal greater than 100 percent of the total available bandwidth of the port.

■ Setting a port's maximum bandwidth to 100 percent allows that partition to use bandwidth that is not used by other partitions. This would apply if one or more of the other partitions were using less than their relative bandwidth setting.

■ If a partition's maximum setting is less than 100 percent, it is limited to that percentage, and the unused bandwidth of other partitions is not available to it if it hits its maximum bandwidth limit.

■ Oversubscription can be especially valuable in an environment with a mixture of NIC and storage and vMotion® traffic. It might be advantageous to have the NIC partition used by vMotion consume all unused bandwidth. Setting the vMotion partition to a maximum bandwidth of 100 percent ensures that the vMotion port dynamically consumes all bandwidth not being used by the other functions.

> **NOTE**
>
> When bandwidth settings exist for both Switch Independent Partitioning and DCBX, *DCBX takes precedence over Switch Independent Partitioning*. DCBX sets the bandwidth for iSCSI and NIC traffic, and then Switch Independent Partitioning sets the bandwidth for the NIC partitions by dividing the NIC bandwidth allocated by DCBX. For more information, see "Interoperation of Bandwidth Settings for DCBX and Switch Independent Partitioning" on page 120.

7. Return to the NIC Partitioning Configuration page.

8. For each partition, use the **Partition** drop-down list to enable, disable, or change the function type. You can only enable one function type on each partition at a time.

> **NOTE**
>
> Enabling one function type automatically disables the other type (or mode).
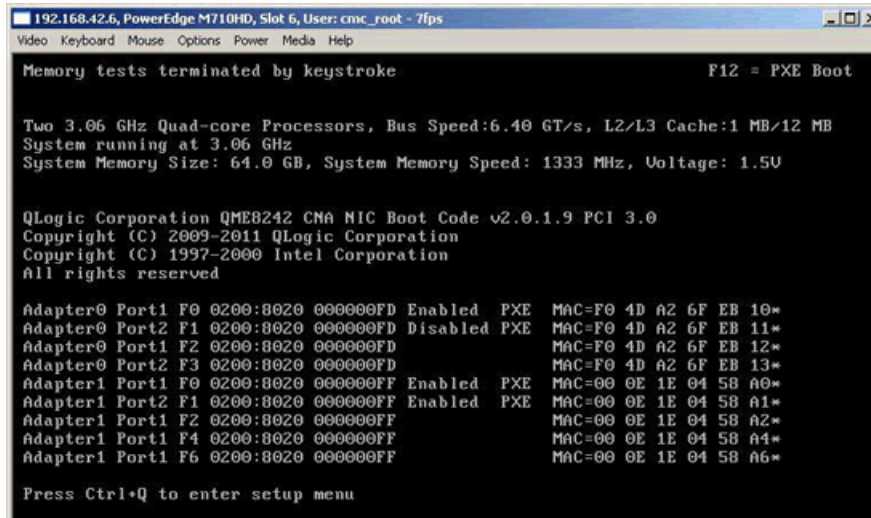
QLogic OptionROM at POST

When you first start host server containing QLogic adapters, the POST starts. Running the POST gives you access to the OptionROM utility.

To set up Switch Independent Partitioning using OptionROM:

1.   When the screen prompts you to enter the setup menu (see Figure 4-8) during the POST, press CTRL+Q to enter OptionROM setup.

Figure 4-8. POST Test Screen Prompt to Enter Setup Menu

2.   Select the adapter you want to manage on the QLogic CNA Function Configuration screen (see Figure 4-9).
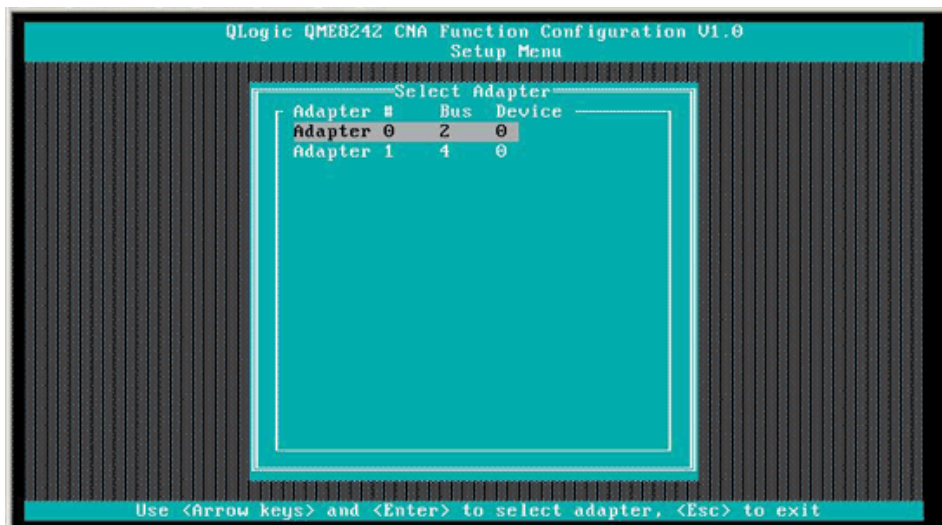
Figure 4-9. Selecting Adapter from Setup Menu

The screen displays a list of functions available to the selected adapter (see Figure 4-10).
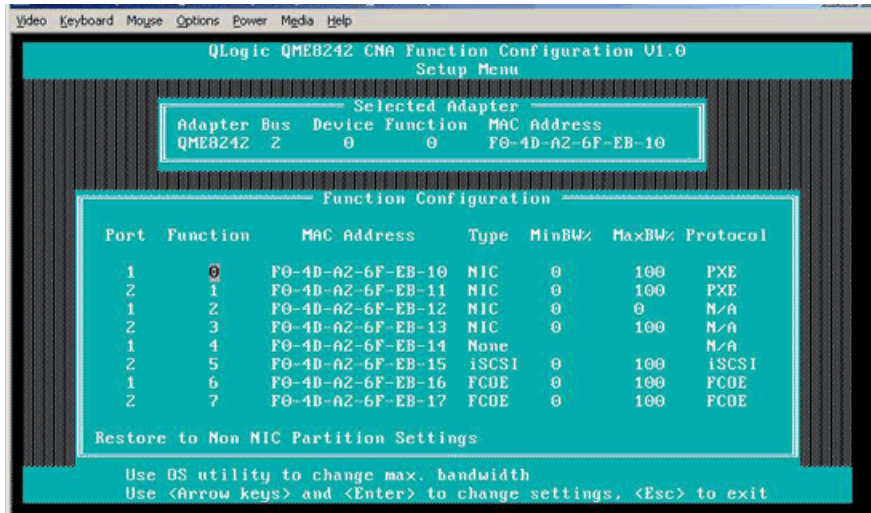
Figure 4-10. Function Configuration Screen

NOTE

For a list of Switch Independent Partitioning configuration options, see "Switch Independent Partitioning Setup" on page 171.

---

3. Move your cursor to the `Type` column for any function type you want to change (see Figure 4-11 and Figure 4-12).
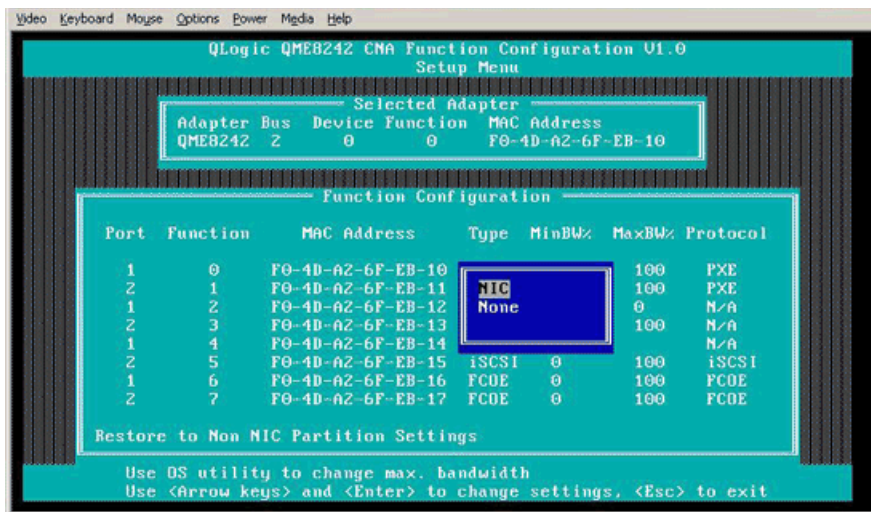
Figure 4-11. Selecting NIC Function Type to Change
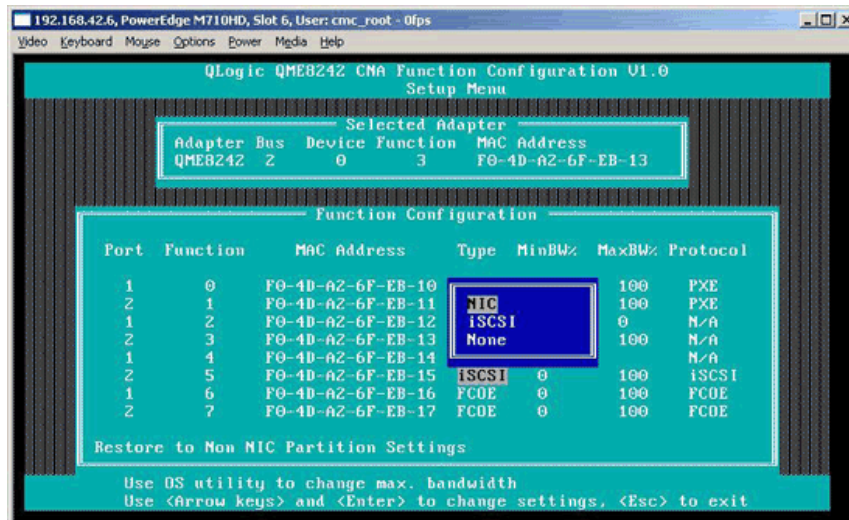
Figure 4-12. Selecting FCOE Function Type to Change

4. Move your cursor to the `MinBW%` column to adjust the minimum bandwidth (see Figure 4-13) on each partition (between `0`–`100` **percent**).
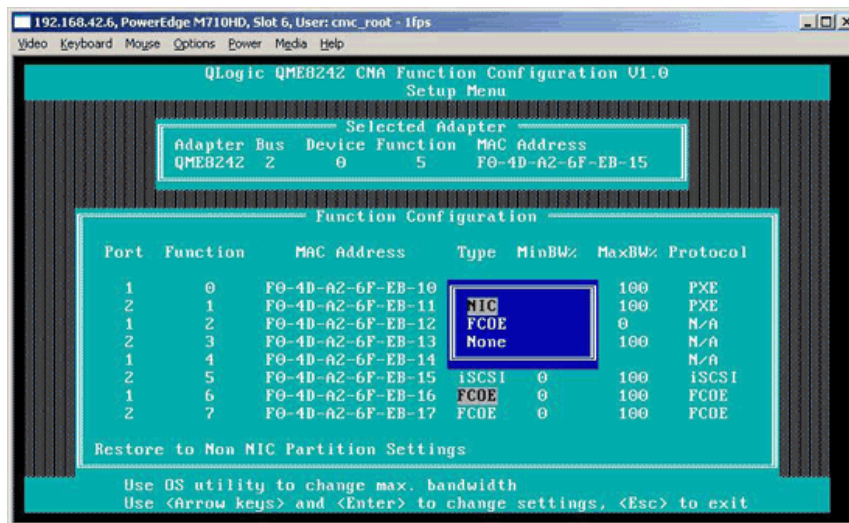
Figure 4-13. Adjusting the Minimum Bandwidth

NOTE

- The minimum bandwidth settings in the OptionROM are equivalent to the relative bandwidth settings in the Dell System Setup.

- The MaxBW% field is read only in this utility. To adjust the maximum bandwidth, use a different utility, such as the Dell System Setup.

- When bandwidth settings exist for both Switch Independent Partitioning and DCBX, *DCBX takes precedence over Switch Independent Partitioning*. DCBX sets the bandwidth for iSCSI and NIC traffic, and then Switch Independent Partitioning sets the bandwidth for the NIC partitions by dividing the NIC bandwidth allocated by DCBX. For more information, see "Interoperation of Bandwidth Settings for DCBX and Switch Independent Partitioning" on page 120.

---

5.   Save your changes (see Figure 4-14).
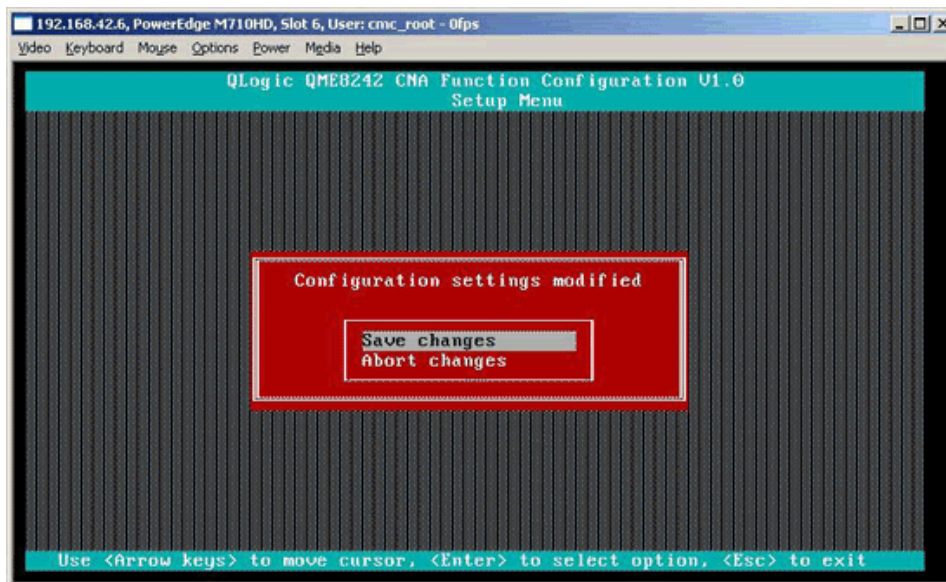
Figure 4-14. Saving Configuration Changes

6.   Reboot the host server after completing the Switch Independent Partitioning configuration.

QConvergeConsole GUI

The QConvergeConsole is a Web-based client/server application that allows for centralized management and configuration of QLogic adapters within the entire network (LAN and SAN). On the server side, QConvergeConsole runs as an Apache Tomcat server Web application. After the application is launched on the Web server, you can connect to the QConvergeConsole's GUI through a browser, either locally on the server or remotely from another computer. Your browser window becomes the client used to connect to servers that host the QLogic adapters and connected storage devices within the network.

You can use the QConvergeConsole GUI tool to configure and manage QLogic adapters installed on either Linux or Windows host servers.

For procedures on installing and starting this management tool, refer to "Installing the QConvergeConsole GUI" on page 237. For help configuring and managing the adapter using this management tool, refer to the QConvergeConsole online help system.

To set up NIC partitions using the QConvergeConsole GUI:

1. Configure NIC Partitions
2. Set Up QoS
3. View eSwitch Configuration

Configure NIC Partitions

You can use QConvergeConsole to configure and manage Switch Independent Partitioning functions for both physical ports through the NIC Partitioning tab. You can enable or disable Switch Independent Partitioning functions on either physical port and must reboot the operating system to apply the changes. When the Switch Independent Partitioning function is enabled, each physical port divides its bandwidth function among four physical functions or physical PCIe functions, configured to support one of the following function types: NIC, FCoE, or iSCSI. QConvergeConsole represents each function type as a personality.

> **NOTE**
>
> For tables that show the default Switch Independent Partitioning function settings, as well as the possible configurations, see "Switch Independent Partitioning Setup" on page 171.

To configure the NIC partitions and change personalities:

1. Expand an adapter node in the QConvergeConsole system tree.

2. Expand the physical Port 1 node and select the NIC port.

   The content pane displays two additional tabs that are not available on NIC ports for physical Port 2.

3.   Select the **NIC Partitioning** (Switch Independent Partitioning) tab. The NIC Partitioning Configuration page displays configuration details that apply to the selected Switch Independent Partitioning configuration and personality options (see Figure 4-15).

Figure 4-15. NIC Partitioning (Switch Independent Partitioning) Configuration Page

4.   Select the physical port you want to configure from the **Physical Port** drop-down list.

5.   If you want to change its function type, select the NIC partition and select the appropriate protocol from the **Function Type** drop-down list.

6.   Click **Save** to save any changes.

The Security Check dialog box might appear. In the **Enter Password** box, type the password and then click **OK**.

7.   Reboot the operating system to apply the changes.

8. Verify that the configured ports have the most current drivers installed.

9. If necessary, update the driver for the port protocol.

Set Up QoS

The QConvergeConsole lets you set the QoS for each partition by setting minimum and maximum percentages of the physical port's bandwidth for each partition.

NOTE

The NIC Partitioning page applies to NIC ports only for Switch Independent Partitioning-enabled QLogic adapters.

To set the QoS:

1. Expand a QLogic adapter node in the QConvergeConsole system tree.

2. Expand the physical Port 1 node and select the NIC port.

   The content pane displays two additional tabs that are not available on NIC ports for physical Port 2.

3. Select the **NIC Partitioning** tab and then click the **Management** subtab.

   The NIC Partitioning Management General page displays configuration details that apply to the selected NIC partition (see Figure 4-16).

Figure 4-16. NIC Partitioning—General Management Page

4. Click the down arrow and select the NIC partition (**NPAR0**, **NPAR1**, **NPAR2**, or **NPAR3**) from the drop-down list.

Information and configuration fields related to the selected NIC partition include:

❑ **Default MAC Address**—The MAC address set at the manufacturer.

❑ **Location**—The logical location in the system: PCI bus number, device number, and function number.

❑ **NPAR PCI Function Number**—The function number (**0**, **1**, **2**, **3**, **4**, **5**, **6**, or **7**) of the eight PCIe function numbers claimed by the adapter.

❑ **NPAR Function Type**—This field correlates to the personality of the selected Switch Independent Partitioning (PCIe) function: **NIC**, **iSCSI**, or **FCoE**.

❑ **Minimum Bandwidth (%)**—Use the up and down arrows to scroll between **0** percent to **100** percent to set the bandwidth you want to guarantee for data sent over the selected partition. Each additional percent increments the bandwidth by 100Mbps. For example, setting the minimum bandwidth to five percent guarantees sending and receiving data over the selected port at 500Mbps.

❑ **Maximum Bandwidth (%)**—The maximum allowed bandwidth is specified as a percentage of the link speed. Use the up and down arrows to scroll between **0** percent to **100** percent to set the maximum bandwidth for data sent over the selected partition. Each additional percent increments the bandwidth by 100Mbps. For example, setting the maximum bandwidth to 100 percent allows for sending and receiving data over the selected partition at up to 10,000Mbps.

> **NOTE**
>
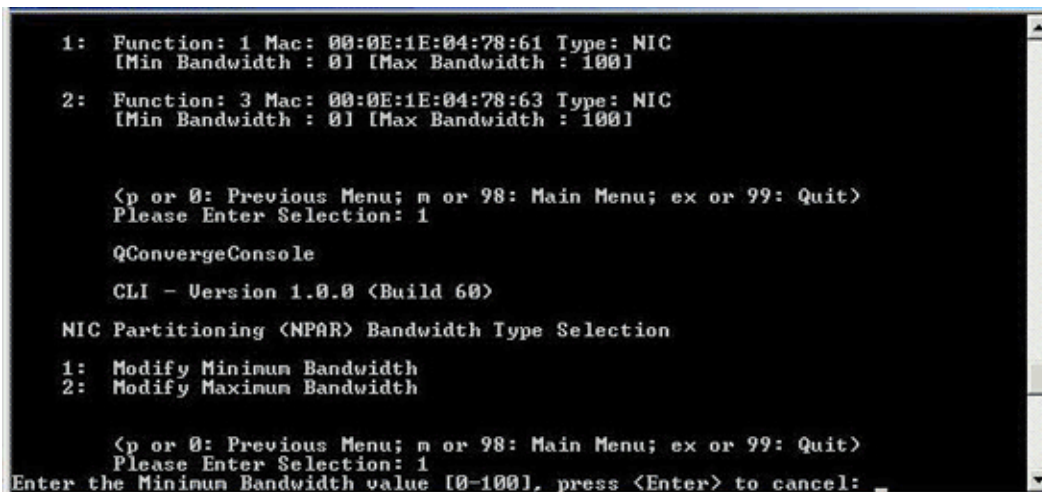> When bandwidth settings exist for both Switch Independent Partitioning and DCBX, *DCBX takes precedence over Switch Independent Partitioning*. DCBX sets the bandwidth for iSCSI and NIC traffic, and then Switch Independent Partitioning sets the bandwidth for the NIC partitions by dividing the NIC bandwidth allocated by DCBX. For more information, see "Interoperation of Bandwidth Settings for DCBX and Switch Independent Partitioning" on page 120.

5. Repeat the previous step to configure the minimum and maximum bandwidth on the other partitions.

6. When you are finished setting changes, click **Save** to save any changes to the advanced parameters of the adapter.

The Security Check dialog box might appear. In the **Enter Password** box, type the password and then click **OK**.

NOTE

The settings are persistent across reboots.

View eSwitch Configuration

QConvergeConsole appears and lets you view the current eSwitch offload settings. The eSwitch configuration is not persistent across reboots and driver reloads; however, you can save the configuration before the reboot/reload and restore it at a later time.

To view the eSwitch configuration:

1. Expand a QLogic adapter node in the QConvergeConsole system tree.
2. Expand the physical port 1 node and select the NIC port. The content pane displays tabs that are not available on NIC ports for physical port 2.
3. Select the **NIC Partitioning** tab.
4. Select the **eSwitch** subtab.

    The NIC Partitioning eSwitch Management page displays offload configuration settings that apply to the selected port (see Figure 4-17).
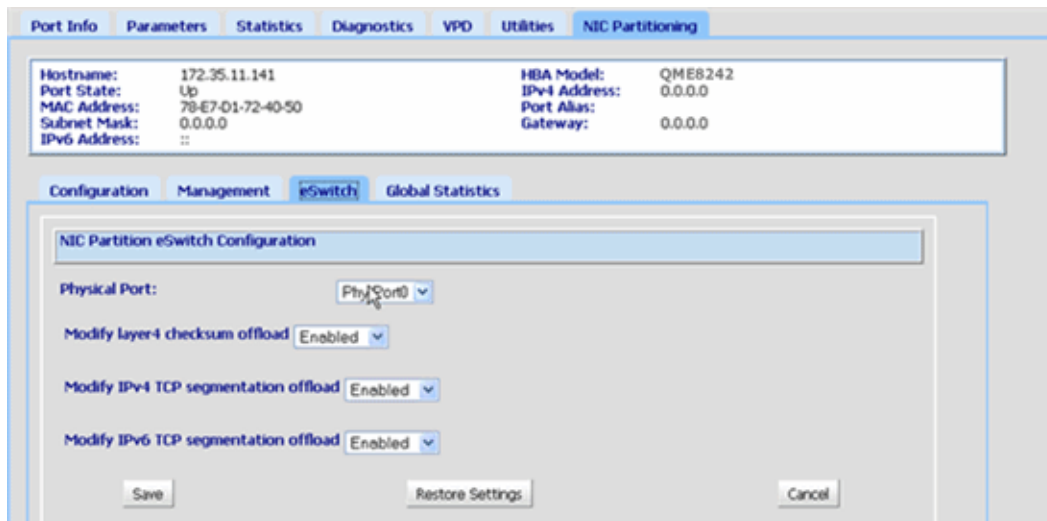
Figure 4-17. NIC Partitioning—eSwitch Management Page

5. Select the port for which you want to set the offload properties from the drop-down menu.

    The listed offload fields apply to the selected port.

6. Click the down arrow next to any of the offload fields provided to change its value to **Enabled** or **Disabled**.

7. Select one of the following command buttons to apply or cancel any changes:

❑ **Save**–Saves changes displayed on the screen.

❑ **Restore Settings**–Restores the default settings.

❑ **Cancel**–Cancels any changes made to this screen before you saved them.

QConvergeConsole CLI

The QConvergeConsole CLI is a management utility that centralizes management and configuration of QLogic adapters within the entire network (LAN and SAN).

You can use the QConvergeConsole CLI tool in either interactive or non-interactive mode to configure and manage QLogic adapters installed on either Linux or Windows host servers.

This section outlines the steps for setting up NIC partitions using the QConvergeConsole CLI in interactive mode. The displayed commands apply to both Linux and Windows operating systems.

> **NOTE**
>
> Throughout this section, the terms *NIC partitioning* and *NPAR* are synonymous with *Switch Independent Partitioning*, and *NIC partition* is synonymous with *switch-independent partition*.

For procedures on installing and starting this management tool, refer to the *QConvergeConsole CLI User's Guide*.

To set up NIC partitions using the QConvergeConsole CLI:

1. Start the QConvergeConsole CLI interface and select `6: NIC Partitioning <NPAR> Information` (see Figure 4-18).
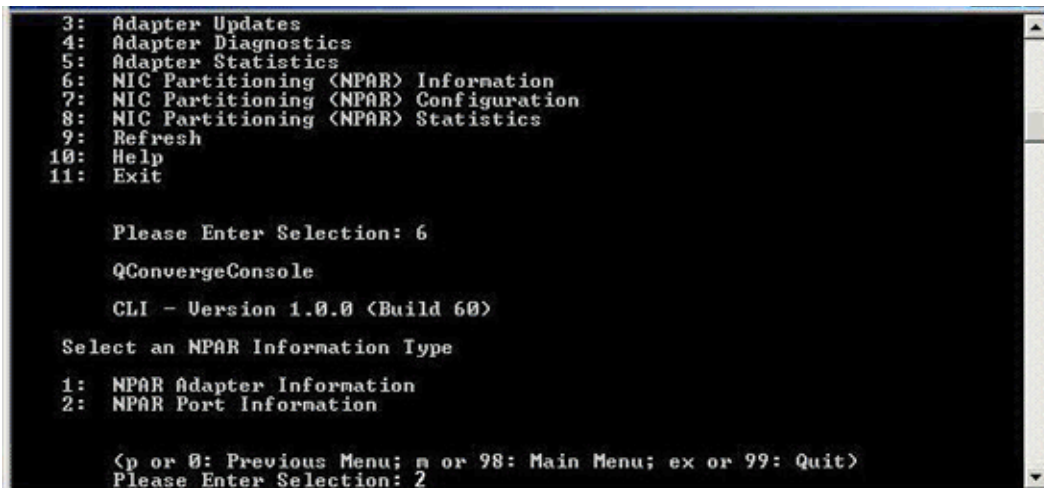
Figure 4-18. Selecting 6 to View NPAR Information Options

2. Select `2: NPAR Port Information` (see Figure 4-19).

Figure 4-19. Selecting 2 to View NPAR Port Information

The NPAR Configuration Selection page displays the current configuration (see Figure 4-20).

Figure 4-20. NPAR Configuration Selection Screen

3.   Return to the main menu after viewing the Switch Independent Partitioning information and select **7: NIC Partitioning <NPAR> Configuration** (see Figure 4-21).

Figure 4-21. Selecting NPAR Configuration

4.   Select **1: NPAR Configuration** to display the NPAR Configuration menu, which provides the following options:

1: Bandwidth Configuration
2: Change PCI Function Personality

5. Configure the bandwidth settings to meet your system requirements.

NOTE

When bandwidth settings exist for both Switch Independent Partitioning and DCBX, *DCBX takes precedence over Switch Independent Partitioning*. DCBX sets the bandwidth for iSCSI and NIC traffic, and then Switch Independent Partitioning sets the bandwidth for the NIC partitions by dividing the NIC bandwidth allocated by DCBX. For more information, see "Interoperation of Bandwidth Settings for DCBX and Switch Independent Partitioning" on page 120.

For example, to change the bandwidth of the function 1 NIC partition:

a. Select **1: Bandwidth Configuration**.

b. Select **1: Function:1**.

c. Select **1: Modify Minimum Bandwidth** (see Figure 4-22).

Figure 4-22. Selecting to Modify Minimum Bandwidth

d. At the prompt, enter the percent value of bandwidth you want committed to the selected function.

e. Enter the percent value of bandwidth you want to limit the selected function to.

f.	Specify whether you want your bandwidth settings to persist across reboots (see Figure 4-23).

Figure 4-23. Setting Bandwidth Changes to Persist

6.	Return to the NIC Partitioning <NPAR> Configuration Selection screen.

7.	Change the personalities of each function to meet your system requirements. For example:

a.	Select **2: Change PCI Function Personality.**

b.	Select the port number, **1** or **2**.

c.	Select the function number.

The command line displays a list of options with choices that apply to the selected function number. This mode prevents you from assigning a function type that does not apply to a given function number.

d.	Set the personality type by selecting the option number that identifies the appropriate function type. Depending on the function number and current state, this could be **Disabled**, **NIC**, **FCoE**, or **iSCSI**.

NOTE

For a list of Switch Independent Partitioning configuration options, see "Switch Independent Partitioning Setup" on page 171.

Figure 4-24 shows the CLI commands leading to the option for changing a function type on a Linux system.

```
CNA Model QME8242 SN: AFE1033C05073
  1. Port   1
     Function   6 [Type: FCoE ] [WWPN: 20-01-00-0E-1E-04-3A-AE] Online
     Function   4 [Type: iSCSI] [iSCSI Name: iqn.2000-04.com.qlogic:isp8214.000E1E043AAC.4] Link Down
     Function   0 [Type: NIC  ] [MAC Address: 00:0E:1E:04:3A:A8] Link Up
  2. Port   2
     Function   7 [Type: FCoE ] [WWPN: 20-01-00-0E-1E-04-3A-AF] Online
     Function   1 [Type: NIC  ] [MAC Address: 00:0E:1E:04:3A:A9] Link Up


    (p or 0: Previous Menu; m or 98: Main Menu; ex or 99: Quit)
    Please Enter Selection: 1

    QConvergeConsole

    CLI - Version 1.1.0 (Build 4)

NIC Partitioning (NPAR) PCI Function Configuration Selection

1:  Function: 2 MAC: 00:0E:1E:04:3A:AA Type: Disabled
2:  Function: 4 MAC: 00:0E:1E:04:3A:AC Type: iSCSI
3:  Function: 6 MAC: 00:0E:1E:04:3A:AE Type: FCoE


    (p or 0: Previous Menu; m or 98: Main Menu; ex or 99: Quit)
    Please Enter Selection: █
```

Figure 4-24. Selecting Function Type on Linux System

8. Return to the main menu and select **8: NIC Partitioning <NPAR> Statistics** to view the statistics.

   Navigate through the menu selections to view eSwitch statistics.

9. After you have finished setting the NIC partitions as needed, you must reboot the host server for the changes to take effect.

Windows Device Manager

On a Windows server that hosts QLogic adapters, you can use the Windows Device Manager to set up NIC partitions. You can also use it to view eSwitch statistics.

To set up Switch Independent Partitioning using the Windows Device Manager:

1. Configure Switch Independent Partitioning
2. Change Personalities
3. Manage Bandwidth
4. View eSwitch Statistics

Configure Switch Independent Partitioning

You can use the NIC Partition Management tab in the device properties page to enable Switch Independent Partitioning and configure the 10GbE physical port into a multifunction storage and networking port.

To set up Switch Independent Partitioning on a QLogic adapter port:

1. Log in to the server that contains installed QLogic adapters.

2. Open the Server Manager and select **Diagnostics**, then select **Device Manager**, and then select **Network Adapters**.

3. Right-click on the first **QLogic 10GbE 2P <*model*>** device and select **Properties** from the shortcut menu (see Figure 4-25).
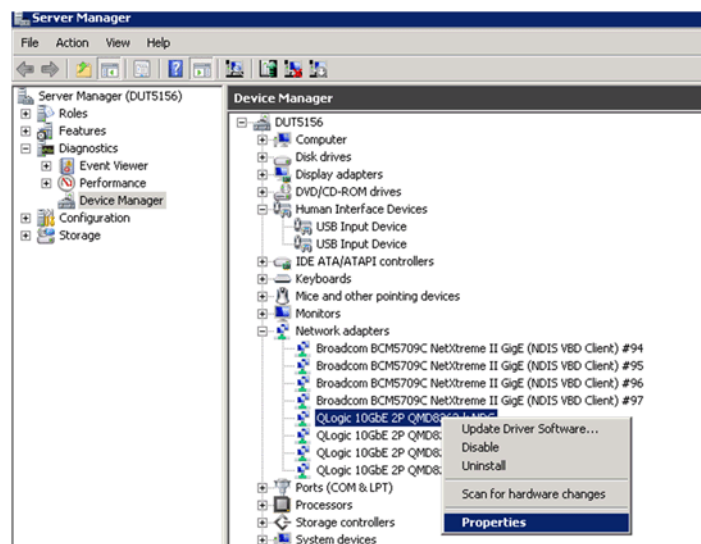
Figure 4-25. Selecting Properties from the Shortcut Menu

4. From the Adapter Properties page, do the following:

    a. Select the **NIC Partition Management** tab.

    b. Right-click on the function number you want to enable.

    c. Select **Enable Partition** (see Figure 4-26).
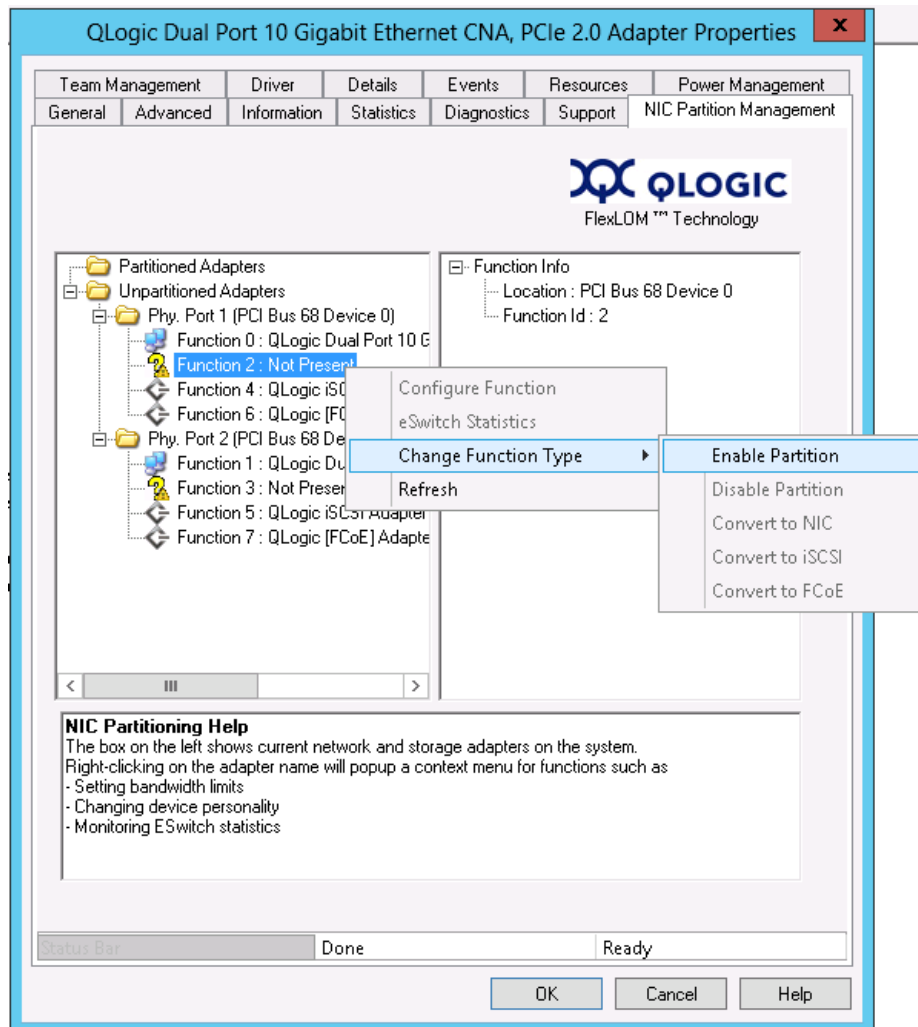
Figure 4-26. Enabling Partition

When partitioning is enabled, the Adapter Properties page appears, as shown in Figure 4-27.
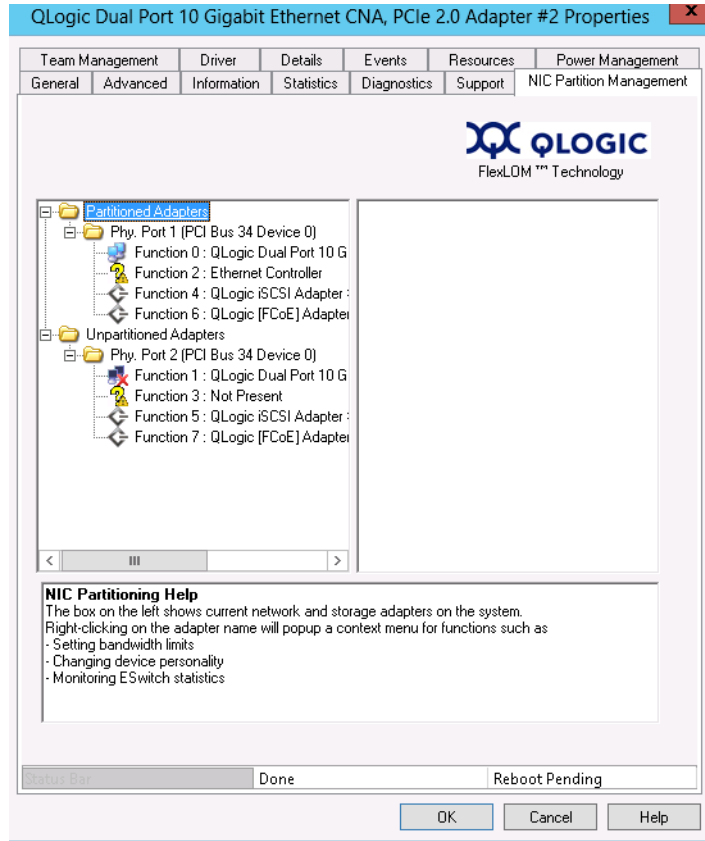
Figure 4-27. Partition Enabled

5.   Click **OK** to close the message box that displays the following information:

   This change requires a reboot. Proceed?

6.   Click **OK** to close the message box that displays the following information:

   Please reboot the system now

7.   Reboot the host server to make the changes take effect.

Change Personalities

To change function types (*personalities*) as needed for your network:

1.   From the Server Manager, select **Diagnostics**, then select **Device Manager**, and then select **Network Adapters**.

2.   Right-click on the appropriate **QLogic 10 Gigabit Ethernet CNA** device to change the function type and select **Properties** from the shortcut menu.

3. On the NIC Partition Management tab, right-click on one of the enabled
    functions, select **Change Function Type**, then select **Convert to
    <Protocol>** from the shortcut menu (see Figure 4-28).
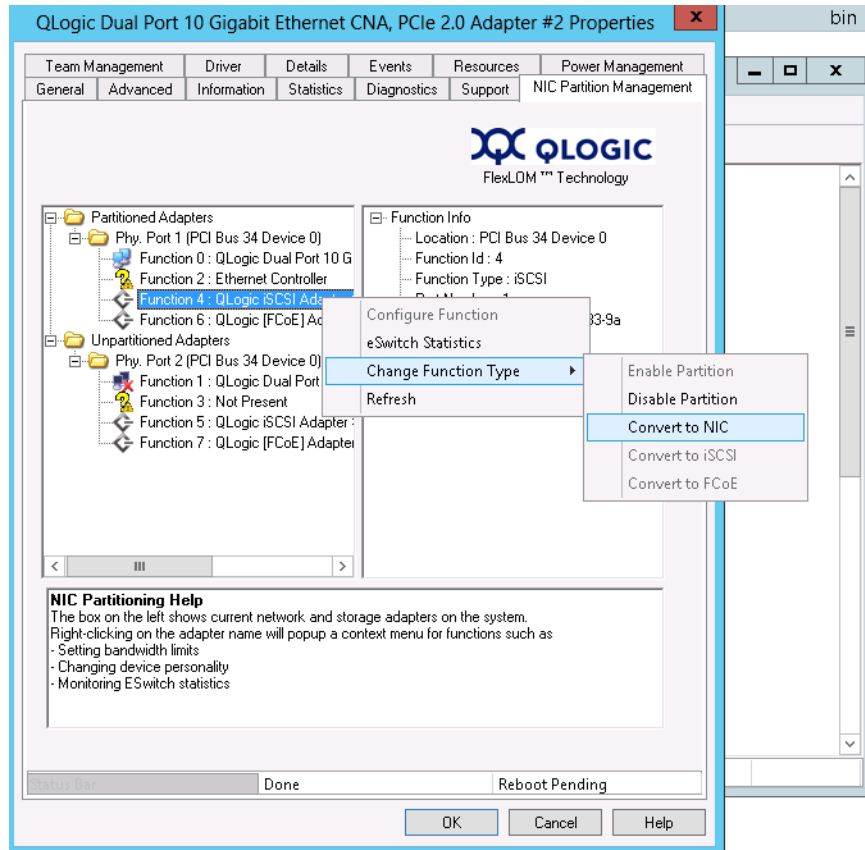
Figure 4-28. Selecting Convert to NIC from Shortcut Menu

4. Repeat these procedures to change the function types as needed.

Manage Bandwidth

Using the NIC Partition Management tab in the Windows device properties page, you can allocate minimum and maximum bandwidth for each NIC function.

> **NOTE**
>
> When bandwidth settings exist for both Switch Independent Partitioning and DCBX, *DCBX takes precedence over Switch Independent Partitioning*. DCBX sets the bandwidth for iSCSI and NIC traffic, and then Switch Independent Partitioning sets the bandwidth for the NIC partitions by dividing the NIC bandwidth allocated by DCBX. For more information, see "Interoperation of Bandwidth Settings for DCBX and Switch Independent Partitioning" on page 120.

1. From the QLogic Adapter Properties page, select the **NIC Partition Management** tab.

2. Right-click the function number for the port you want to configure and select **Configure Function** from the shortcut menu (see Figure 4-29).
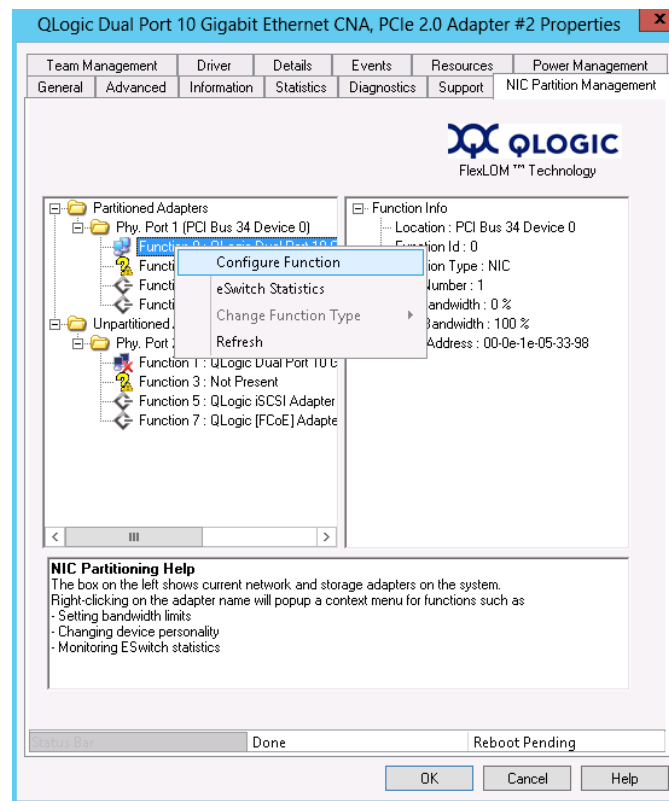
Figure 4-29. Selecting Configure Function for Function 0

3. Use the Configure Function dialog box to set the minimum and maximum bandwidth percentages, **New Minimum BW** and **New Maximum BW** (see Figure 4-30).
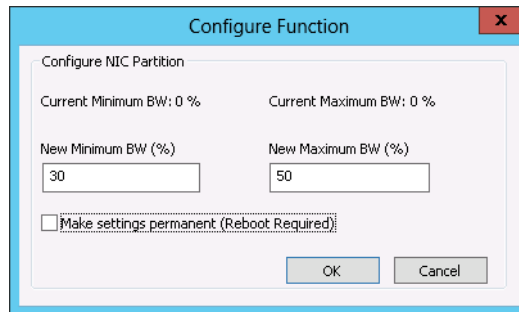
Figure 4-30. Entering New Bandwidth Values

> **NOTE**
>
> ETS only specifies the division of bandwidth between FCoE and non-FCoE traffic. It does not specify the bandwidth allocated to the NIC or iSCSI partitions. When the switch sets ETS values, the ETS bandwidth parameters take precedence. The FCoE partition is allocated the bandwidth specified for FCoE in the ETS parameters. The non-FCoE bandwidth is divided between the NIC and iSCSI partitions in the proportion specified by the Switch Independent Partitioning management user interface. In other words, when ETS is in effect, the NIC and iSCSI bandwidth values specified by the Switch Independent Partitioning management UI are no longer a percentage of the total bandwidth. Instead, they are a percentage of the non-FCoE bandwidth.

4. If needed, select the **Make settings permanent** check box to retain the new settings.

> **NOTE**
>
> If you do not select this option, the bandwidth values revert to the default settings after you reboot the host server.

5. Click **OK** to save your changes.

The new bandwidth values appear in the right pane of the NIC Partition
Management property sheet (see Figure 4-31).
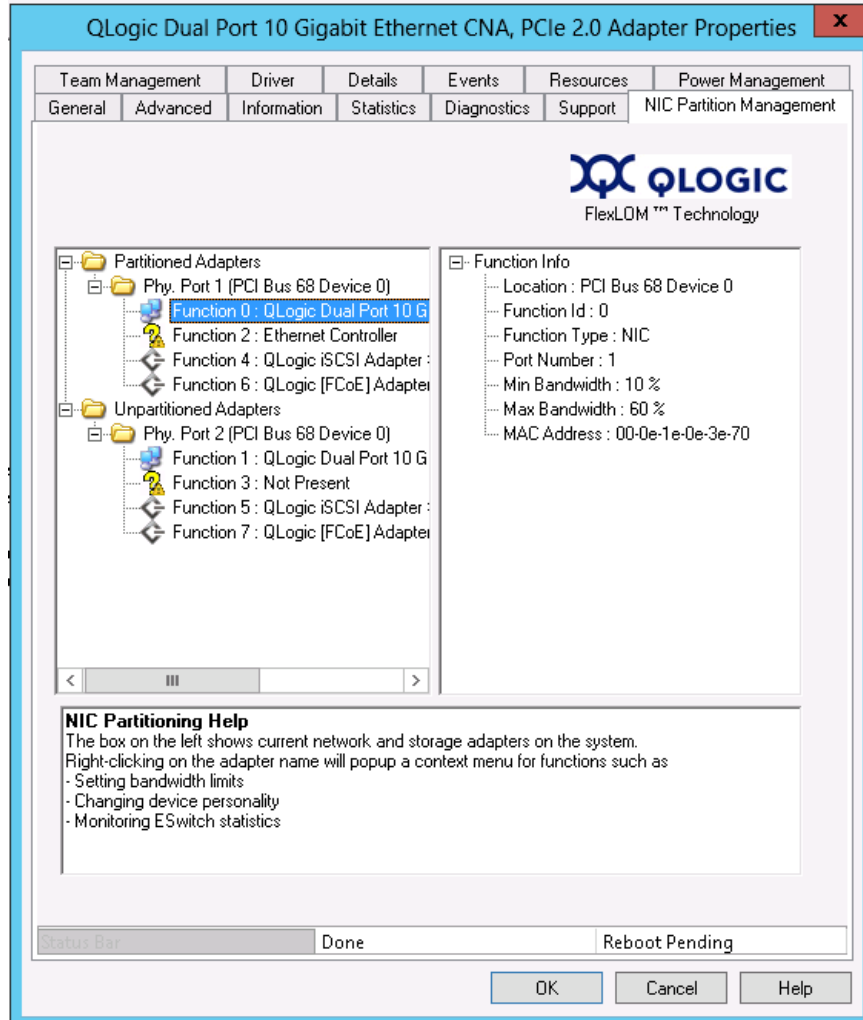
Figure 4-31. NIC Partition Management Property Sheet

6.   Click **OK** at the bottom of the Properties page to close it.

View eSwitch Statistics

You can use the Window Device Manager's NIC Partition Management window to view eSwitch statistics for enabled partitions.

To display eSwitch statistics:

1. From the QLogic Adapter Properties page, select the **NIC Partition Management** tab.

2. Right-click the function number for the port you want to review and select **eSwitch Statistics** from the shortcut menu.

   A pop-up window displays the statistics (see Figure 4-32).
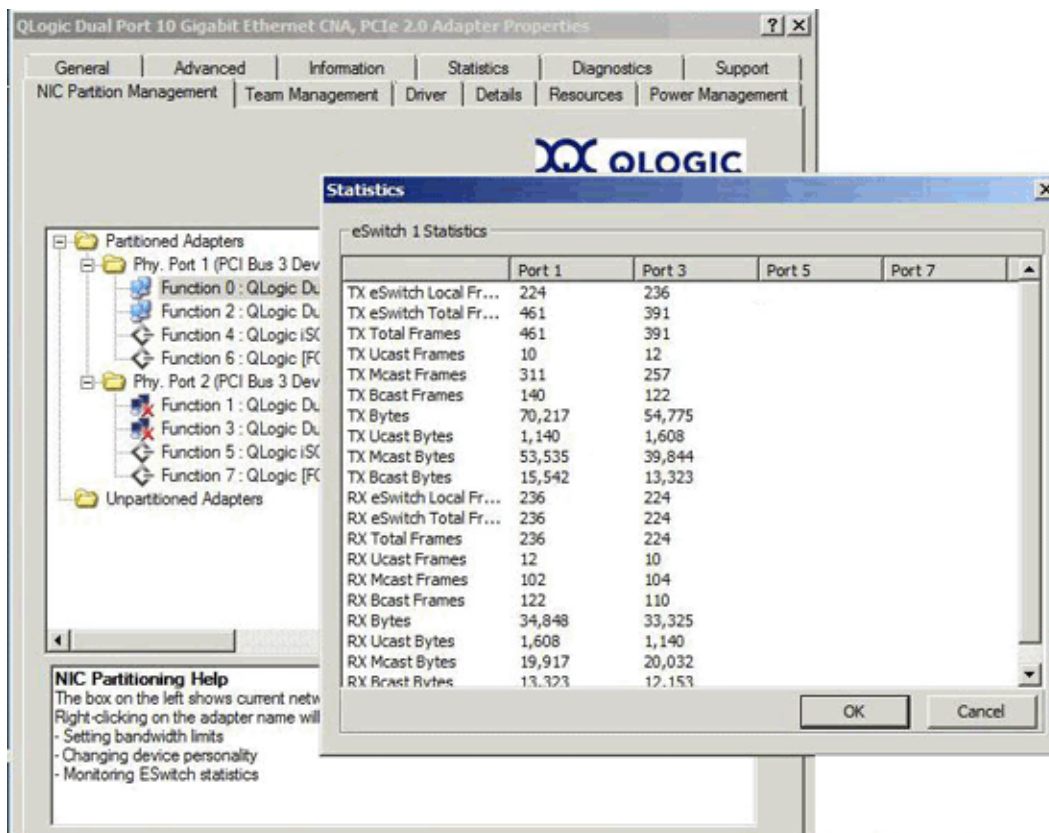
Figure 4-32. eSwitch Statistics for Function 0

3. After reviewing the statistics, click **OK** or **Cancel** to close the pop-up window.

CIM Provider and vCenter Server Plug-in for VMware ESX/ESXi

The QConvergeConsole vCenter Server Plug-in provides a QConvergeConsole tab you can use to manage the QLogic adapter in the VMware ESX/ESXi environment.

Use the QConvergeConsole tab to do the following:

- Enable or disable Switch Independent Partitioning
- Enable or disable functions
- Configure function personalities
- Configure the relative bandwidth of NIC functions
- View the eSwitch statistics

The default configuration of a QLogic adapter is a NIC function, iSCSI function, and FCoE function. There is an additional NIC function that is disabled by default but that you can enable as an option.

To configure Switch Independent Partitioning (called NPAR or NIC partitioning by the plug-in), select a NIC port in the system tree and click **NPAR** in the content pane. The content pane appears as shown in Figure 4-33.
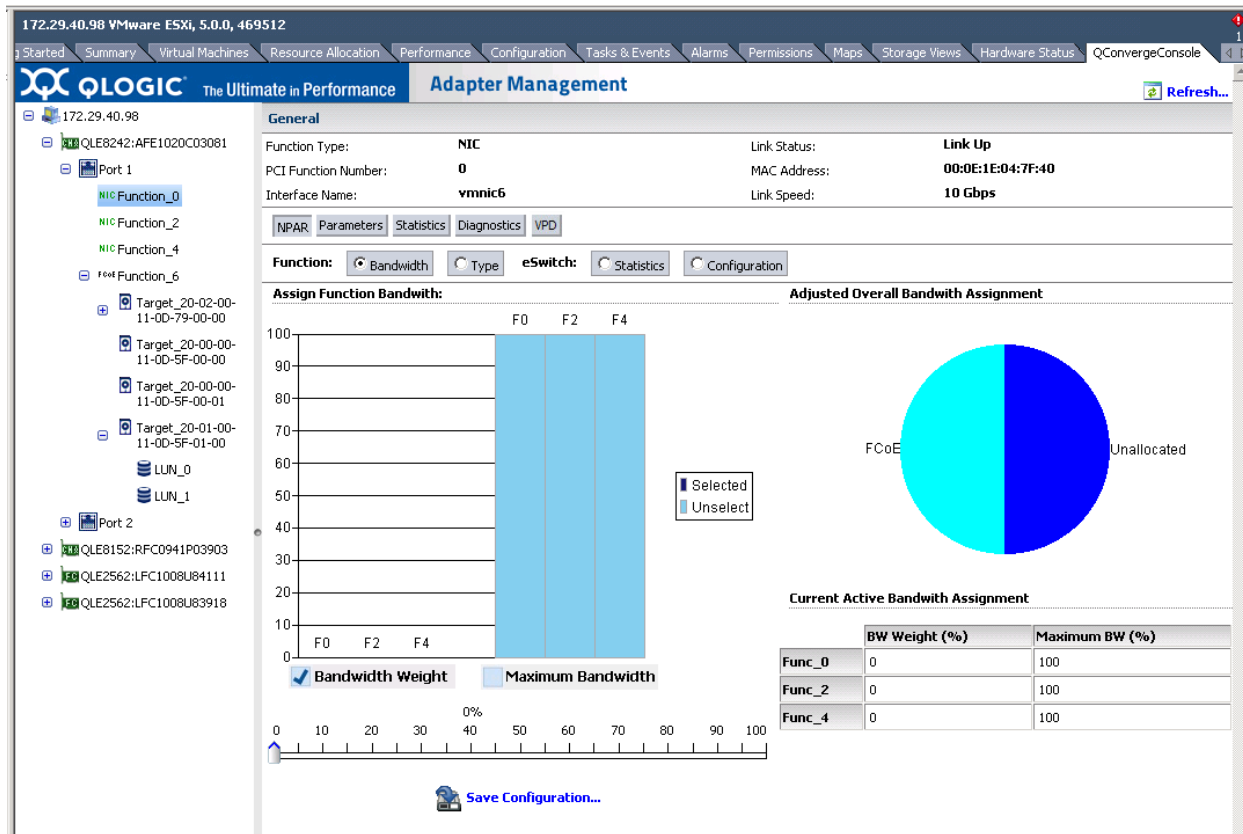
Figure 4-33. Converged Network Adapter NIC Function—NPAR Configuration

The content pane varies depending on which **Function** is selected:

- **Bandwidth:** This setting allows you to display and set the bandwidth allocation for the NIC function. For detailed information, refer to "Bandwidth Allocation" on page 168.

- **Type:** This setting displays the current function type and allows you to change the function type. For detailed information refer to "Function Type" on page 169.

Bandwidth Allocation

Selecting the **Bandwidth** option displays the current bandwidth allocation settings in the content pane, and allows you to change the settings as you require, as shown in Figure 4-34.
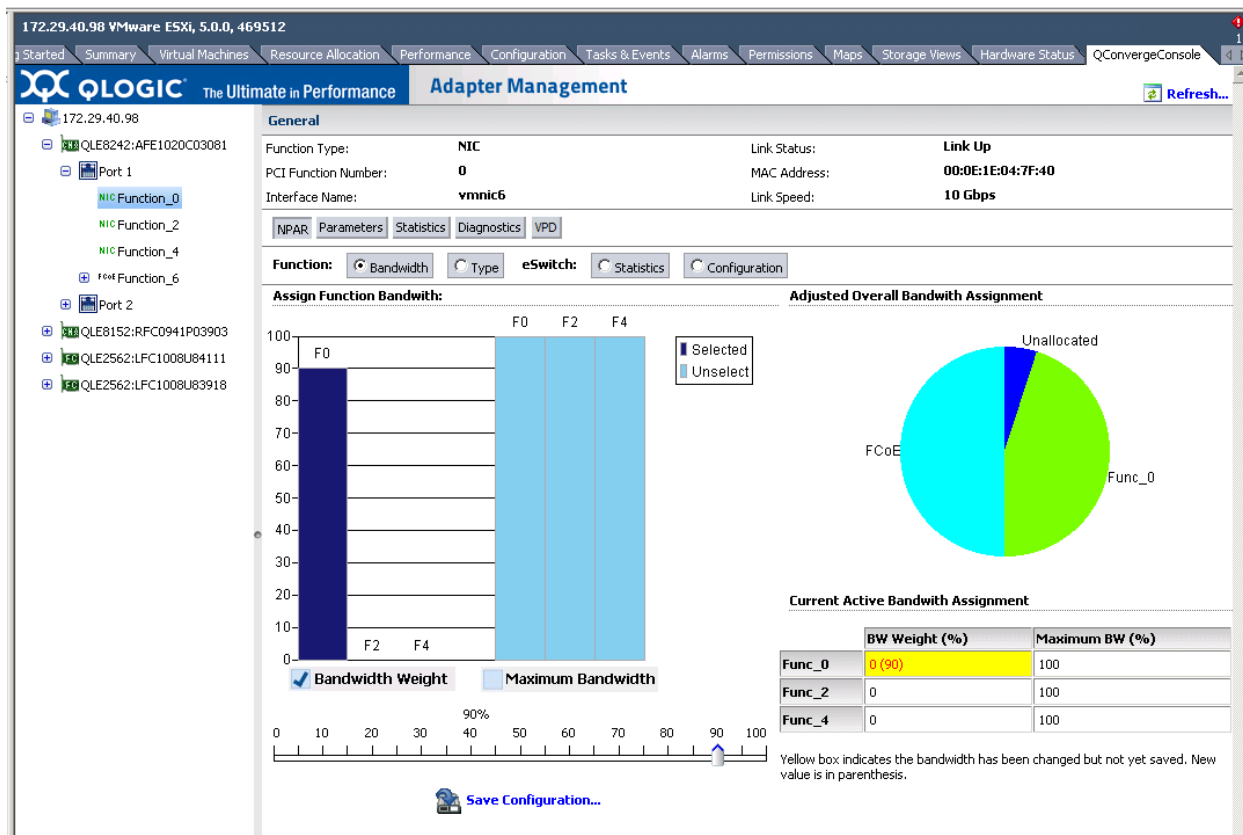
Figure 4-34. Bandwidth Configuration

The content window is divided into four sections:

- **Assign Function Bandwidth:** This section contains a slider control used to set the parameters **Bandwidth Weight** and **Maximum Bandwidth**. To set a parameter, check the appropriate box and drag the slider to the desired setting. To save the new settings, click **Save Configuration**.

■ **Adjusted Overall Bandwidth Assignment:** This is a pie chart diagram that shows the amount of the total bandwidth assigned to the NIC function.

■ **Current Active Bandwidth Assignment:** This lists the current settings for the Bandwidth Assignment and Maximum Bandwidth parameters. A yellow background indicates that the new value (in parentheses) has not been saved yet.

To save your changes, click **Save Configuration**.

Function Type

Selecting the **Type** option displays the current function type (NIC) and allows you to change the function type, as shown in Figure 4-35 and Figure 4-36. To change the function type, select the new type and click **Save Configuration**. The new type takes effect upon the next system reboot.

> **NOTE**
>
> On 8200 Series Converged Network Adapters:
> - ■ PCI functions 0 and 1 are always set to NIC and *cannot* be reconfigured.
> - ■ PCI functions 2 and 3 are configurable as either NIC or Disabled.
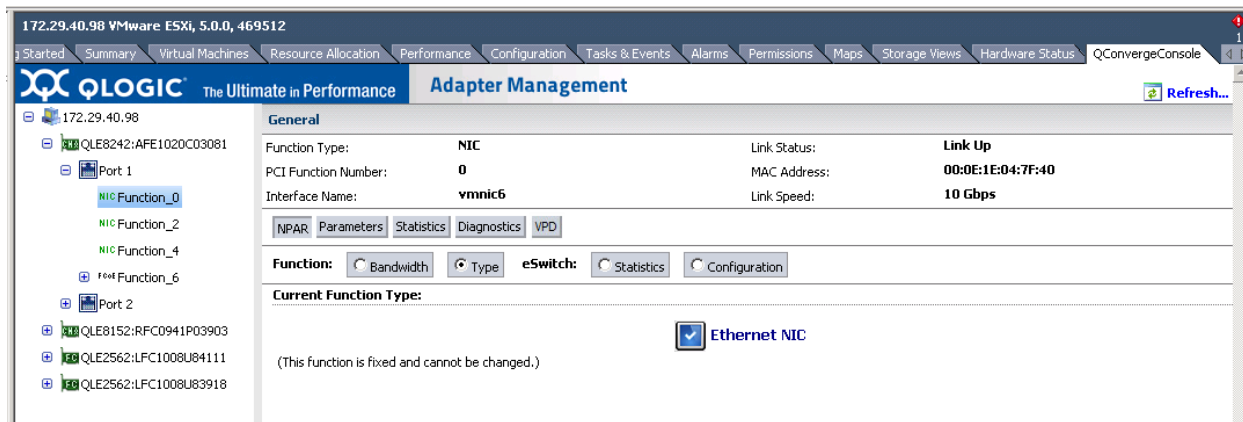
Figure 4-35. Converged Network Adapter NIC Function 0 or 1—Function Type
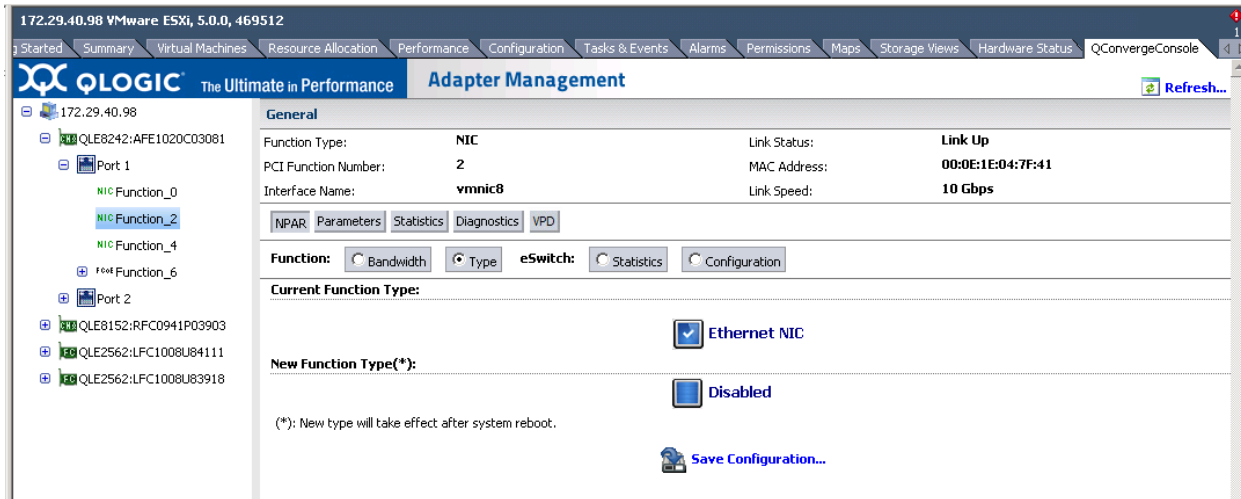
Figure 4-36. Converged Network Adapter NIC Function 2 or 3—Function Type

Switch Independent Partitioning Setup

This section provides Switch Independent Partitioning reference tables you can use when configuring NIC partitions using the various tools available.

■ Default Settings

■ Configuration Options

■ Switch Independent Partitioning Configuration Parameters and Setup Tools

Default Settings

Before configuring NIC partitions, the adapter appears as a simple dual-port 10GbE adapter with Switch Independent Partitioning settings shown in Table 4-5.

Table 4-5. Default Configuration

Function Number	Function Type	Physical Port Number	Minimum Bandwidth (%)	Maximum Bandwidth (%)	Default Function Type
0	NIC	0	0	100	Enabled as NIC
1	NIC	1	0	100	Enabled as NIC

Configuration Options

Depending on your system requirements and operating environment, you can set up the adapter port partitions to support different function types. Table 4-6 shows the available function types and configurable parameters.

Table 4-6. Configuration Options

Function Number	Function Type	Physical Port Number	Bandwidth (%)		Default Function Type
			Minimum [a]	Maximum [b]	
0	NIC	0	0	100	NIC
1	NIC	1	0	100	NIC
2	Disabled/NIC	0	0	100	NIC
3	Disabled/NIC	1	0	100	NIC
4	iSCSI/NIC/Disabled	0	0	100	iSCSI
5	iSCSI/NIC/Disabled	1	0	100	iSCSI
6	FCoE/NIC/Disabled	0	0	100	FCoE
7	FCoE/NIC/Disabled	1	0	100	FCoE

[a] Minimum Bandwidth: Minimum guaranteed bandwidth, specified as a percentage of the link speed. The total across all partitions will add up to less than the maximum link bandwidth. The queue's rate is allowed to exceed the specified value up to the maximum rate, if excess bandwidth is available on the physical port link.

[b] Maximum bandwidth: Maximum allowed bandwidth, specified as a percentage of the link speed. The queue's rate is not allowed to exceed the specified value, even if excess bandwidth is available on the physical port link. The total across all partitions cannot be greater than the maximum link bandwidth.

Switch Independent Partitioning Configuration Parameters and Setup Tools

Table 4-7 identifies which parameters you can configure using each of the available management tools.

Table 4-7. Switch Independent Partitioning Configuration Parameters and Setup Tools

Tools/Configurable Switch Independent Partitioning Parameters	Function Type [a]	Minimum Bandwidth [b] (Range 0–100%)	Maximum Bandwidth [c] (Range 0–100%)
Dell System Setup, Lifecycle Controller, or other HII browser	Yes	Yes, configurable	Yes, configurable
QLogic NIC OptionROM (Hot Key CTRL+Q during POST)	Yes	Yes, configurable	Not configurable, read only
QLogic QConvergeConsole GUI/CLI for supported Windows and Linux operating systems	Yes	Yes, configurable but only for NIC partitions, not for storage (iSCSI/FCoE) partitions	Yes, configurable but only for NIC partitions, not for storage (iSCSI/FCoE) partitions
QLogic Windows Device Manager–NIC Property Page	Yes	Yes, configurable but only for NIC partitions, not for storage (iSCSI/FCoE) partitions	Yes, configurable but only for NIC partitions, not for storage (iSCSI/FCoE) partitions
QLogic QConvergeConsole Plug-ins for vSphere	Yes	Yes, configurable but only for NIC partitions, not for storage (iSCSI/FCoE) partitions	Yes, configurable but only for NIC partitions, not for storage (iSCSI/FCoE) partitions

[a] These changes require system reboot to take effect. Refer to Table 4-6 for available function type options of each partition.

[b] For FCoE, the DCBX/ETS negotiated bandwidth overwrites the manually configured bandwidth.

[c] For FCoE, the DCBX/ETS negotiated bandwidth overwrites the manually configured bandwidth.

> **NOTE**
>
> Table 4-8 applies to QME8262-k only.

Table 4-8. Switch Independent Partitioning Wake-on-LAN (WoL) and PXE Boot Support

Switch Independent Partitioning Function	WoL	PXE Boot
Function 0	Yes	Yes
Function 1	Yes	Yes
Function 2	No	No
Function 3	No	No
Function 4	No	No
Function 5	No	No
Function 6	No	No
Function 7	No	No

5 Boot Configuration

Overview

This section provides the following information about boot configuration for the QLogic adapter:

- **Boot from SAN**

    Booting servers from SANs can provide significant benefits in today's complex data center environments. One of the driving forces behind SANs is the need to deliver mission-critical data quickly, at any time, without interruptions or delays.

- **Dell System Setup**

    The Dell System Setup allows you to configure a network adapter.

- **PXE Boot Setup**

    The PXE allows a workstation to boot from a server on a network before booting the operating system on the local hard drive.

- **iSCSI Configuration Using Fast!UTIL**

    *Fast!*UTIL is the QLogic iSCSI BIOS Configuration utility used to configure the adapter for iSCSI.

- **iBFT Boot Setup**

    iBFT is a component of the *Advanced Configuration and Power Interface Specification* (ACPI) 3.0b standard that provides operating systems a standard way to boot from a software-initiated iSCSI protocol. This component uses software-based initiators.

- **DHCP Boot Setup (iSCSI)**

    DHCP boot provides a method for the iSCSI initiator to acquire target parameters from a DHCP server.

Boot from SAN

This section provides the following information on boot from SAN:

- General Boot from SAN
- Windows Boot from SAN
- Linux Boot from SAN
- ESX Boot from SAN

Additional information can be found in the driver readme and release notes.

General Boot from SAN

The following high-level boot from SAN instructions apply to all OSs: Linux, Windows, and ESX:

1. Set up the boot order to disable boot from the local disk or disconnect internal hard drives.

2. Configure a single path to the storage array when installing Microsoft Windows.

   For multiple QMD8262-k/QLE8262/QME8262-k port configurations (for redundancy), only one adapter port should be connected to the SAN during installation. The same recommendation applies to the storage controller ports.

3. Enable the adapter port BIOS.

4. Specify a boot LUN from the QLogic adapter *Fast!*UTIL BIOS configuration utility.

Windows Boot from SAN

This section provides the following procedures:

- Creating a Driver Disk
- Windows 2008 Boot From SAN

Creating a Driver Disk

Before you begin boot from SAN for Windows, create a driver disk as follows:

1. From the Dell Web site, download the appropriate driver package

2. Extract the drivers using the QLogic Installer or the Dell DUP:

   QLogic setup installer: Issue the following command from the directory where the installer is located:

   ```
   setup.exe /s /z"/extractdrivers=\"c:\temp\"
   /logfile=\"C:\temp\log.txt\""
   ```

Dell DUP: Issue the following command to extract the drivers to the
appropriate path/location:

```
<package_name> /s /e=<path>
```

Windows 2008 Boot From SAN

For Windows 2008, follow these steps to perform an initial OS installation with the
adapter as boot or as add-on.

NOTE

- The following procedure requires a USB Flash drive; see "Creating a
  Driver Disk" on page 176.

- Ensure that the target SAN device is available and configured before
  beginning the procedure.

To perform the Windows 2008 boot from SAN, follow these steps.

1. Start the Windows installation procedure using the setup CD.

2. When the following message appears, click **Custom (Advanced)**:

   ```
   Which type of installation do you want?
   ```

3. When the following message appears, click **Load Driver**:

   ```
   Where do you want to install Windows?
   ```

4. At the **Load Driver** window, insert the USB Flash drive in a USB port and
   then click **OK**.

5. To continue with the standard installation procedure, click **Next**.

Linux Boot from SAN

This section provides the following procedures:

- Red Hat Enterprise Linux Boot from SAN
- SUSE Linux Enterprise Server (Novell) Boot from SAN

Red Hat Enterprise Linux Boot from SAN

Follow these steps to install boot from SAN for Red Hat Enterprise Linux for all
supported versions:

1. Insert the current Linux Red Hat product CD #1 in the CD drive and then
   boot the system.

   The system boots from the CD and stops at the boot prompt.

2. At the boot prompt, type **Linux dd**, and then press ENTER.

The Driver Disk message box displays the following prompt:

Do you have a driver disk?

3. Click **YES** and then press ENTER.

4. In the **Driver Disk Source** window, select the driver source:

    ❑ If the driver file is on a disk, select **fd0**, then press ENTER.

    ❑ If the driver file is on a CD, select **hd*x*** (where *x* is the CD drive letter) and then press ENTER.

    The **Insert Driver Disk** window opens.

5. Insert the iso-dd-kit disk into either the disk drive or the CD drive, depending on the option selected in .

6. Click **OK** and then press ENTER.

    The SCSI driver is loaded automatically.

    The **Disk Driver** window opens and prompts you to install more drivers.

7. Click **NO** and then press ENTER.

8. Insert the current Linux Red Hat product CD #1 in the CD drive (remove the iso-dd-kit CD first, if necessary) and then press ENTER.

9. Follow the on-screen instructions to complete the installation.

SUSE Linux Enterprise Server (Novell) Boot from SAN

Follow these steps to install boot from SAN for SUSE Linux Enterprise Server (Novell®), all supported versions:

1. Insert the SLES CD #1 in the CD drive, and then boot the system.

2. On the main installation window, press one of the following keys:

    ❑ Press F5 for SLES 11 SP4
    ❑ Press F6 for SLES 11 and SLES 11 SP1 and SLES 12

    The system prompts you to select either **Yes**, **No**, or **File**.

3. Select **Yes**.

4. Select the installation option and then press ENTER.

    The system prompts you to choose a driver update medium.

5. Ensure that the dd-kit is inserted in the appropriate drive, either the CD-ROM or floppy drive, select the drive (for dd-kit ISO images, select the CD-ROM), and then press ENTER.

    The driver update is loaded to the system.

    If the driver update was successful, the **Driver Update OK** message appears.

6. Press ENTER.

7. If the system prompts you to update another drive, click **BACK** and then press ENTER.

   The following message appears:

   `Make sure that CD number 1 is in your drive.`

8. Insert the SLES CD #1 in the drive and then click **OK**.

9. Follow the on-screen instructions to complete the installation.

ESX Boot from SAN

For VMware ESX, follow these steps to install the driver for devices as part of a new ESX installation.

> **NOTE**
>
> You need the ESX installation DVD to begin.

1. Place the ESX installation DVD in the host DVD drive.

2. Restart the host.

3. Read and accept the terms of the license agreement.

4. Select a keyboard type.

5. When prompted for custom drivers, select **Yes** to install custom drivers.

6. To eject the ESX installation DVD, click **Add**.

7. Place the driver CD in the ESX host DVD drive.

8. Select the driver module from which to import drivers to the ESX host and then click **Next**.

9. On the **Load the system drivers?** message box, click **Yes**.

10. After loading the driver module, continue installing ESX.

11. After the driver is installed, you are prompted to swap the driver CD with the ESX installation DVD.

Dell System Setup

The Dell System Setup allows you to configure a network adapter. The System Setup contains the following pages related to the QLogic adapter:

■ System Setup Main menu (see Figure 5-1)

    ❑ Device Settings (see Figure 5-2)

        ■ Main Configuration

        ❑Device and Firmware Image Information

        ❑NIC Partitioning (Switch Independent Partitioning) Configuration

            ■Global Bandwidth Allocation

            ■Partition 1 Configuration

            ■Partition 2 Configuration

            ■Partition 3 Configuration

            ■Partition 4 Configuration

        ❑NIC Configuration

        ❑iSCSI Configuration

            ■iSCSI General Parameters

            ■iSCSI Initiator Parameters

            ■iSCSI First Target Parameters

            ■iSCSI Second Target Parameters

        ❑FCoE Configuration

The first three pages are used to access the Dell System Setup, select the QLogic adapter, and access the adapter's Main Configuration page. This page provides access to all available configuration options for the selected adapter.

Accessing Dell System Setup

When you first start the host server that contains QLogic adapters, the POST starts. Running POST gives you access to the Dell System Setup.

To access the Dell System Setup:

1. While running POST, press F2. The Main menu for the Dell System Setup opens.

NOTE

    Depending on your server model and System Setup version, the screens you see might differ from those shown.

2. Select **Device Settings** (see Figure 5-1).
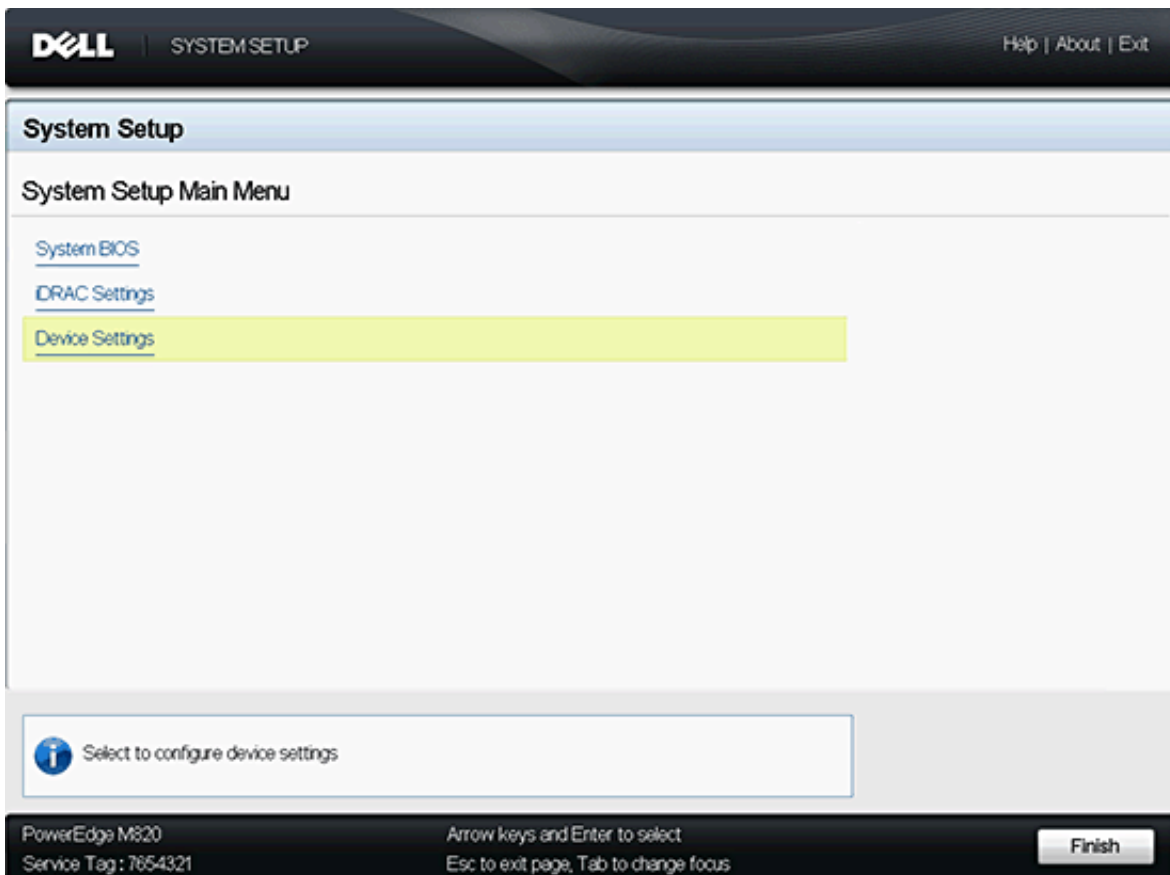
Figure 5-1. System Setup Main Menu Page

3.   In the Device Settings screen, select the adapter that you want to configure
     or display information about (see Figure 5-2).
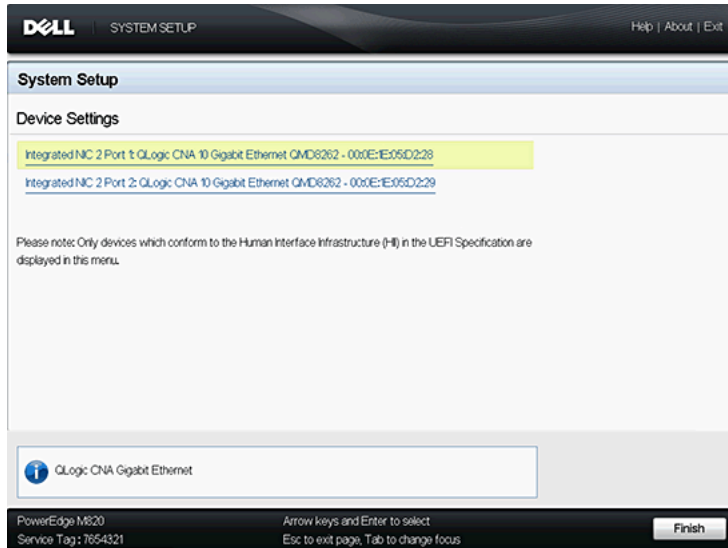
Figure 5-2. Selecting the Device to Configure

The next screen that appears (see Figure 5-3) is the Main Configuration
page for the selected adapter.
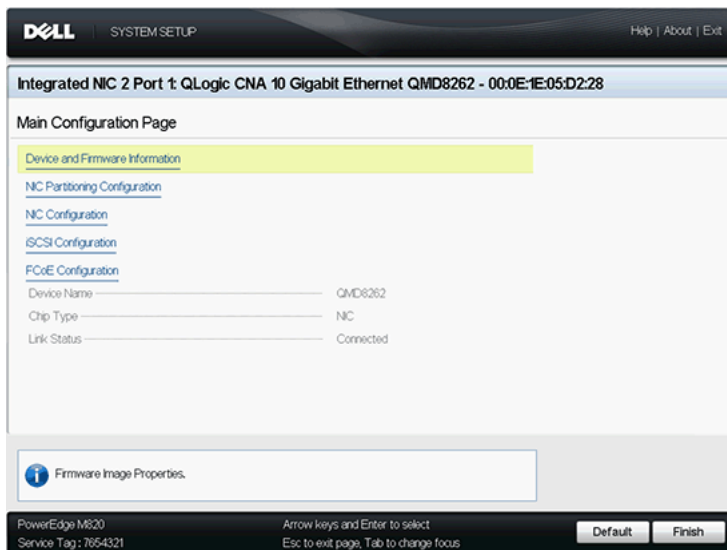
Figure 5-3. Main Configuration for Selected Adapter

Main Configuration

The Main Configuration page (see Figure 5-3 on page 182) displays information about the selected network adapter and provides the following options. Each option is described in detail in the sections that follow:

- Device and Firmware Image Information (read-only)

- NIC Configuration

- iSCSI Configuration

- FCoE Configuration

- NIC Partitioning (Switch Independent Partitioning) Configuration

Device and Firmware Image Information

This page displays device and firmware image information (see Figure 5-4). The information on this page is read only and cannot be modified by the user. The page displays the following information:

- **Device Name**: This file is the model number of the adapter.

- **Chip Type**: This file identifies the type of chip in the adapter.

- **Family Version**: This file is the multiboot image version of the Flash kit.

- **EFI Version**: This file the version number of the extensible firmware interface (EFI).
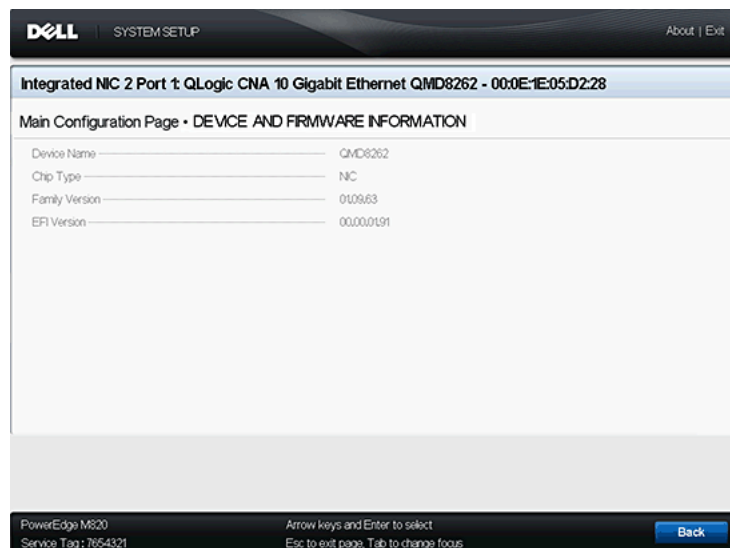
Figure 5-4. Device and Firmware Image Information

NIC Configuration

The NIC Configuration page (see Figure 5-5) allows the user to set the following:

■ **Legacy Boot Protocol**: Select **PXE**, **iSCSI**, or **None** to control the network boot protocol. The configuration and enablement of iSCSI and FCoE are controlled separately.

■ **Wake on LAN**: This option enables or disables server power-on using an in-band magic packet.

■ **Link Speed**: This option is the link speed of the NIC. This field is read only and cannot be modified by the user.
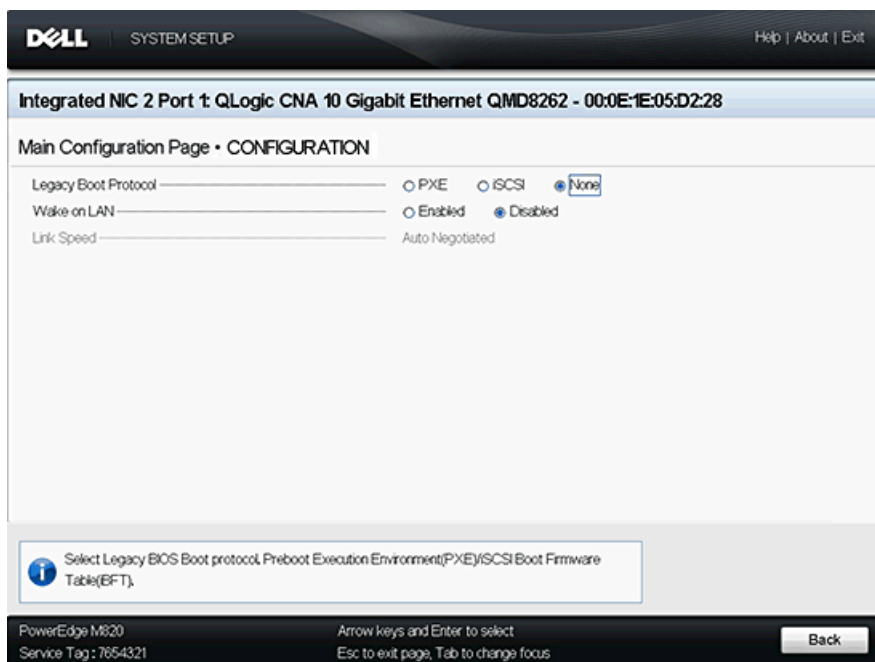
Figure 5-5. NIC Configuration

iSCSI Configuration

The iSCSI Configuration page (see Figure 5-6) provides the following choices for iSCSI configuration:

- iSCSI General Parameters
- iSCSI Initiator Parameters
- iSCSI First Target Parameters
- iSCSI Second Target Parameters
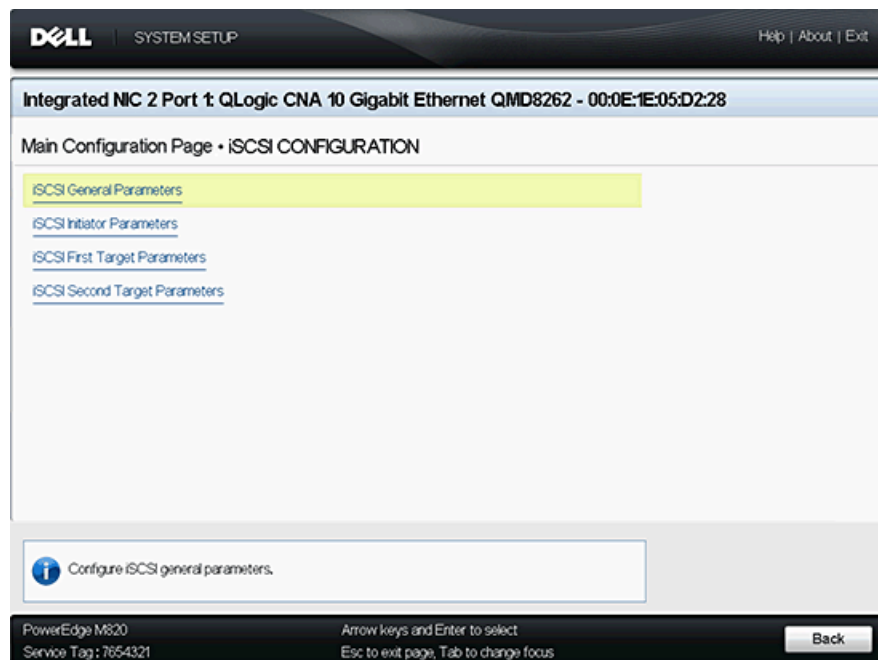
Figure 5-6. iSCSI Configuration

iSCSI General Parameters

The iSCSI General Parameters page (see Figure 5-7) lets you set the following:

- **TCP/IP Parameters via DHCP**: Select **Enabled** or **Disabled**. When set to **Enabled**, the adapter uses the DHCP to obtain its IP address, subnet mask, and gateway IP address.

- **iSCSI Parameters via DHCP**: Select **Enabled** or **Disabled**. When set to **Enabled**, the initiator acquires its IP address from a DHCP server. The initiator also acquires all boot parameters for an iSCSI target if the boot parameters are specified on the DHCP server. This setting requires advanced users to manually input vendor-specific information on the DHCP server and in the **DHCP Vendor ID** field of the iSCSI General Parameters.

- **CHAP Authentication**: Select **Disabled** to disable CHAP authentication. Select **Target** or **Target/Initiator** to enable CHAP or bidirectional CHAP and to configure the CHAP name and CHAP secret.

- **CHAP Mutual Authentication**: This option indicates if the mutual CHAP is **Enabled** or **Disabled**.

- **Boot to Target**: Select **Enabled** or **Disabled** for boot to target.

- **DHCP Vendor ID**: This field allows the user to provide the DHCP vendor ID when **iSCSI Parameters via DHCP** is enabled. The DHCP vendor ID is the identification string sent to DHCP server.
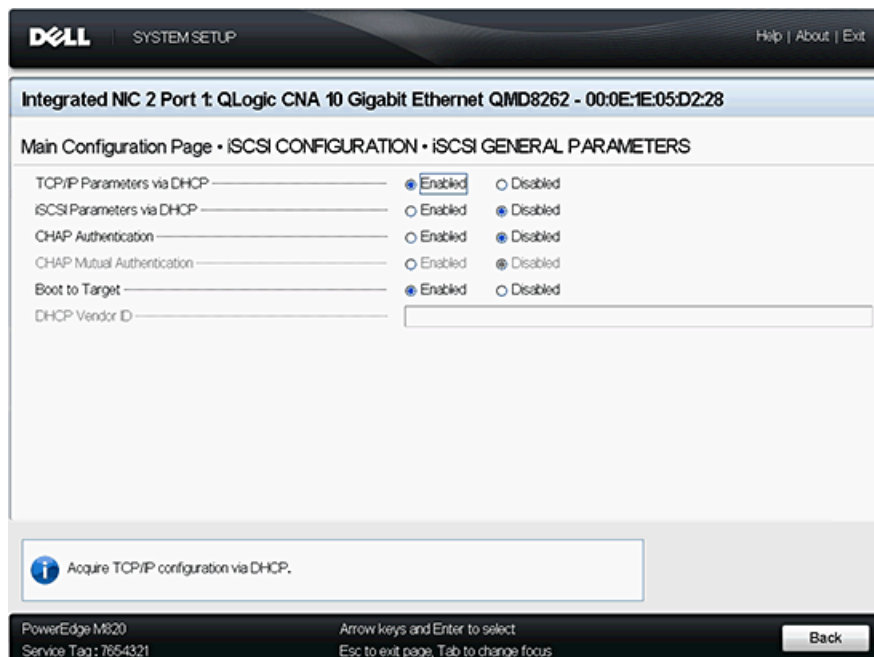
Figure 5-7. iSCSI General Parameters

iSCSI Initiator Parameters

The iSCSI Initiator Parameters page (see Figure 5-8 and Figure 5-9) lets you set the following:

- **IPv4**: This field indicates whether or not the iSCSI initiator uses the IPv4 protocol. If **Enabled**, the following parameters can be set:

  - ❑ **IPv4 Address**: When **TCP/IP Parameter via DHCP** is set to **Disabled**, this field must contain a valid IP address.

  - ❑ **IPv4 Subnet Mask**: When **TCP/IP Parameter via DHCP** is set to **Disabled**, this field must contain a valid subnet mask.

  - ❑ **IPv4 Default Gateway**: When **TCP/IP Parameter via DHCP** is set to **Disabled**, this field must contain a valid gateway IP address.

- **IPv6**: This field indicates whether or not the iSCSI initiator uses the IPv6 protocol. If **Enabled**, the following parameters can be set:

  - ❑ **IPv6 Link Local Address Auto Mode**: This enables or disables the automatic assignment of the IPv6 link local address. The default is **Disabled** (Manual assignment).

  - ❑ **IPv6 Link Local Address**: When **IPv6 Link Local Address Auto Mode** is **Disabled**, this field must contain the IPv6 link local address of the iSCSI initiator.

  - ❑ **IPv6 Auto-Configuration**: This option is specific to IPv6. It controls the source of the initiator IP address: DHCP if **Enabled** or Static assignment if **Disabled**.

  - ❑ **IPv6 Routable Address 1**: If **IPv6 Auto-Configuration** is **Disabled**, this field must contain the IPv6 address of the iSCSI initiator.

  - ❑ **IPv6 Routable Address 2**: If **IPv6 Auto-Configuration** is **Disabled**, this field may contain the IPv6 address of the iSCSI initiator.

  - ❑ **IPv6 Default Gateway**: This field specifies the IPv6 default gateway of the iSCSI initiator.

- **iSCSI Initiator Name**: You can configure the iSCSI name of the initiator. By default, the iSCSI name (containing two parts, the standard product model and serial number) is used.

- **CHAP ID**: This field allows you to configure the bidirectional CHAP name.

- **CHAP Secret**: This field allows you to configure the bidirectional CHAP secret.
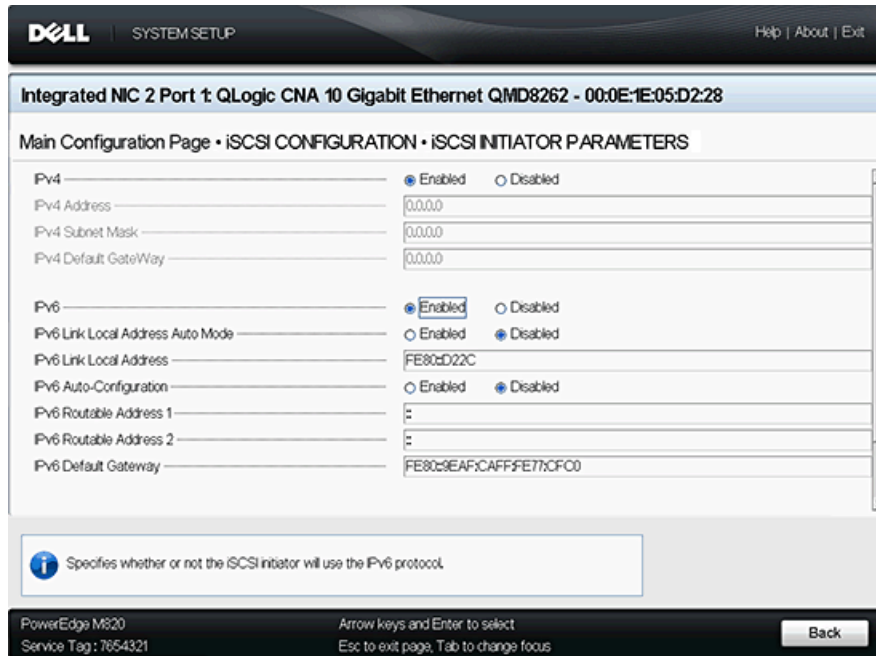
Figure 5-8. iSCSI Initiator Parameters—Start of Page
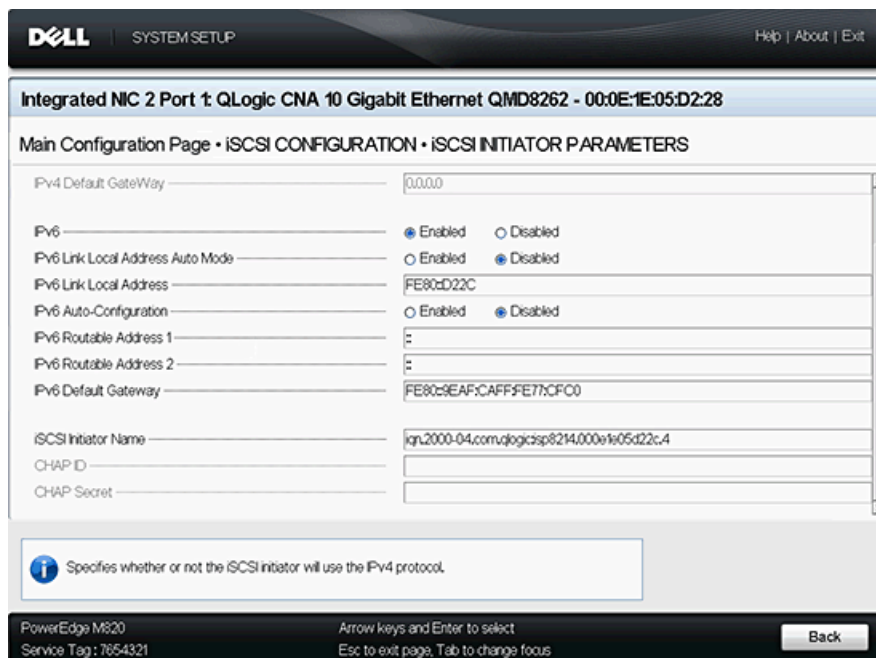
Figure 5-9. iSCSI Initiator Parameters—End of Page

iSCSI First Target Parameters

The iSCSI First Target Parameters page (see Figure 5-10) lets you set the following:

- **IP Version**: This option indicates whether **IPv4** or **IPv6** is selected.

- **IPv4 Address**: If IPv4 is selected, this field let you specify the IPv4 address of the intended iSCSI boot target.

- **IPv6 Address**: If IPv4 is selected, this field lets you specify the IPv6 address of the intended iSCSI boot target.

- **Target Port**: This field lets you specify the port used by the target. It must be set to 3260. By default, iSCSI uses port 3260. Do not modify this field unless it is necessary.

- **Boot LUN**: This field lets you specify the LUN of the boot device.

- **Target Name**: This field lets you specify the iSCSI target name of an intended boot target.

- **CHAP ID**: This field lets you configure the CHAP name. This field is dependent upon the CHAP configuration on the iSCSI General Parameters page.

- **CHAP Secret**: This field lets you configure the CHAP secret. This field is dependent upon the CHAP configuration on the iSCSI General Parameters page.

Figure 5-10. iSCSI First Target Parameters

iSCSI Second Target Parameters

The iSCSI Second Target Parameters page (see Figure 5-11) lets you set the following:

■ **IP Version**: This option indicates whether **IPv4** or **IPv6** is selected.

■ **IPv4 Address**: If IPv4 is selected, this field let you specify the IPv4 address of the intended iSCSI boot target.

■ **IPv6 Address**: If IPv4 is selected, this field lets you specify the IPv6 address of the intended iSCSI boot target.

■ **Target Port**: This field lets you specify the port used by the target. It must be set to 3260. By default, iSCSI uses port 3260. Do not modify this field unless it is necessary.

■ **Boot LUN**: This field lets you specify the LUN of the boot device.

■ **Target Name**: This field lets you specify the iSCSI target name of an intended boot target.

■ **CHAP ID**: This field lets you configure the CHAP name. This field is dependent upon the CHAP configuration on the iSCSI General Parameters page.

■ **CHAP Secret**: This field lets you configure the CHAP secret. This field is dependent upon the CHAP configuration on the iSCSI General Parameters page.

Figure 5-11. iSCSI Second Target Parameters

FCoE Configuration

The FCoE Configuration page (see Figure 5-12) lets you set the following:

■ **Connect**: Select **Enabled** to enable OS boot from an FCoE storage device, or **Disabled** to disable OS boot from an FCoE storage device.

■ **Boot from LUN**: The boot device LUN. This is a 16-bit value. This parameter is selectable only if the **Boot** parameter is set to **Enabled**.

■ **Boot from Target**: The boot device worldwide port name. This is a 64-bit value. This parameter is selectable only if the **Boot** parameter is set to **Enabled**.
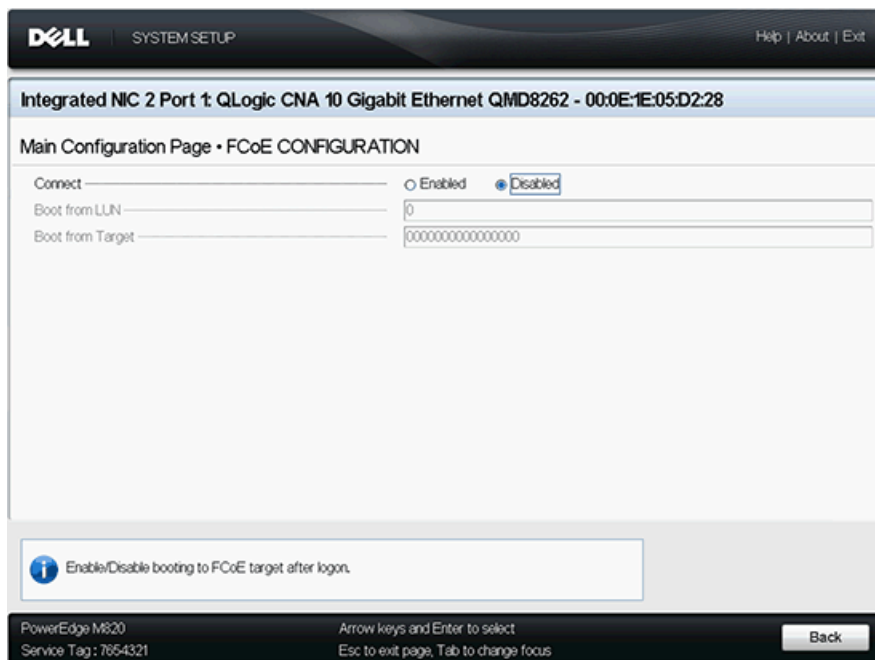
Figure 5-12. FCoE Configuration

NIC Partitioning (Switch Independent Partitioning) Configuration

The NIC Partitioning (Switch Independent Partitioning) Configuration page (see Figure 5-13) provides the following choices for Switch Independent Partitioning configuration:

■ Global Bandwidth Allocation
■ Partition 1 Configuration
■ Partition 2 Configuration
■ Partition 3 Configuration
■ Partition 4 Configuration
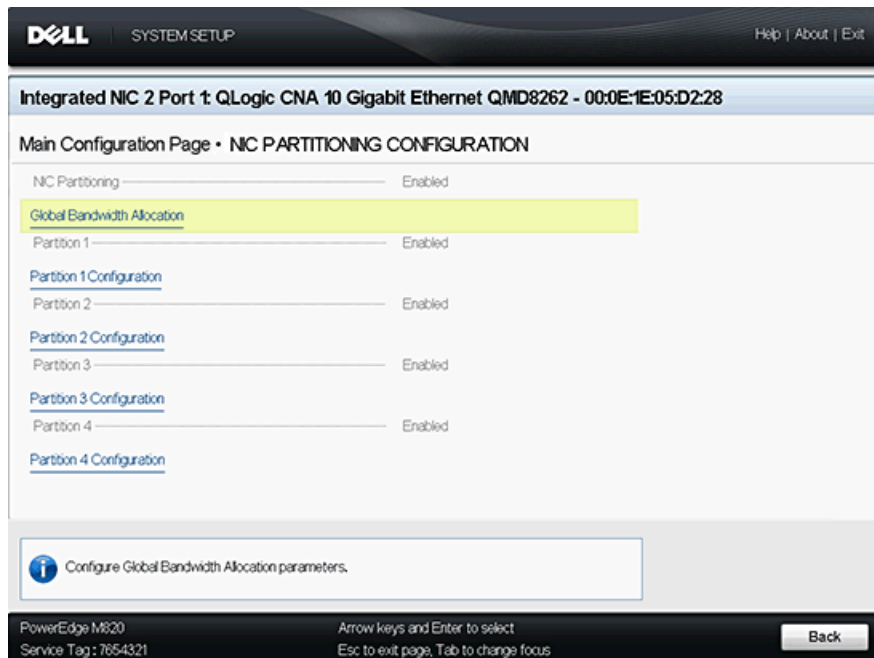
Figure 5-13. NIC Partitioning (Switch Independent Partitioning) Configuration

Global Bandwidth Allocation

The Global Bandwidth Allocation page (see Figure 5-14) lets you change a partition's relative bandwidth weighting and maximum bandwidth if it has been enabled. For more information on bandwidth allocation, refer to "Configuration Options" on page 172.

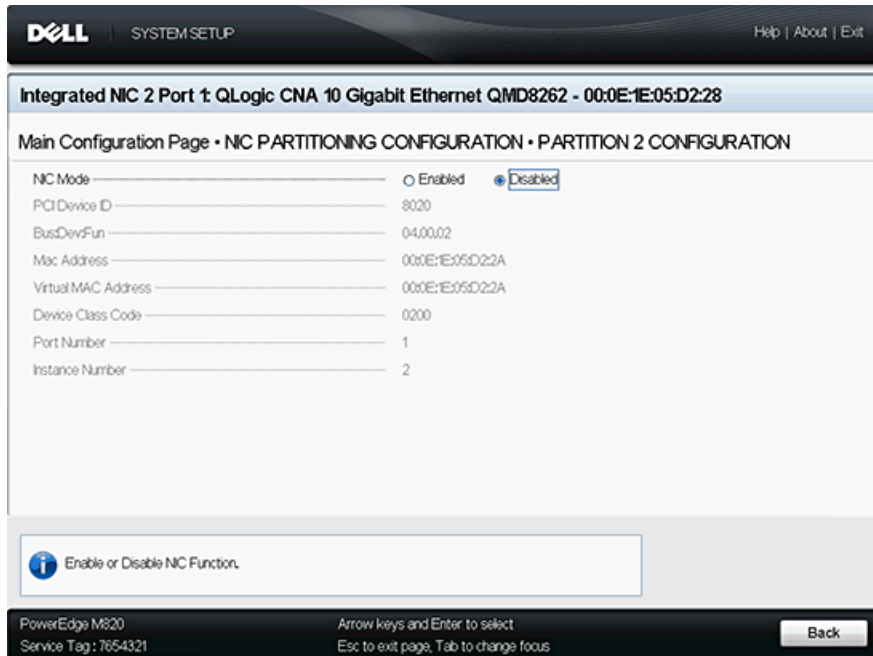*Figure 5-14. Global Bandwidth Allocation*

Partition 1 Configuration

The Partition 1 Configuration page (see Figure 5-15) has only one selection, **Enabled** for **NIC Mode**.

Figure 5-15. Partition 1 Configuration

Partition 2 Configuration

The Partition 2 Configuration page (see Figure 5-16) lets you set **NIC Mode** to **Enabled** or **Disabled**.
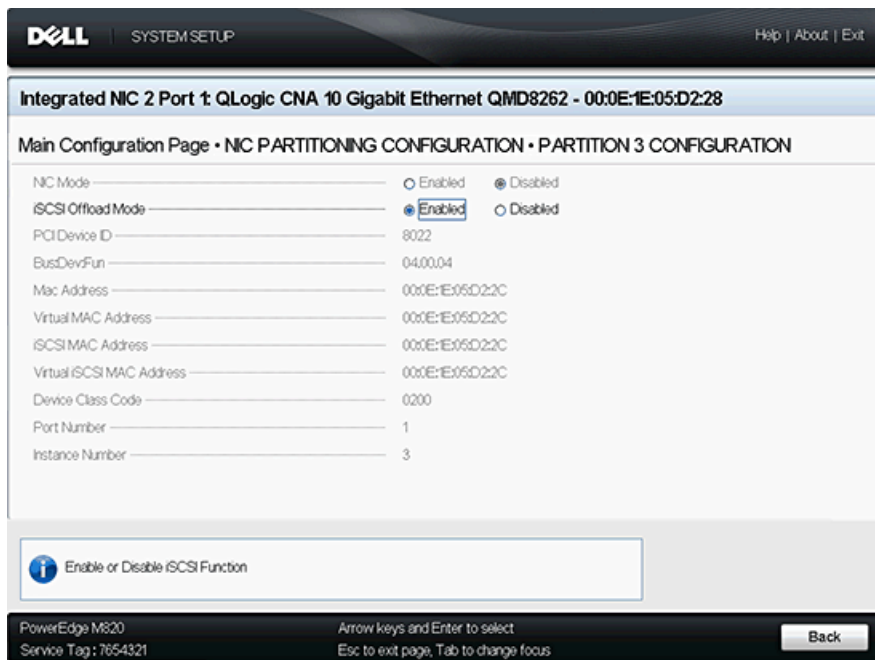

Figure 5-16. Partition 2 Configuration

Partition 3 Configuration

The Partition 3 Configuration page (see Figure 5-17) lets you set **NIC Mode** to **Enabled** or **Disabled**.

If you select **Disabled** for **NIC Mode**, you can select **Enabled** or **Disabled** for **iSCSI Offload Mode**.

Figure 5-17. Partition 3 Configuration

Partition 4 Configuration

The Partition 4 Configuration page (see Figure 5-18 and Figure 5-19) lets you set **NIC Mode** to **Enabled** or **Disabled**.

If you select **Disabled** for **NIC Mode**, you can select **Enabled** or **Disabled** for **iSCSI Offload Mode**.
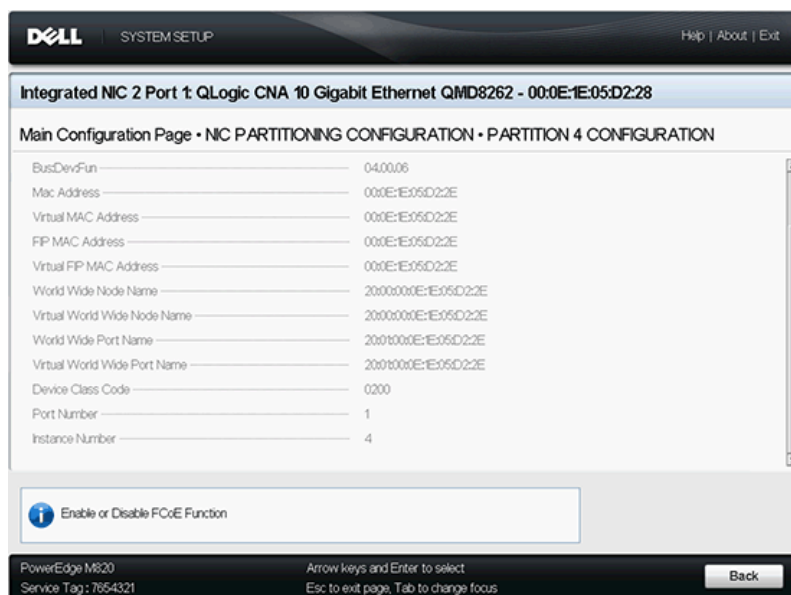

Figure 5-18. Partition 4 Configuration—Start of Page

*Figure 5-19. Partition 4 Configuration—End of Page*

PXE Boot Setup

The PXE allows a workstation to boot from a server on a network before booting the operating system on the local hard drive.

Configuring PXE Boot

This section provides procedures for configuring the ProductLine to perform PXE boot. The example uses function 1 and NIC 1.

To configure PXE boot:

1. During POST, press the CTRL+Q keys to enter the **QLogic 8200 Series CNA Function Configuration** window.

2. On the CNA Function Configuration main window, ensure that **Protocol** is set to **PXE** (Figure 5-20).
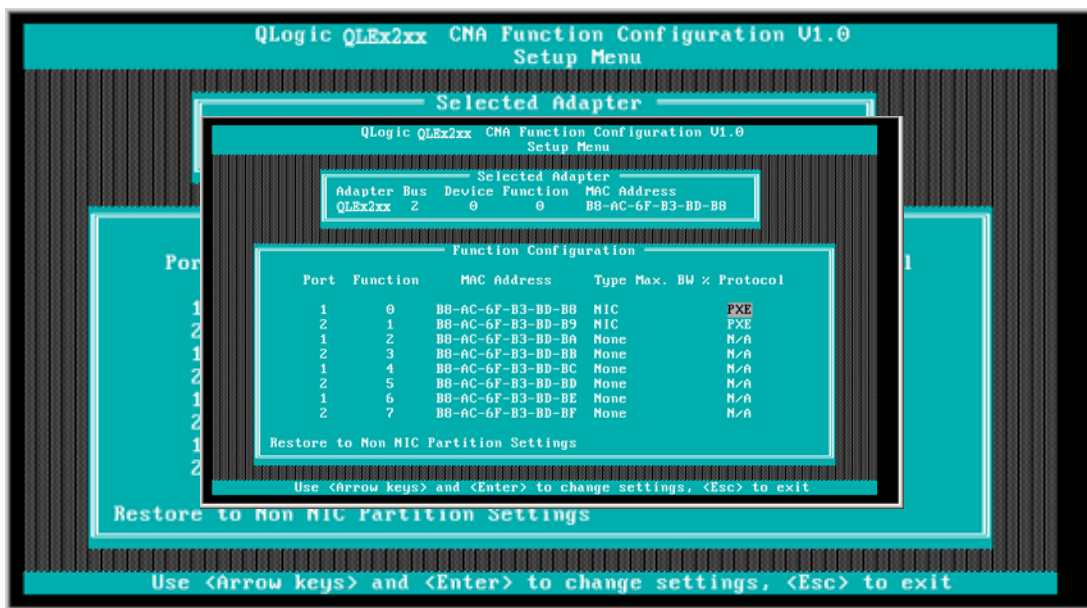
Figure 5-20. QLogic 82xx CNA Function Configuration

3. Press the ESC key to exit.

4. Select **Save changes** to save your edits, exit, and reboot the system.

5. During POST, press the F1 or F2 key to enter the BIOS system.

6. Select **Boot Settings**, and then press ENTER.

7. Select the **Boot Sequence** option, and then press ENTER.

8. Select the **QLogic PXE** entry as the first boot option.

9. Press the ESC key, and then select **Save changes and exit**.

   The system reboots.

10. After the system reboot, follow the window prompt for PXE boot server for the installation of OS of your choice.

    The system attempts to boot from the PXE. For example:

```
Attempting Boot From NIC
QLogic PXE v2.0.x.x PCI x.x Px
Copyright (C) 2009-2014 QLogic Corporation
Initializing...
  CLIENT MAC ADDR: xx xx xx xx xx xx
  CLIENT IP: xx.xx.xx.xx MASK: xx.xx.xx.xx
  DHCP IP: xx.xx.xx.xx
```

> **NOTE**
>
> Ensure that the 8200 Series Adapters' drivers are added in the OS image from the PXE server to complete the OS install.

iSCSI Configuration Using *Fast!*UTIL

QLogic's *Fast!*UTIL provides one method of configuring the
QMD8262-k/QLE8262/QME8262-k adapter for iSCSI.

- Accessing Fast!UTIL
- Configuring Host Adapter Settings
- Configuring iSCSI Boot Settings
- Configuring DHCP iSCSI Boot for IPv4

Accessing *Fast!*UTIL

To access *Fast!*UTIL for PXE, iSCSI, and FCoE, respectively, press CTRL+Q during
the QLogic adapter BIOS initialization PXE, iSCSI, or FCoE QLogic banner display.

It might take a few seconds for the *Fast!*UTIL menu to appear. The *Fast!*UTIL
Options menu for iSCSI contains the following selections:

```
Configuration Settings
Scan iSCSI Devices
iSCSI Disk Utility
Ping Utility
Loopback Test
Reinit Adapter
Select Host Adapter
Exit Fast!UTIL
```

*Fast!*UTIL prompts you to select the adapter you want to configure. After changing
the settings, exit and reboot your system to load the new *Fast!*UTIL parameters.

Configuring Host Adapter Settings

On the *Fast!*UTIL Options menu, select **Configuration Settings** and then select
Host Adapter Settings to configure the following:

- **Initiator IPv4 and IPv6 Address Through DHCP**
  When set to **Yes**, the adapter uses the DHCP to obtain its IP address,
  subnet mask, and gateway IP address. Press ESC three times to save the
  changes.

  > **NOTE**
  >
  > These values are not populated when you save the settings for the IP
  > address, Subnet Mask, Gateway, IPv6 Routable Address 1, and IPv6
  > Routable Address 2.

- **Initiator IPv4/IPv6 Address**
  When DHCP is set to **No**, this field must contain a valid IP address.

- **Subnet Mask**
  When DHCP is set to **No**, this field must contain a valid subnet mask.

- **Gateway IP Address**
  When DHCP is set to **No**, this field must contain a valid gateway IP address; otherwise, the system under configuration can communicate only with other nodes on its LAN.

- **Initiator iSCSI Name**
  Press ENTER to configure the iSCSI name of the initiator. Specifying the Initiator iSCSI Name is optional; however, if not specified, the iSCSI function uses the default iSCSI qualified name (IQN) that is programmed during manufacture. Use this option to specify an IQN other than the default.

- **Initiator CHAP Name**
  Press ENTER to configure the bidirectional (BIDI) CHAP name.

- **Initiator CHAP Secret**
  Press ENTER to configure the BIDI CHAP secret.

Configuring iSCSI Boot Settings

On the *Fast!*UTIL Options menu, select **Configuration Settings** to access the **iSCSI Boot Settings** and specify the drive you want to boot from. Press C to clear the boot device information from the primary and alternate boot locations.

> **NOTE**
>
> For an alternate method of iSCSI boot from SAN, use the fields in the iBFT. For details, see "iBFT Boot Setup" on page 213.

iSCSI boot setup information includes the following:

- Boot Device Primary and Alternate
- Adapter Boot Mode
- Primary and Alternate Boot Device Settings
- Configuring the iSCSI Boot Parameters
- Configuring QLogic iSCSI Boot
- Booting

Boot Device Primary and Alternate

After configuring a device (through **Primary/Alternate Boot Device Settings**), press ENTER on these locations to view a list of available devices. To select an iSCSI boot device, highlight the device and then press ENTER.

Adapter Boot Mode

- **Disable**—Select this option to disable the ROM BIOS on the adapter, freeing space in upper memory.

- **Manual**—Select this option to manually configure the parameters for the iSCSI boot device in the **Primary/Alternate Boot Device Settings** area. Manual parameters include the following:

  - ❏ **Initiator IPv4/IPv6 Address via DHCP**—When set to **Yes**, the adapter uses DHCP to obtain its IP address, subnet mask, and gateway IP address. Press ESC three times to save the changes. These values are not populated when you save the settings for the IP address, subnet mask, gateway, IPv6 routable addresses, and initiator IPv4 or IPv6 address. When DHCP is set to **No**, this field must contain a valid IP address.

  - ❏ **Subnet Mask**—When DHCP is set to **No**, this field must contain a valid subnet mask.

  - ❏ **Gateway IP Address**—When DHCP is set to **No**, this field must contain a valid gateway IP address.

  - ❏ **Initiator iSCSI Name**—Press ENTER to configure the iSCSI name of the initiator. By default, the iSCSI name (containing two parts, the standard product model and serial number) is used. If you change the iSCSI name, you need to write to the adapter's Flash memory to make the name change permanent.

  - ❏ **Initiator CHAP Name**—Press ENTER to configure the bidirectional CHAP name.

  - ❏ **Initiator CHAP Secret**—Press ENTER to configure bidirectional CHAP secret. You can access the iSCSI Boot Settings area from the **Configuration Settings** menu. Use these options to specify the drive from which you want to boot. Press C to clear the boot device information from the primary and alternate boot locations.

- **DHCP** (applies to IPv4 only)—Select this option to have the initiator acquire its IP address from a DHCP server. The initiator also acquires all boot parameters for an iSCSI target if the boot parameters are specified on the DHCP server. This setting requires advanced users to manually input vendor-specific information on the DHCP server and in the **Vendor ID** field of the DHCP Boot Settings area.

Primary and Alternate Boot Device Settings

- **Security Settings**—Press ENTER to access Primary Boot Security Settings. Press ENTER to enable or disable CHAP and bidirectional CHAP and to configure the CHAP name and CHAP secret. (Depending on your configuration, it might not be necessary to configure this option.)

- **Target IP**—When the Adapter Boot Mode is set to **Manual**, you must use this field to specify the IP address of the intended iSCSI boot target.

- **Target Port**—Enter the port used by the target. (By default, iSCSI uses port 3260. Do not modify this unless it is necessary.)

- **Boot LUN**—Enter the LUN of the boot device. (Depending on your configuration, it might not be necessary to change this option.)

- **iSCSI Name**—Enter the iSCSI name of an intended boot target. (Depending on your configuration, it might not be necessary to configure this option.)

- **DHCP Dynamic iSCSI Boot Configuration**—Select this option to have the initiator acquire its IP address from a DHCP server. The initiator also acquires all boot parameters for an iSCSI target if the boot parameters are specified on the DHCP server. This setting requires advanced users to manually input vendor-specific information on the DHCP server and in the **Vendor ID** field of the DHCP Boot Settings area.

  - ❑ **Vendor ID** is the identification string sent to DHCP server.

  - ❑ **Alternate Client ID** provides an alternative for DHCP clients to use instead of the **dhcp-client-identifier** option. The alternate client ID is a unique identifier string sent to the DHCP server. It provides an additional filtering option to set device-specific boot parameters. If this field is not set, a default identifier string consisting of the devices MAC address is used.

  - ❑ **Configure Parameters** for the secondary adapter.

  - ❑ **Select Host Adapter** if you have multiple adapters in your system, select, view, and configure the settings of a specific adapter.

> **NOTE**
>
> These options apply only to disk devices; they do not apply to tape drives and other nondisk devices.

Configuring the iSCSI Boot Parameters

This section discusses how to configure a QLogic iSCSI adapter for booting from a SAN. On a system set up to boot from a SAN, the system BIOS must designate the iSCSI adapter as the boot controller. To locate the target boot LUN, you must enable the QLogic BIOS on the iSCSI adapter.

To enable the QLogic iSCSI adapter to boot from a SAN:

1. During server POST, press CTRL+Q to enter the QLogic iSCSI *Fast!*UTIL BIOS.

2. Select the I/O port to configure.

   By default, the **Adapter Boot** mode is set to **Disable**.

3. From the *Fast!*UTIL Options menu, select **Configuration Settings** and then select **iSCSI Boot Settings**.

4. Before you can set SendTargets, set the **Adapter Boot** mode to **Manual**.

5. Select **Primary Boot Device Settings**.

6. Enter the discovery **Target IP** and **Target Port**.

7. You can leave the **Boot LUN** and **iSCSI Name** fields blank only if one iSCSI target and one LUN are at the specified address from which to boot. Otherwise, you must specify these fields to ensure that you do not boot from a volume for some other system. After the target storage system is reached, these fields are populated after a rescan.

8. Save your changes.

9. From the **iSCSI Boot Settings** menu, select the primary boot device. An auto rescan of the host bus adapter finds new target LUNs.

10. Select the iSCSI target.

> **NOTE**
>
> If more than one LUN exists within the target, you can choose a specific LUN ID by pressing ENTER after you locate the iSCSI device. Return to the **Primary Boot Device Setting** menu. After the rescan, the **Boot LUN** and **iSCSI Name** fields are populated.

11. Change the value of **Boot LUN** to the LUN ID you want.

12. Save your changes and restart the system.

For additional details and the latest information on QLogic adapter configuration settings, see the QLogic adapter readme file on the QLogic Web site.

For information on configuring iSCSI boot parameters, see and .

Boot Protocol Configuration

Boot protocol primary and alternate boot device settings include the following:

- **Security Settings**—Press ENTER to access Primary Boot Security Settings. Press ENTER to enable or disable CHAP and bidirectional CHAP and to configure the CHAP name and CHAP secret. (Depending on your configuration, it might not be necessary to configure this option.)

- **Target IP**—When the Adapter Boot Mode is set to **Manual**, you must use this field to specify the IP address of the intended iSCSI boot target.

- **Target Port**—Enter the port used by the target.

  By default, iSCSI uses port 3260. Do not modify this value unless it is necessary.

- **Boot LUN**—Enter the LUN of the boot device. (Depending on your configuration, it might not be necessary to change this option.

- **iSCSI Name**—Enter the iSCSI name of an intended boot target. (Depending on your configuration, it might not be necessary to configure this option).

Static iSCSI Boot Configuration

Use the QLogic iSCSI *Fast!*UTIL BIOS to configure static iSCSI boot parameters. Refer to .

Dynamic iSCSI Boot Configuration

Dynamic iSCSI boot configuration primary and alternate boot device settings include the following:

- **Security Settings**—Press ENTER to access Primary Boot Security Settings. Press ENTER to enable or disable CHAP and bidirectional CHAP and to configure the CHAP name and CHAP secret. (Depending on your configuration, it may not be necessary to configure this option.)

- **DHCP Dynamic iSCSI Boot Configuration**—Select this option to have the initiator acquire its IP address from a DHCP server. The initiator also acquires all boot parameters for an iSCSI target if the boot parameters are specified on the DHCP server This setting requires advanced users to manually input vendor-specific information the DHCP server and in the **Vendor ID** field of the DHCP Boot Settings area.

- **Vendor ID** specifies the identification string sent to DHCP server.

- **Alternate Client ID** is the additional filtering option to set the boot device.

- **Configure Parameters** for the secondary adapter.

- **Select Host Adapter** if you have multiple adapters in your system, select, view, and configure the settings of a specific adapter.

Configuring Parameters for a Secondary Adapter

If login to the primary boot target fails, the BIOS should attempt to log in to the secondary target using the same technique. BIOS attempts to log in to boot targets configured on different ports, depending on their configuration. iSCSI ports can reside on physical interfaces and may exist on separate adapters.

When any combination of these adapters (QMD8262-k, QLE8262, or QME8262-k) are present in a system, *Fast!*UTIL lists all of the iSCSI interface ports. Select a specific port of a secondary adapter to configure its parameters.

Configuring QLogic iSCSI Boot

To configure QLogic iSCSI boot:

1. During POST, press CTRL+Q to enter the QLogic iSCSI BIOS *Fast!*UTIL utility.

    Figure 5-21 shows the main Options menu.
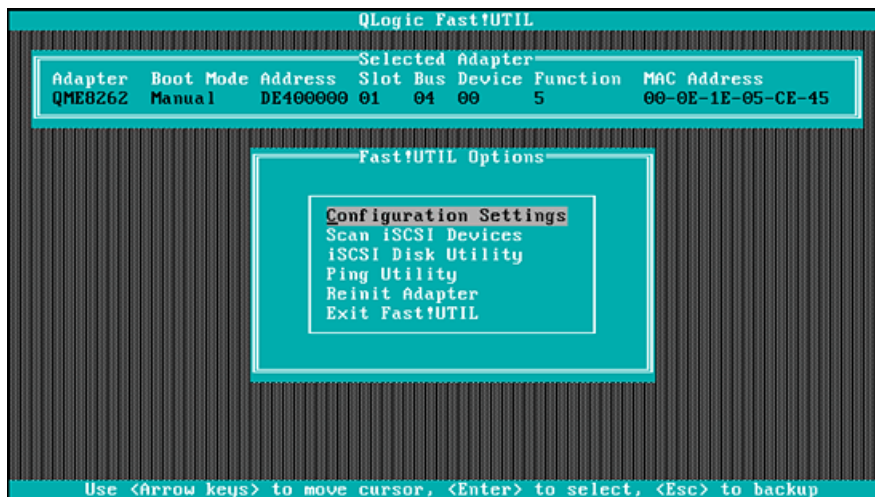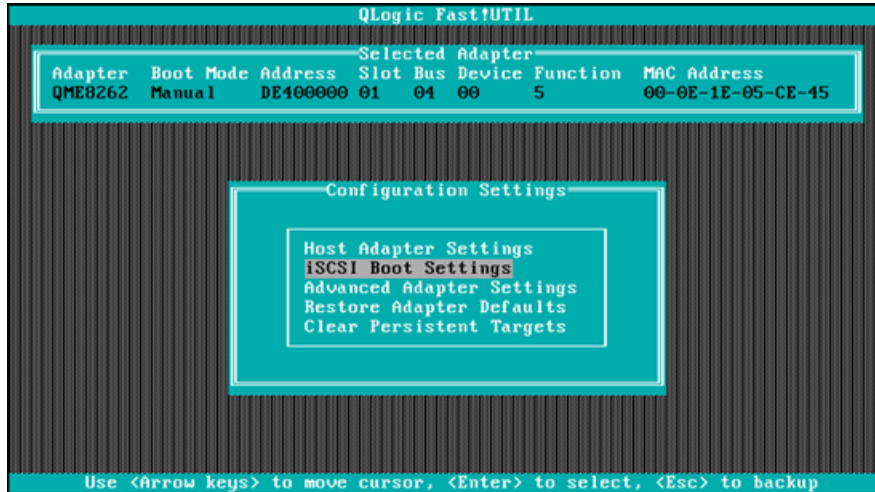

Figure 5-21. Fast!UTIL: Options Menu

2. On the Options menu, select **Configuration Settings**.

    The Configuration Settings window opens, as shown in Figure 5-22.

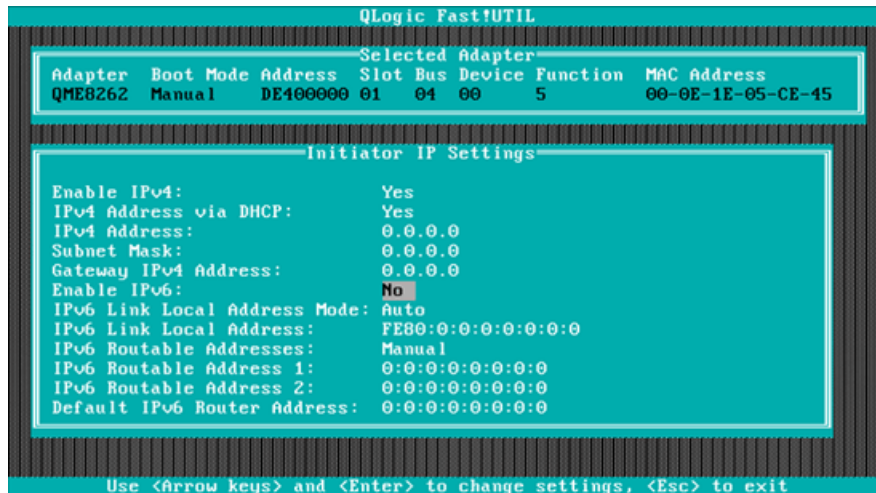***Figure 5-22. Fast!UTIL: Configuration Settings Window***

3. On the Configuration Settings menu, select **Host Adapter Settings**.

    The **Host Adapter Settings** window opens, as shown in Figure 5-23.

Figure 5-23. Fast!UTIL: Host Adapter Settings Window

4. Select **Initiator IP Settings**.

   The Initiator IP Settings window opens, as shown in Figure 5-24.

*Figure 5-24. Fast!UTIL: Initiator IP Settings Window*

5. Specify the IP version(s) by setting Enable IPv4 and Enable IPv6 as needed (you can enable IPv4 only, IPv6 only, or both), and then specify additional information as needed for the selected IP version(s):

   ❑ For IPv4: Initiator IP address, subnet, and gateway

   ❑ For IPv6: Initiator link local and routable addresses and default router address

   When you are finished updating the settings, return to the Host Adapter Settings window (see Figure 5-23).

6. (Optional) On the Host Adapter Settings window, specify the **Initiator iSCSI Name**.

   This option is not dependent upon DHCP.

7. Return to the Configuration Settings menu and then select **iSCSI Boot Settings** to configure the target settings shown in Figure 5-25.
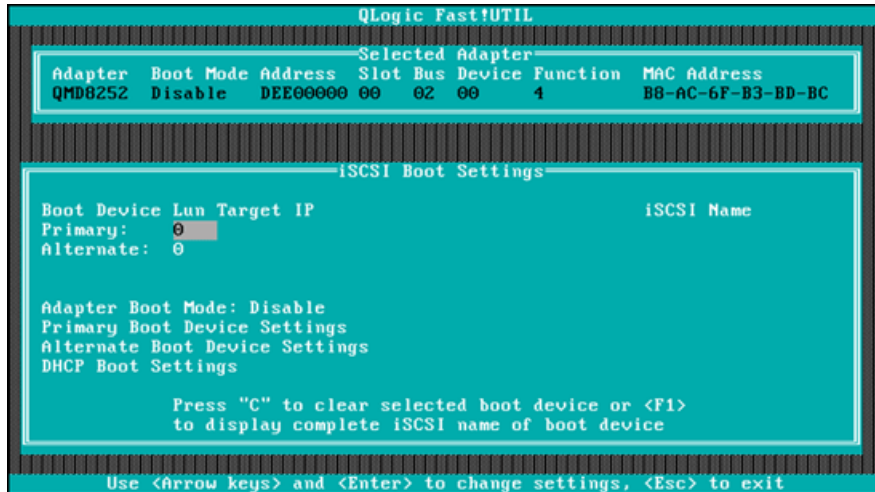
Figure 5-25. Fast!UTIL: iSCSI Boot Settings Window

    a. On the iSCSI Boot Settings window, select **Adapter Boot Mode** and set it to **Manual**.

    b. On the iSCSI Boot Settings window, select **Primary Boot Device Settings**.

8. On the Primary Boot Device Settings window (see Figure 5-26), specify the target parameters.
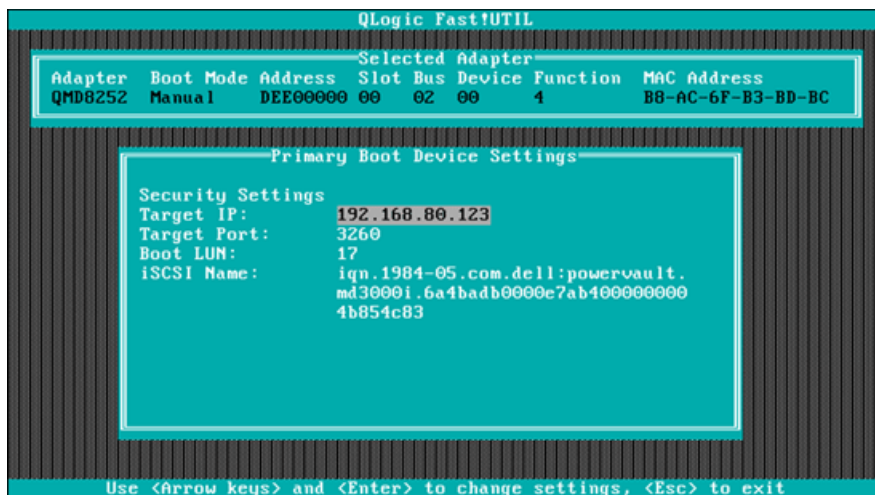
*Figure 5-26. Fast!UTIL: Primary Boot Device Settings Window*

a.   To scan for the specified target, highlight the primary LUN **Target IP** and then press ENTER.

b.   Select the target from the list of discovered targets on the **Select iSCSI Device** window, as shown in Figure 5-27.
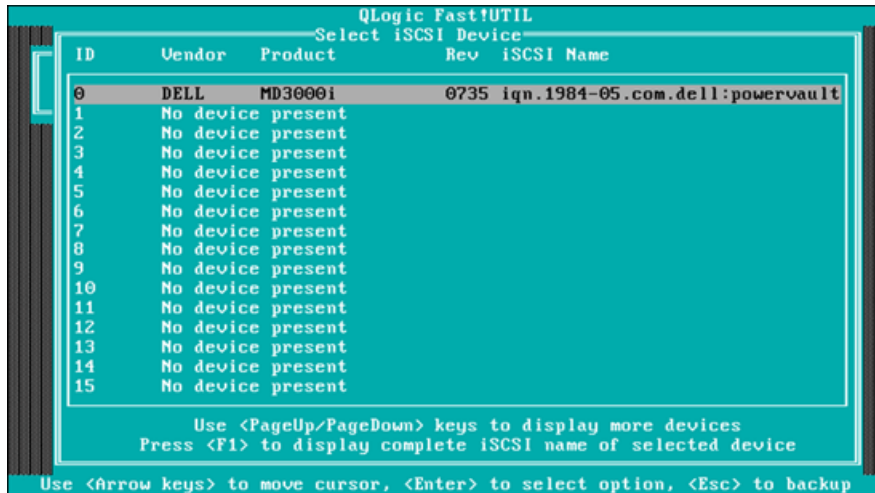

Figure 5-27. Fast!UTIL: Select iSCSI Device Window

c.   On the **Select LUN** window, select the LUN to set the target as the primary iSCSI boot device.

9.   Press ESC and then select **Save changes**.

10. Press ESC again and then select **Reboot System**.

After rebooting, the iSCSI BIOS loads for the target configured previously in the QLogic iSCSI *Fast!*UTIL BIOS utility.

11. During POST, press F2 to enter the Dell System Setup menu, as shown in Figure 5-28.
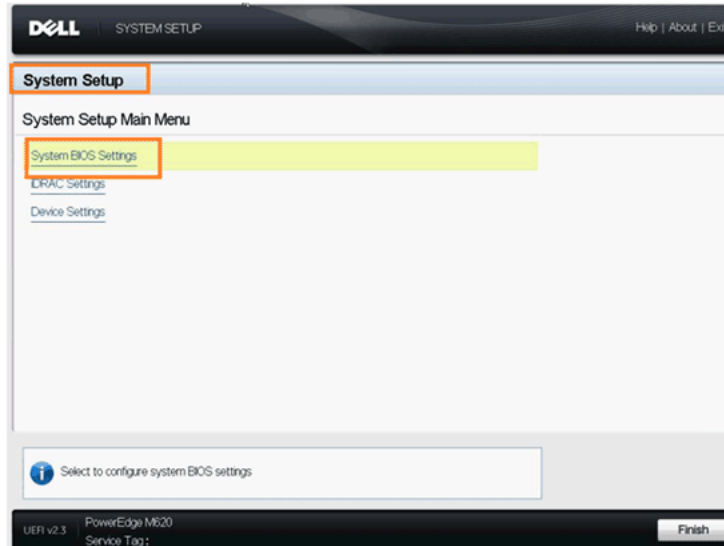
Figure 5-28. Dell System Setup

12. Select **System BIOS Settings**, then select **Boot Settings**, then select **BIOS Boot Settings**, and then select **Hard-Disk Drive Sequence**, as shown in Figure 5-29.
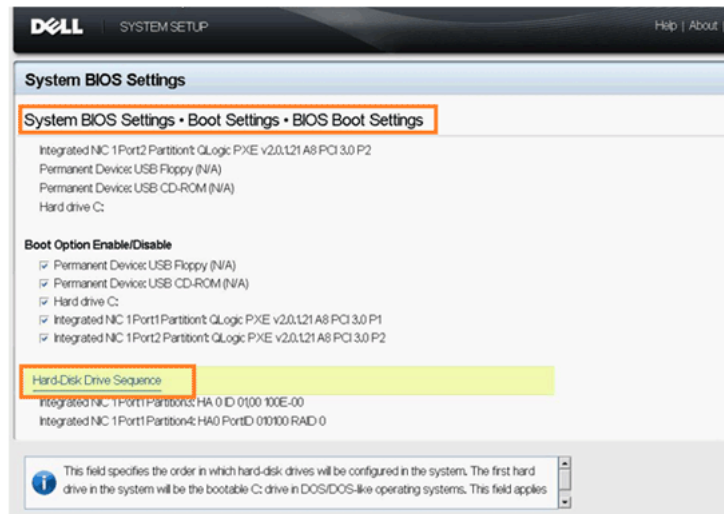
Figure 5-29. Selecting iSCSI Boot Sequence

13. In the pop-up window, use the UP ARROW and DOWN ARROW or the **+** and **–** buttons to move the iSCSI target to the top of the list, as shown in Figure 5-30 (where the iSCSI target is configured on Port 1, Partition 3). Then click **OK**.
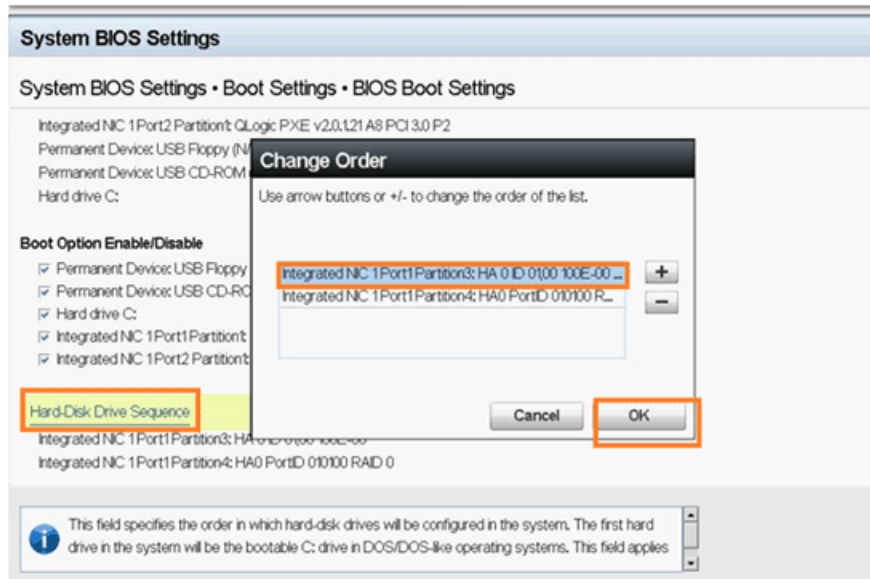
Figure 5-30. Setting the iSCSI Boot Sequence

14. Select **Save changes and exit**.

15. Follow the manufacturer's OS installation instructions.

Booting

The iSCSI *Fast!*UTIL BIOS includes an iSCSI Boot Settings option, which is accessed from the Configuration Settings menu. When you enable this option, you can specify the drive from which you want to boot. After this is enabled, the option forces the system to boot from the selected iSCSI drive, subject to the settings in system BIOS (QLogic BIOS does not override system BIOS). When you disable this option, the system looks for a boot device (as selected in the system BIOS).

iBFT Boot Setup

For an alternate method of iSCSI boot from SAN, use the fields in the iBFT. iBFT is a component of the *Advanced Configuration and Power Interface Specification* 3.0b standard that provides operating systems a standard way to boot from software-initiated iSCSI protocol. To view the iBFT specification, visit the following URL:

http://www.microsoft.com/whdc/system/platform/firmware/ibft.mspx

This section provides details for setting up the adapter using iBFT, including:
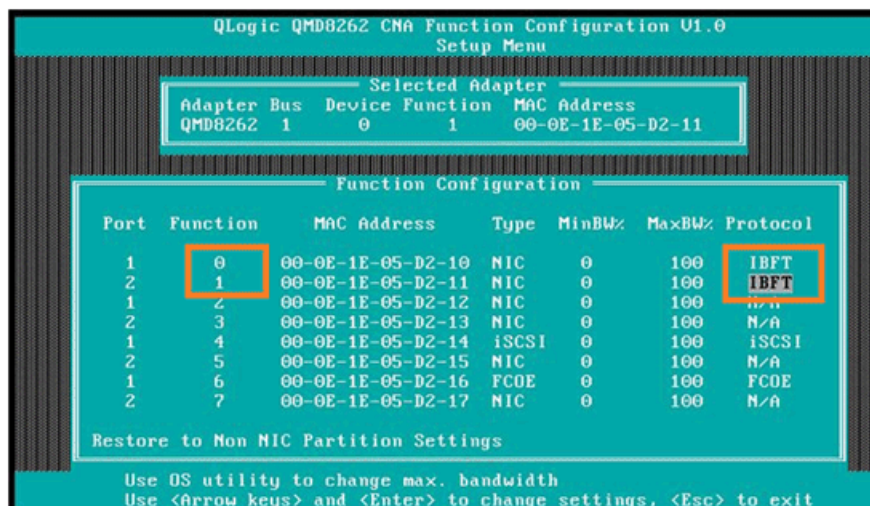
■ Enabling iBFT Boot
■ Booting to a Target Disk

Enabling iBFT Boot

Follow these steps to enable iBFT boot in the Dell BIOS system.

1. Boot the server, and when the NIC boot code banner appears, press CTRL+Q.

   The QLogic CNA Function Configuration utility opens, as shown in Figure 5-31.

Figure 5-31. CNA Configuration Utility

2. Check that the protocol for functions 0 and 1 is set to **iBFT**. If necessary, change the setting(s), as shown in Figure 5-32, and then press ENTER.
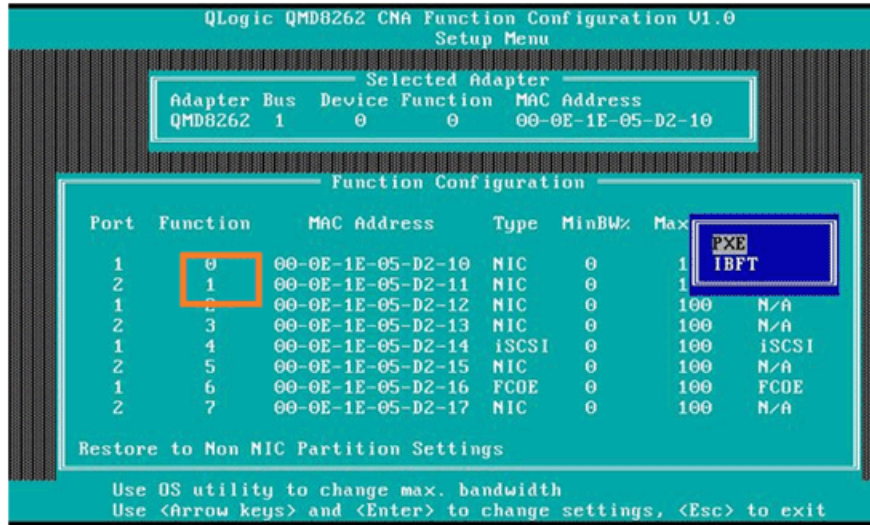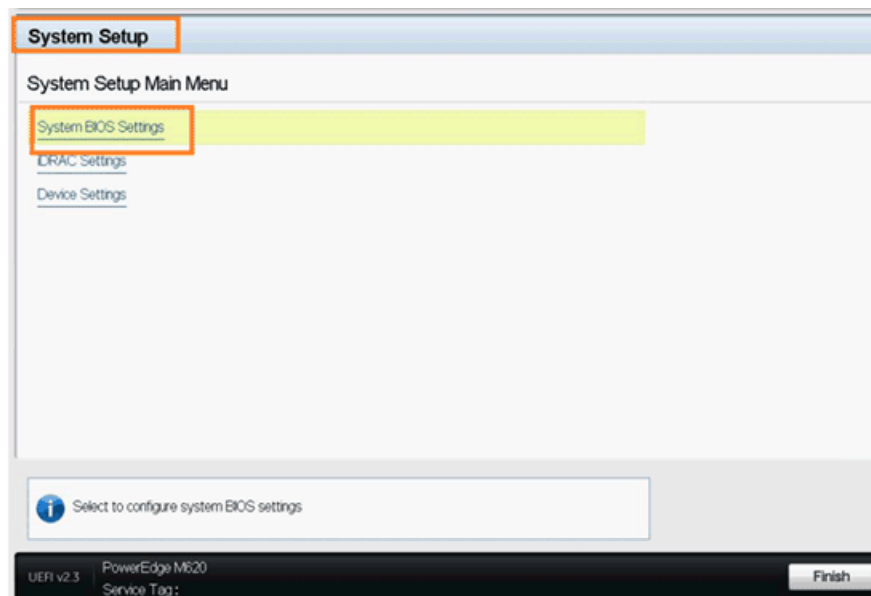

Figure 5-32. Enabling iBFT Boot

3. Press ESC and save the settings.
4. Reboot the system.

Booting to a Target Disk

To boot to the target disk, see the boot target vendor's instructions for the hardware setup.

1. During POST, press F2 to enter the Dell System Setup menu, as shown in .

Figure 5-33. Dell System Setup

2.   Select **System BIOS Settings**, then select **Boot Settings**, then select **BIOS Boot Settings**, and then select **Hard-Disk Drive Sequence**, as shown in Figure 5-34.
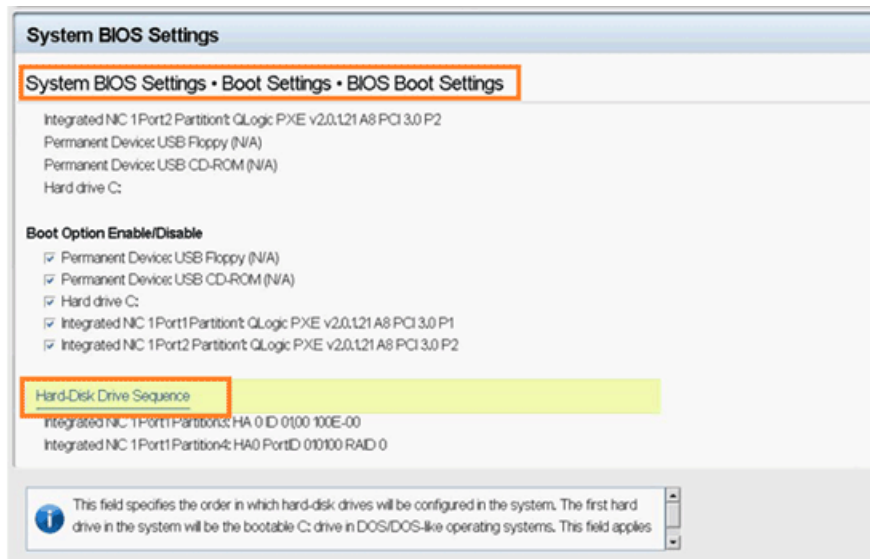
Figure 5-34. Selecting iSCSI Boot Sequence

3.   In the pop-up window, use the UP ARROW and DOWN ARROW or the **+** and **−** buttons to move the iSCSI target to the top of the list, as shown in Figure 5-35 (where the iSCSI target is configured on Port 1, Partition 3). Then click **OK**.
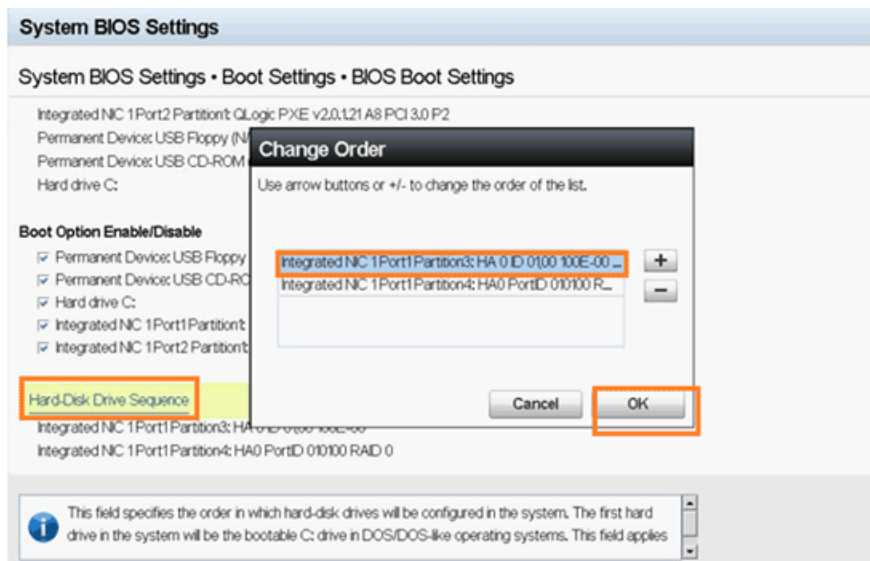
Figure 5-35. Setting the iSCSI Boot Sequence

4.   Select **Save changes and exit**.

5.   Reboot the system.

6.   The Option Rom shows the iSCSI target login information, as shown in .


```
Adapter1 Port1 F0 0200:8020 00000003 Enabled  PXE  MAC=84 2B 2B E2 71 34
Adapter1 Port2 F1 0200:8020 00000003 Enabled  PXE  MAC=84 2B 2B E2 71 35

Press Ctrl+Q to enter setup menu
.....
Connecting to iSCSI target with NIC B8:AC:6F:B4:1F:40

         IP Address: 192.168.90.144
        Subnet Mask: 255.255.240.0
    Default Gateway: 0.0.0.0
     Initiator Name: iqn.2000-04.com.qlogic:isp82xx.B8AC6FB41F40.0
        Target Name: iqn.1984-05.com.dell:powervault.md3000i.6a4badb0000e7ab400000
0004b854c83
  Target IP Address: 192.168.95.121
    Target TCP Port: 3260
    Target Boot Lun: 0

Log into iSCSI target...Succeeded


_____

Adapter Port   Lun    Vendor   Product       Product
Number  Number Number ID       ID            Revision
0       1      0      DELL     MD3000i       0735
```

Figure 5-36. Connecting to the iSCSI Target

7.   Continue with OS installation (refer to the OS documentation).

DHCP Boot Setup (iSCSI)

To configure the DHCP server to support iSCSI boot, first ensure that your DHCP server is set up and then refer to the following procedure.

> **NOTE**
>
> This release does not support DHCP iSCSI boot for IPv6. Refer to future readme and release notes for IPv6 support notification.

Configuring DHCP iSCSI Boot for IPv4

> **NOTE**
>
> Although this section describes how to configure a Windows 2003 or Windows 2008 DHCP server and QLogic adapter to allow DHCP to work, additional steps are required. For example, you must also configure the IP addresses of the Ethernet interfaces on the DHCP server, configure the iSCSI storage box, and configure the network. These configuration steps are outside the scope of this guide.

To configure the QLogic adapter to obtain all of its IP and iSCSI initiator and target settings from the DHCP server:

1. Ensure that the latest BIOS and firmware are installed.

2. Reboot the server, and when the QLogic iSCSI BIOS screen appears, press CTRL+Q to enter the adapter BIOS utility, *Fast!*UTIL.

3. On the initial Select Host Adapter screen, select the adapter ports you want to configure for SAN boot.

   These ports are the same MAC addresses for which you created reservations on the DHCP server.

4. Press ENTER.

   If both ports are to be configured, perform the following steps for the first port, complete its configuration, and then repeat the steps for the second port.

5. On the *Fast!*UTIL Options screen, select **Configuration Settings** and then press ENTER.

6. On the Configuration Settings screen, select **iSCSI Boot Settings** and then press ENTER.

7. On the iSCSI Boot Settings screen, select **Adapter Boot Mode** and then press ENTER.

8. When presented with the various Boot Modes, select **DHCP using VendorID** and then press ENTER.

9. Select **DHCP Boot Settings** and then press ENTER.

10. On the DHCP Boot Settings screen, select **Vendor ID** and then press ENTER.

11. Enter the **Vendor ID** (class) that you defined earlier in the DHCP server configuration steps and then press ENTER.

   The vendor ID name is case sensitive and is limited to 10 characters in length.

12. Press ESC until the Configuration Settings Modified screen is displayed.

13. Select **Save changes** and then press ENTER.

14. If all components are configured correctly, you can either:

   ❑ Reinitialize the adapter from the BIOS

   ❑ Reboot the server

The adapter can now obtain all of its IP and initiator and target information from the DHCP server.

DHCP Option 17, Root Path

DHCP option 17 specifies the path name that contains the client's root disk. The path is formatted as a character string consisting of characters from the network virtual terminal (NVT) ASCII character set. The data is formatted as a string using the DHCP Root Path Option (Option 17):

```
"iscsi:"<servername | serveripaddress>":"<protocol>":"<port>":"
<LUN>":"<targetname>
```

DHCP Option 43 (Adding Vendor Options)

Define a vendor class called "QLogic" as part of the "Default User Class" and then add the iSCSI specific information. Two of these options—201 and 202—contain iSCSI boot target information for primary and secondary targets, respectively. BIOS must parse the Vendor Specific Information buffer for the Primary Boot Target (Option 201). The third option, iSCSI Initiator (Option 203), contains the iSCSI initiator IQN information.

DHCP Vendor Class Option 201, Primary Boot Target IQN and Boot Parameters

Format the data as a string using the DHCP vendor-defined Primary Boot Target IQN and Boot Parameters Option (Option 201):

```
"iscsi:"<serveripaddress>":"<protocol>":"<port>":"<LUN>":"
<targetname>
```

Example string value (no spaces):

```
iscsi:192.168.95.121:6:3260:7:iqn.1984-05.com.dell:powervault.md30
00i.6a4badb0000e7ab4000000004b854c83
```

Parameters

DHCP Vendor Class Option 202, Secondary Boot Target IQN and Boot

Format the data as a string using the DHCP vendor-defined Secondary Boot Target IQN and Boot Parameters Option (Option 202):

```
"iscsi:"<serveripaddress>":"<protocol>":"<port>":"<LUN>":"
<targetname>
```

Example string value (no spaces):

```
iscsi:192.168.95.109:6:3260:9:iqn.1992-01.com.lsi:1535.600a0b80006
7fe9a000000004b9fd854
```

DHCP Vendor Class Option 203, Initiator IQN

Format the data as a string using the DHCP vendor-defined Initiator IQN Option (Option 203):

```
"<initiatorname>"
```

Example string value (no spaces):

```
iqn.2000-04.com.qlogic:isp8214.00e1e049e02.4
```

A Troubleshooting

This appendix provides the following troubleshooting information:

- Diagnosing Problems
- NIC Troubleshooting
- iSCSI Troubleshooting
- FCoE Troubleshooting
- ESX Troubleshooting

Diagnosing Problems

Network activity indicators and diagnostic utilities help you to verify that the hardware and software are working properly. If the installed adapter cannot communicate over the network, the flowcharts shown in this appendix can help diagnose the problem with the adapter.
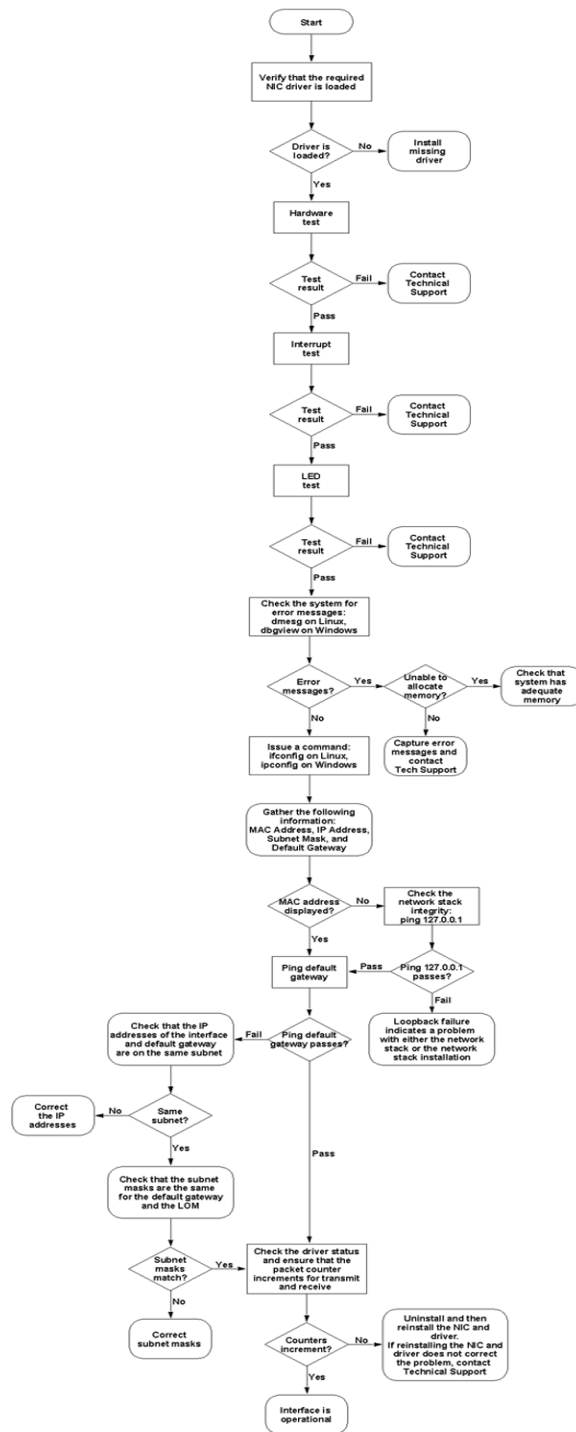
NIC Troubleshooting

Figure A-1. NIC Diagnostics Flowchart
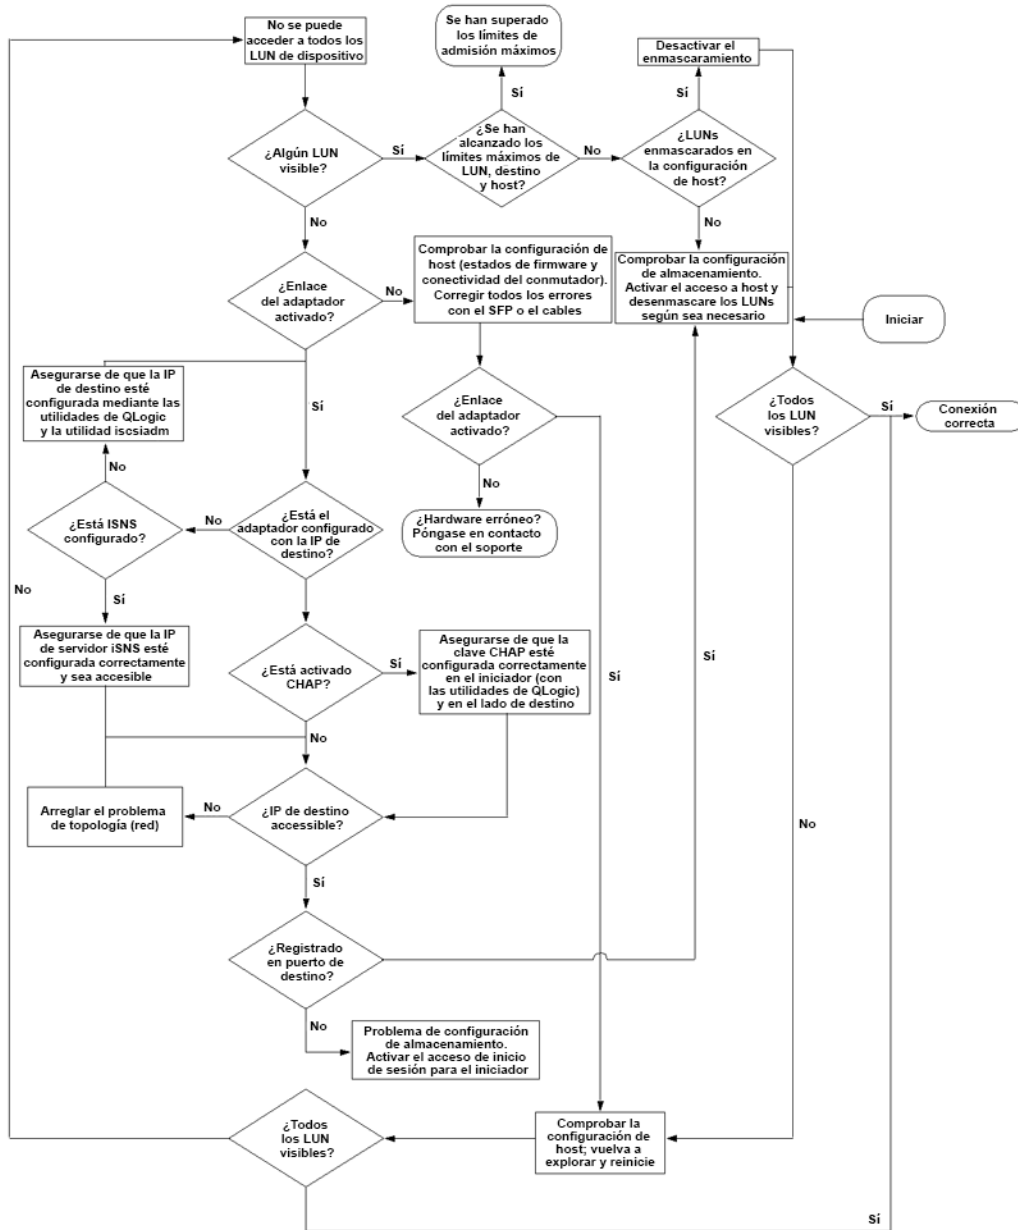
iSCSI Troubleshooting

Figure A-2. iSCSI Diagnostics Flowchart

FCoE Troubleshooting

> **NOTE**
>
> If most of the IP packet traffic is not TCP or UDP, the FCoE FIP session might be dropped. If you experience this problem, turn off RSS.
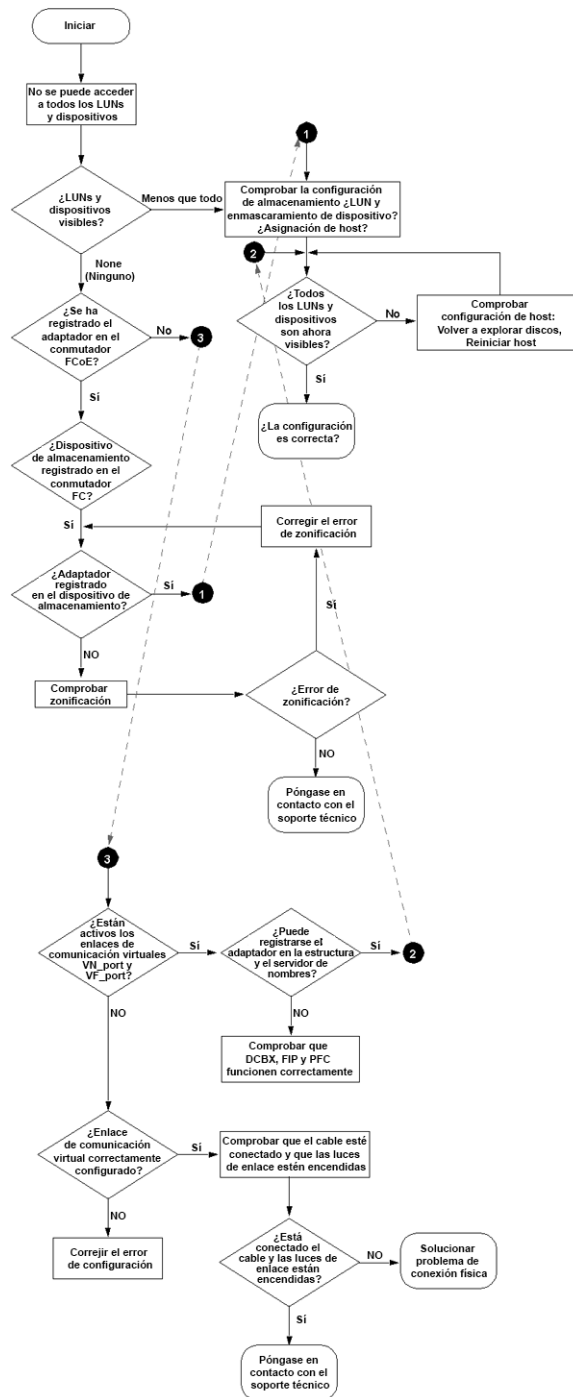
Figure A-3. FCoE Diagnostics Flowchart

ESX Troubleshooting

For debugging and troubleshooting networking issues on ESX, refer to the VMware document, *VI3 Networking: Advanced Troubleshooting*, located here:

http://www.vmware.com/files/pdf/technology/vi_networking_adv_troubleshooting.pdf

If the troubleshooting procedures in this document do not resolve the problem, please contact Dell for technical assistance (refer to the "Getting Help" section in your Dell system documentation).

B   Specifications

This appendix provides specifications for the following products:

- QMD8262-k Specifications
- QLE8262 Specifications
- QME8262-k Specifications

QMD8262-k Specifications

- Physical Characteristics
- Power Requirements
- Standards Specifications
- Interface Specifications
- Environmental Specifications

Physical Characteristics

Table B-1. Physical Characteristics

Adapter	Description
Type	Blade network daughter card
Length	3.00 inches
Width	2.45 inches

Power Requirements

Table B-2. Power Requirements

Voltage Rail	Voltage	Current
12V	12.0V	2mA
12V Aux	12.0V	0.784A
3.3V	N/A	N/A
3.3V Aux	3.3V	12.5mA
1.0/1.2V Core	1.0V	9.860A

Standards Specifications

The QMD8262-k adapter supports the following standards specifications:

- IEEE: 802.3ae (10Gb Ethernet)
- IEEE: 8021q (VLAN)
- IEEE: 802.3ad (Link Aggregation)
- IEEE: 802.1p (Priority Encoding)
- IEEE: 802.3x (Flow Control)
- IEEE: 802.1Qbb (Priority Based Flow Control)
- IEEE: 802.1Qaz (Enhanced Transmission Selection)
- IPv4 Specification (RFC791)
- IPv6 Specification (RFC2460)
- TCP/UDP Specification (RFC793/768)
- ARP Specification (RFC826)
- SCSI-3 Fibre Channel Protocol (SCSI-FCP)

- Fibre Channel Tape (FC-TAPE) Profile
- SCSI Fibre Channel Protocol-2 (FCP-2)
- Second Generation FC Generic Services (FC-GS-2)
- Third Generation FC Generic Services (FC-GS-3)

Interface Specifications

Table B-3. Interface Specifications

Port Type	Media
10G-BASE-KR	Dell PE M1000e KR Midplane Revision 1.1

Environmental Specifications

Table B-4. Environmental Specifications

Condition	Operating	Non-Operating
Temperature Ranges (for Altitude ≤900 m or 2952.75 ft)	10°C to 35°C (50°F to 95°F)	–40°C to 65°C (–40°F to 149°F)
Temperature Ranges (for Altitude >900 m or 2952.75 ft)	10°C to Note[a] °C (50°F to Note[b] °F)	–40°C to 65°C (–40°F to 149°F)
Temperature Gradient Max. per 60 min.	10 °C	20 °C
Humidity Percent Ranges—Noncondensing	20% to 80%* (Max. Wet bulb temp.= 29°C)[c]	5% to 95%+ (Max. Wet bulb temp. = 38°C)[c]
Humidity Gradient Max. per 60 min.	10%	10%
Altitude Ranges—Low Limits	–15.2m (–50ft)	–15.2m (–50ft)
Altitude Ranges—High Limits	3,048m (10,000ft)	10,668m (35,000ft)
Airborne Contaminants—ISA-71 Level	G1[c]	G1[c]

[a] Use the following formulas to calculate the maximum operating temperature (in °C) for a specific altitude. Use the first formula if the altitude is stated in meters and the second formula if the altitude is stated in feet.

$$35 - \frac{(MaximumAltitude_{meters} - 900)}{300} °C \quad or$$
$$35 - \frac{(MaximumAltitude_{feet} - 2952.75)}{984.25} °C$$

[b] Use the following formulas to calculate the maximum operating temperature (in °F) for a specific altitude. Use the first formula if the altitude is stated in meters and the second formula if the altitude is stated in feet.

$$95 - \frac{(MaximumAltitude_{meters} - 900) \times 1.8}{300} °F \quad or$$
$$95 - \frac{(MaximumAltitude_{feet} - 2952.75) \times 1.8}{984.25} °F$$

[c] Maximum corrosive contaminant levels measured at ≤50% relative humidity; see Table 3 in ISA-71.04-1985.

QLE8262 Specifications

- Physical Characteristics
- Power Requirements
- Standards Specifications
- Interface Specifications
- Environmental Specifications

Physical Characteristics

Table B-5. Physical Characteristics

Adapter	Description
Type	Low-profile PCIe card
Length	6.6 inches
Width	2.54 inches

Power Requirements

Table B-6. Power Requirements

Voltage Rail	Voltage	Current
12V	12V	1.4A
3.3V	3.3V	0A
3.3V AUX	3.3V	5mA

Standards Specifications

For the standards specifications for the QLE8262 adapter, refer to "Standards Specifications" on page 228.

Interface Specifications

Table B-7. Interface Specifications

Port Type	Media
10G-BASE-SFP+	SFF-8431 Specification for SFP+, Revision 3.2

Environmental Specifications

For the environmental specifications for the QLE8262 adapter, refer to Table B-4 on page 230.

QME8262-k Specifications

- Physical Characteristics
- Power Requirements
- Standards Specifications
- Interface Specifications
- Environmental Specifications

Physical Characteristics

Table B-8. Physical Characteristics

Adapter	Description
Type	Mezzanine card
Length	3.307 inches
Width	3.465 inches

Power Requirements

Table B-9. Power Requirements

Voltage Rail	Voltage	Current
12V	12V	1.3A
3.3V	3.3V	0A
3.3V AUX	3.3V	1mA

Standards Specifications

For the standards specifications for the QME8262-k adapter, refer to "Standards Specifications" on page 228.

Interface Specifications

Table B-10. Interface Specifications

Port Type	Media
10G-BASE-KR	Dell PE M1000e KR Midplane Revision 1.1

Environmental Specifications

For the environmental specifications for the QME8262-k adapter, refer to Table B-4 on page 230.

C QConvergeConsole GUI

This appendix provides the following information about the QConvergeConsole GUI:

- Introduction to QConvergeConsole
- Downloading QConvergeConsole Documentation
- Downloading and Installing Management Agents
- Installing the QConvergeConsole GUI
- What Is in the QConvergeConsole Help System

NOTE

For information on installing the QConvergeConsole GUI, refer to the *QConvergeConsole GUI Installation Guide*. All procedural information for the QConvergeConsole GUI utility is covered in the utility's online help system.

Introduction to QConvergeConsole

The QConvergeConsole GUI is a Web-based client and server GUI management tool that provides centralized management and configuration of QLogic adapters within the entire network (LAN and SAN).

On the server side, the QConvergeConsole GUI runs as an Apache Tomcat™ application server. After the QConvergeConsole GUI is launched on the application server, you can connect to the QConvergeConsole GUI through a browser, either locally on the server or remotely from another computer. Your browser window becomes the client used to connect to servers that host the QLogic adapters and connected storage devices within the network.

The QConvergeConsole Web GUI displays the main window used to connect to servers that host QLogic adapters. Each host server can include a combination of the following QLogic adapters:

- 4Gb and 8Gb Fibre Channel Adapters (QL*x*24*xx* and QL*x*25*xx*)
- 8Gb Fibre Channel-to-PCIe Mezzanine Card (QME2572)
- 16Gb Fibre Channel Adapters (QLE266x-DEL)
- 16Gb Fibre Channel Mezzanine Card (QME2662-DEL)
- 10Gb Converged Network Adapter (QLE8152)
- 10GbE Blade Network Daughter Card (QMD8262)
- 10GbE Monolithic Server Standup Card (QLE8262)
- 10GbE Blade Mezzanine Card (QME8262-k)
- FabricCache™ Adapters (QLE10000)

Note the following:

- The QConvergeConsole GUI works with the Windows Server 2008, Windows Server 2012, Red Hat® Linux® Advanced Server and Enterprise Server, Novell® SUSE® Linux Enterprise Server (SLES®), and Oracle Linux—Unbreakable Enterprise Kernel (UEK). Multiple operating system support allows control of heterogeneous environments.

- The QConvergeConsole GUI runs on the following Web browsers: Google Chrome™, Mozilla® Firefox®, Microsoft® Internet Explorer®, Opera®, Apple® Safari®. For details on browser versions, see the QConvergeConsole *Readme*, which you can download at http://driverdownloads.qlogic.com.

NOTE

Some auxiliary help features, such as the Search and Index tools, may not work properly in every Web browser.

Downloading QConvergeConsole Documentation

To download the *QConvergeConsole GUI Installation Guide*, go to at
http://driverdownloads.qlogic.com.and click **Downloads**.

Downloading and Installing Management Agents

To manage the adapters on a local or remote host, the *management agents* (also
called *agents*) used by the host's adapters must already be installed on the host.

To connect to a remote host, at least one management agent used by any of the
host's adapters must be installed and running on the host. Only protocols whose
agents are installed and running are shown under the host's system tree node.

The management agents are:

- `netqlremote` for Ethernet Networking (NIC)
- `qlremote` for Fibre Channel and FCoE
- `iqlremote` for iSCSI
- `qlremote_fca` for FabricCache Adapter

The management agents are *not* automatically installed by the QConvergeConsole
GUI installer and must be downloaded and installed separately. You can download
the agents from the QLogic Web site Downloads page, or use the built-in Agent
Installer. See "Installing the Agents Using the Built-in Agent Installer" on page 236.
for more information.

Installing the Agents from the QLogic Web Site

To obtain the agents from the QLogic Web site and install them:

Windows and Linux (all versions):

1.   Go to the QLogic Downloads page at http://driverdownloads.qlogic.com and download the following for each adapter on the host server:

   ❑   SuperInstaller
   ❑   *Readme* and *Release Notes*

2.   Install the agents by running the SuperInstaller. Refer to the *Readme* and *Release Notes* for detailed information on running the SuperInstaller.

Installing the Agents Using the Built-in Agent Installer

To access the Agent Installer within QConvergeConsole:

1.   Click **Help** in the QConvergeConsole menu bar, and then select **Download Agent Installers**.

2.   Follow the on-screen prompts. For detailed installation instructions, refer to the QConvergeConsole online help (select **Help** from the menu bar, and then select **Browse Contents**; search for *agent installer*).

Installing the QConvergeConsole GUI

Refer to the installation procedure for your operating system.

- Installing QConvergeConsole in a Windows Environment
- Installing QConvergeConsole in a Linux Environment
- Installing QConvergeConsole in Silent Mode

Installing QConvergeConsole in a Windows Environment

The QConvergeConsole Installer for Windows is a self-extracting utility that installs QConvergeConsole and related files. Follow these steps to install QConvergeConsole on a Windows system.

NOTE

Before you begin, read the *QConvergeConsole Release Notes* and *Readme* documents for important, up-to-date product information.

To install QConvergeConsole in a Windows environment:

1. Locate the folder where you downloaded the installation file, and then double-click the `QConvergeConsole_Installer_<version>.exe` file.

2. If an Open File security dialog box appears, click **Run**.

   The InstallAnywhere utility prepares to install QConvergeConsole.

3. On the Tomcat Port Number Setup dialog box, enter a valid port number in the range 1025–65536 or accept the default port number, 8080. Then click **Next** to continue.

4. In the next dialog box, select **Yes** if you want a localhost-only installation. If you do not want a localhost-only installation, select **No**.

   NOTE

   The localhost-only option installs QConvergeConsole locally so that you must run it locally (remote connection is not possible). To disable the option, you must uninstall QConvergeConsole and then re-install it, selecting **No** in this step.

   ---

5. On the Pre-Installation Summary dialog box, verify that the information is correct and then click **Install** to proceed.

   During the installation, the installer notifies you of the status. A message box prompts you to confirm that you want to enable the SSL feature, which allows you to run QConvergeConsole in a secured environment. Enable SSL to allow data to be encrypted within a secure layer before passing it to and from the server.

6. To enable the SSL feature, click **Yes**. To disable SSL, click **No**.

7. On the Install Complete dialog box, click **Done** to exit the installer.

You have installed QConvergeConsole on your server.

Installing QConvergeConsole in a Linux Environment

You have the option of installing QConvergeConsole in a Linux environment using either a GUI or CLI method. To install from the CLI, see .

You can graphically install QConvergeConsole on a Linux host by running a Virtual Network Computing (VNC®) program on a Windows host. The TightVNC Viewer, a VNC program, can remotely control QConvergeConsole installation on a Linux host.

To install QConvergeConsole in a Linux environment:

1. Open the directory containing the installer file (one of the following):

    ```
    QConvergeConsole_Installer_Linux_x32_.bin
    QConvergeConsole_Installer_Linux_x64_.bin
    ```

2. In the terminal window, verify that the installer file has permission to execute by issuing the `ls -lt` command at the prompt.

3. To install QConvergeConsole, double-click the QConvergeConsole `.bin` file.

    A message box prompts whether you want to run or view the contents of the `.bin` file.

4. Click **Run in Terminal**.

    The installer extracts the files and installs QConvergeConsole in the `/opt/QLogic_Corporation/QConvergeConsole` directory.

    The InstallAnywhere window prepares to install QConvergeConsole.

5. On the installation Introduction dialog box, click **Next**.

6. On the Tomcat Port Number Setup dialog box, enter a valid port number in the range 1025–65536 or accept the default port number, 8080. Then click **Next** to continue.

7. In the next dialog box, select **Yes** if you want a localhost-only installation. If you do not want a localhost-only installation, select **No**.

> **NOTE**
>
> The localhost-only option installs QConvergeConsole locally so that you must run it locally (remote connection is not possible). To disable the option, you must uninstall QConvergeConsole and then re-install it, selecting **No** in this step.

8. On the Pre-Installation Summary dialog box, read the information, and then click **Install**.

    During the installation, the installer notifies you of the status.

9. On the Install Complete dialog box, click **Done** to exit the installer.

You have installed QConvergeConsole on your server.

Installing QConvergeConsole in Silent Mode

You can install the QConvergeConsole from a command prompt in silent mode.

To run the installer in silent mode, issue the following command:

```
<installer_file_name> -i silent
-DUSER_INPUT_PORT_NUMBER=<A_NUMBER>
```

What Is in the QConvergeConsole Help System

To access the QConvergeConsole help system while the GUI utility is running, click the **Help** menu and then click **Browse Contents**. The help system provides topics containing details of the following:

- **Getting Started** shows how to start using QConvergeConsole and the help system. It also describes how to initiate automatic e-mail notifications and set security on adapters installed for a selected host.

- **Managing Host Connections** describes how to connect to a host, create and use a host group, view host information, and disconnect from one or more host servers.

- **Displaying Host Information** describes how to view general host information.

- **Displaying Adapter Information** describes how to view general adapter information.

- **Displaying Device Information** shows how to view information about a device (disk or tape).

- **Using Reports** describes the different types of reports and how to generate, view, and save them using QConvergeConsole's report facility.

- **Managing Fibre Channel and FCoE Adapters** describes how to do the following for a Fibre Channel Adapter (2400/2500/2600 Series) or Converged Network Adapter (8100/8200/8300 Series):

  - ❑ Display and edit adapter-specific information for the selected adapter
  - ❑ View and modify adapter configuration and settings
  - ❑ Use the provided utilities and wizards to update the adapter

- **Managing Fibre Channel and FCoE Ports** describes how to do the following for a port on a Fibre Channel Adapter or a Converged Network Adapter:

  - ❑ View and modify the parameters and statistics for the selected port

  - ❑ Use the advanced utilities, manage virtual ports, and run diagnostics

- **Managing iSCSI Ports** shows and describes the parameters on the iSCSI port management pages for the 8200/8300 Series Converged Network Adapters.

- **Managing Ethernet (NIC) Ports** shows and describes the parameters and configuration options provided on the NIC port management pages for the 3200 Series Intelligent Ethernet Adapters, as well as the 8100/8200/8300 Series Converged Network Adapters.

- **Managing FabricCache Adapters and Ports** shows and describes how to display and edit information parameters for 10000 Series FabricCache Adapters (FCA) and ports, as well as how to configure port parameters.

- **Managing FabricCache Adapter Storage** shows and describes how to use the system tree on the Storage tab, how to display clusters, how to use the logical and physical views, how to display information for FCA cluster adapters, ports, SSDs, and LUNs, and how to configure LUN caching.

- **Using the Wizards** describes how to use QConvergeConsole's wizards to update the following: Flash, drivers, NIC parameters, Fibre Channel/FCoE parameters, parameters files, passwords, and agents.

- **Troubleshooting** describes tools you can use to resolve the problems that can occur when installing and using the adapter. It also provides the procedures for tracing QConvergeConsole GUI and agent activity, if directed to do so by your authorized service provider.

- **Frequently Asked Questions** provides the answers to frequently asked questions about QConvergeConsole.

- **Technical Support** explains how to obtain technical support, gives instructions for downloading software and documentation, and provides contact information.

D Regulatory Information

This appendix provides the following information for the QMD8262-k, QLE8262, and QME8262-k products:

- Warranty
- Regulatory and Compliance Information

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.

<div style="border:1px solid black; display:inline-block; padding:10px;">

Láser clase I

</div>

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
Produit der Laser Klasse 1	**Vorsicht**—Laserstrahlung der Klasse 1 bei geöffneter Abdeckung Direktes Ansehen mit optischen Instrumentenvermeiden
Luokan 1 Laserlaite	**Varoitus**—Luokan 1 lasersäteilyä, kun laite on auki Älä katso suoraan laitteeseen käyttämällä optisiainstrumenttej

Agency Certification

The following sections contain a summary of EMI and EMC test specifications performed on the models listed below to comply with emission, immunity, and product safety standards:

- QMD8262-k (CU0310419)
- QLE8262 (CU0310414)
- QME8262-k (CU0310410)

EMI and EMC Requirements

FCC Part 15 compliance: 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.

ICES-003 compliance: Class A

This Class A digital apparatus complies with Canadian ICES-003.Cet appareil numériqué de la classe A est conformé à la norme NMB-003 du Canada.

CE Mark 2004/108/EC EMC Directive compliance:

EN55022:2010 Class A

EN55024:2010
EN61000-3-2: Harmonic Current Emission
EN61000-3-3: Voltage Fluctuation and Flicker

Immunity Standards
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: 2010-04 Class A

AS/NZS CISPR22: Class A

KCC: Class A

Korea RRA Class A Certified

Product Name/Model: Fibre Channel 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/telecom-munications 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: ProductLine

UL60950-1 (2nd Edition), 2007-03-3-27
UL CSA C22.2 60950-1-07 (2nd Edition)

Use only with listed ITE or equivalent.

Complies with 21 CFR 1040.10 and 1040.11.

2006/95/EC low voltage directive: ProductLine

TUV:
TUV EN60950-1:2006+A11+A1+A12

CB Certified to IEC 60950-1 2nd Edition