User's Guide

Converged Network Adapter

8200 and 3200 Series

SN0054671-00 A

Information furnished in this manual is believed to be accurate and reliable. However, QLogic Corporation assumes no responsibility for its use, nor for any infringements of patents or other rights of third parties which may result from its use. QLogic Corporation reserves the right to change product specifications at any time without notice. Applications described in this document for any of these products are for illustrative purposes only. QLogic Corporation makes no representation nor warranty that such applications are suitable for the specified use without further testing or modification. QLogic Corporation assumes no responsibility for any errors that may appear in this document.

Document Revision History	
Revision A, March 31, 2011	
Changes	**Sections Affected**
Initial release	All

Table of Contents

4 Configuring FCoE

List of Figures

List of Tables

Preface

This guide provides detailed instructions on the installation, configuration, and troubleshooting of 8200 and 3200 Series Adapters for Windows®, Linux®, and VMware®. It also provides details on the use of QLogic adapter features to enhance the value of server virtualization using VMware ESX/ESXi 4.0. Such features include virtual adapter configuration using N_Port ID virtualization (NPIV) and boot-from-SAN configuration.

What's in This Guide

This guide is organized into the following sections and appendices:

- This Preface describes the intended audience, related materials, document conventions used, license agreements, technical support, and legal notices.
- The Quick Start section provides high-level hardware and software installation instructions for advanced users.
- The Product Overview provides a product introduction and specifications; information about the multi-boot image (CNA Function Configuration Utility); and detailed installation instructions for the SuperInstaller and vCenter Plug-in for ESX/ESXi.
- The Configuring NIC section describes installing the NIC driver and agent across operating systems (OSs); configuring the PXE boot; how to use driver parameters, VLANs, and teaming/bonding; and troubleshooting NIC issues.
- The Configuring iSCSI section describes installing the iSCSI driver and agent across OSs; configuring iSCSI; how to use the iSCSI initiator, ISNS, and iSCSI boot; and troubleshooting iSCSI issues.
- The Configuring FCoE section describes installing the Fibre Channel over Ethernet (FCoE) driver and agent across OSs; setting HBA parameters; configuring the boot device, NPIV, driver parameters, and selective LUNs; and troubleshooting FCoE issues.
- The Glossary describes many of the terms used in this guide.

Intended Audience

This guide is intended for those responsible for deploying QLogic Fibre Channel and Converged Network Adapters on Windows, Linux, and VMware: users ranging from end users, such as data center managers and system administrators, to the test and development community.

Related Materials

For additional information, refer to the 8200 and 3200 Series Adapters Read Me and Release Notes files, the *QConvergeConsole User's Guide*, and the *QConvergeConsole CLI User's Guide*, available on the QLogic Web site, Downloads page: http:\\driverdownloads.qlogic.com.

Documentation Conventions

This guide uses the following documentation conventions:

- **NOTE:** provides additional information.
- ***CAUTION!*** indicates the presence of a hazard that has the potential of causing damage to data or equipment.
- ***WARNING!!*** indicates the presence of a hazard that has the potential of causing personal injury.
- Text in blue font indicates a hyperlink (jump) to a figure, table, or section in this guide, and links to Web sites are shown in underlined blue. For example:
  - Table 9-2 lists problems related to the user interface and remote agent.
  - See "Installation Checklist" on page 3-6.
  - For more information, visit www.qlogic.com.
- Text in **bold** font indicates user interface elements such as a menu items, buttons, check boxes, or column headings. For example:
  - Click **Start**, point to **Programs**, point to **Accessories**, and then click **Command Prompt**.
  - Under **Notification Options**, select the **Warning Alarms** check box.
- Text in `Courier` font indicates a file name, directory path, or command line text. For example:
  - To return to the root directory from anywhere in the file structure: Type `cd /root` and press <ENTER.
  - Enter the following command: **`sh ./install.bin`**

- Key names and key strokes are indicated with UPPERCASE:
  - Press the CTRL+P keys.
  - Press the UP ARROW key.
- Text in *italics* indicates terms, emphasis, variables, or document titles. For example:
  - For a complete listing of license agreements, refer to the *QLogic Software End User License Agreement*.
  - What are *shortcut keys*?
  - To enter the date type *mm/dd/yyyy* (where *mm* is the month, *dd* is the day, and *yyyy* is the year).
- Topic titles between quotation marks identify related topics either within this manual or in the online help, which is also referred to as *the help system* throughout this document.

License Agreements

Refer to the *QLogic Software End User License Agreement* for a complete listing of all license agreements affecting this product.

Technical Support

Customers should contact their authorized maintenance provider for technical support of their QLogic products. QLogic-direct customers may contact QLogic Technical Support; others will be redirected to their authorized maintenance provider.

Visit the QLogic support Web site listed in Contact Information for the latest firmware and software updates.

QLogic Technical Support for products under warranty is available during local standard working hours excluding QLogic Observed Holidays. For customers with extended service, consult your plan for available hours.

For details about available service plans, or for information about renewing and extending your service, visit the Service Program web page at http://www.qlogic.com/services.

Training

QLogic offers training for technical professionals for all iSCSI, Converged Network, InfiniBand, and Fibre Channel products. From the main QLogic web page at www.qlogic.com, click the **Support** tab at the top, then click **Training and Certification** on the left. The QLogic Global Training Portal offers online courses, certification exams, and scheduling of in-person training.

Technical Certification courses include installation, maintenance and troubleshooting QLogic products. Upon demonstrating knowledge using live equipment, QLogic awards a certificate identifying the student as a certified professional. You can reach the training professionals at QLogic by e-mail at training@qlogic.com.

Contact Information

Support Headquarters	QLogic Corporation 4601 Dean Lakes Blvd. Shakopee, MN 55379 USA
QLogic Web Site	www.qlogic.com
Technical Support Web Site	http://support.qlogic.com
Technical Support E-mail	support@qlogic.com
Technical Training E-mail	training@qlogic.com

For Support phone numbers, see the Contact Support link at support.qlogic.com.

Knowledge Database

The QLogic knowledge database is an extensive collection of QLogic product information that you can search for specific solutions. We are constantly adding to the collection of information in our database to provide answers to your most urgent questions. Access the database from the QLogic Support Center: http://support.qlogic.com.

Legal Notices

Warranty

For warranty details, please check the QLogic Web site at http://www.qlogic.com/Support/Pages/Warranty.aspx.

Laser Safety

FDA Notice

This product complies with DHHS Rules 21CFR Chapter I, Subchapter J. This product has been designed and manufactured according to IEC60825-1 on the safety label of laser product.

CLASS I LASER

Class 1 Laser Product
Appareil laser de classe 1
Produkt der Laser Klasse 1
Luokan 1 Laserlaite

Caution—Class 1 laser radiation when open
Do not view directly with optical instruments

Attention—Radiation laser de classe 1
Ne pas regarder directement avec des instruments optiques Vorsicht—Laserstrahlung der Klasse 1 bei geöffneter Abdeckung
Direktes Ansehen mit optischen Instrumenten vermeiden

Varoitus—Luokan 1 lasersäteilyä, kun laite on auki
Älä katso suoraan laitteeseen käyttämällä optisia instrumenttej

Agency Certification

The following sections contain a summary of EMC/EMI test specifications performed on the QLogic adapters to comply with radiated emission, radiated immunity, and product safety standards.

EMI and EMC Requirements

FCC Part 15 compliance: Class A (QLE8*xxx*) 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 (QLE8*xxx*) 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:
EN55022:2006+A1:2007/CISPR22:2006: Class A (QLE8*xxx)*

EN55024:1998
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: 2009-04 Class A (QLE8*xxx*)

AS/NZS CISPR22: Class A (QLE8*xxx*)

CNS 13438: Class A (QLE8*xxx*)

This is a Class A product. In a domestic environment, this product may cause radio interference, in which case, the user may be required to take adequate measures.

警告使用者:
這是甲類的資訊產品，在居住的環境中使用時，
可能會造成射頻干擾，在這種情況下，
使用者會被要求採取某些適當的對策。

MIC: 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/telecommunications equipment)	As this equipment has undergone EMC registration for business purpose, the seller and/or the buyer is asked to beware of this point and in case a wrongful sale or purchase has been made, it is asked that a change to household use be made.

Korean Language Format— Class A

A급 기기 (업무용 정보통신기기)

이 기기는 업무용으로 전자파적합등록을 한 기기이오니 판매자 또는 사용자는 이 점을 주의하시기 바라며, 만약 잘못판매 또는 구입하였을 때에는 가정용으로 교환하시기 바랍니다.

Product Safety Compliance

UL, cUL product safety: 8200 and 3200 Series Adapters

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: 8200 and 3200 Series Adapters

TUV:
EN60950-1:2006+A11 2nd Edition
EN60825-1:1994+A1+A2
EN60825-2:2004+A1
IEC60950-1 2nd Edition (2005) CB

CB Certified to IEC 60950-1 2nd Edition

Quick Start

Installation Instructions

This Quick Start section describes how to install and configure your new QLogic converged network adapter in four simple steps:

- Step 1. Verify the Package Contents
- Step 2. Install the Adapter Hardware
- Step 3. Install the Adapter Drivers
- Step 4. Install QConvergeConsole™

CAUTION!

Keep the adapter in the antistatic bag until installation. The adapter contains parts that can be damaged by electrostatic discharge (ESD). Before handling the adapter, use standard methods to discharge static electricity. Place the adapter on the bag when examining it. Retain the bag for future use.

Step 1. Verify the Package Contents

The QLogic adapter is shipped with the following items:

- Low-profile bracket

For more information on the 8200 Series Adapter package contents, go to the following link, select **Show+** under the QLogic 8200 Series, and then select **Ordering Information**.

http://www.qlogic.com/Products/adapters/Pages/ConvergedNetworkAdapters.aspx

For more information on the 3200 Series Adapter package contents, go to the following link, select **Show+** under the QLogic 3200 Series, and then select **Ordering Information.**

http://www.qlogic.com/Products/adapters/Pages/IntelligentEthernetAdapters.aspx

Step 2. Install the Adapter Hardware

To install the adapter hardware, you need to open the computer and locate the appropriate bus slot. If necessary, consult your computer system manual for instructions on how to remove the computer cover.

Follow these steps to install the adapter hardware:

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 (or larger) bus slot.
4. Pull out the slot cover (if any) by removing the screw or releasing the lever.
5. Install the low-profile bracket (if required).
6. Grasp the adapter by the top edge and seat it firmly into the appropriate slot (see Figure i).

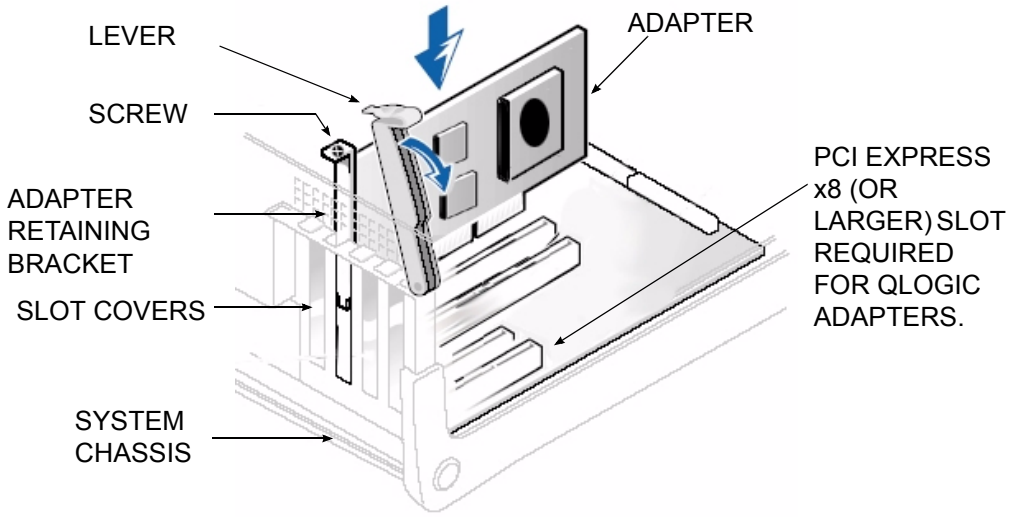

Figure i Illustration of Server Motherboard and Slots

7. Refasten the adapter's retaining bracket using the existing screw or lever.
8. Close the computer cover.

9. Plug the appropriate Ethernet cable (either copper or optical) into the adapter.
   - Optical models ship with optical transceivers already installed. The 8200 and 3200 Series Adapters will only operate with optical transceivers sold by QLogic.
   - For copper connectivity, see the list of approved SFP+ direct-attach copper cables on QLogic's Web site: http://www.qlogic.com/Resources/Documents/LineCards/Copper_Cables_Support_Matrix_Line_Card.pdf
10. Plug in the power cable and turn on the computer.

Step 3. Install the Adapter Drivers

To install the adapter drivers:

1. Go to the QLogic Driver Downloads/Documentation page at http://driverdownloads.qlogic.com.
2. Locate the appropriate driver package for the adapter product, as follows:
   - 8200 Series Adapters: Click **QLogic Products**. Select **Converged Network Adapters**, the appropriate adapter model, and your operating system. Click **Go**. The product download page opens.
   - 3200 Series Adapters: Click **QLogic Products**. Select **Intelligent Ethernet Adapters**, the appropriate adapter model, and your operating system. Click **Go**. The product download page opens.
3. On the download page under **Drivers**, select the appropriate driver and download it to your system.
4. Follow the installation instructions included in the Readme file for the downloaded driver.

Step 4. Install QConvergeConsole™

To install QConvergeConsole:

1. Go to the QLogic Driver Downloads/Documentation page at http://driverdownloads.qlogic.com.
2. Locate the QConvergeConsole package for the adapter product, as follows:
   - 8200 Series Adapters: Click **QLogic Products**. Select **Converged Network Adapters**, the appropriate adapter model, and your operating system. Click **Go**. The product download page opens.
   - 3200 Series Adapters: Click **QLogic Products**. Select **Intelligent Ethernet Adapters**, the appropriate adapter model, and your operating system. Click **Go**. The product download page opens.
3. On the download page under **Tools**, select the appropriate **QConvergeConsole** package (GUI or CLI) and download it to your system.
4. Follow the instructions in the QConvergeConsole readme file for installing the downloaded software.

Additional Resources

- To obtain the most current drivers, management tools, user instructions, and documentation, please visit the QLogic Web site at http://driverdownloads.qlogic.com and download the latest versions.
- See What's in This Guide for descriptions of user instructions provided in this document.
- For important product information, including warranty, laser safety, and agency certification, and see the Legal Notices section.
- For descriptions and procedures related to QConvergeConsole, use the built-in help system.

1 Product Overview

What is a Converged Network Adapter?

A Converged Network Adapter is a multifunction adapter that combines the capabilities of a Fibre Channel adapter and an Ethernet NIC. A Converged Network Adapter provides simultaneous Fibre Channel and Ethernet traffic over a shared 10Gb Ethernet link.

What is an Intelligent Ethernet Adapter?

The Intelligent Ethernet Adapter is a multifunction adapter that supports one Ethernet functions per port (default) and can be expanded to four Ethernet functions per port.

Functionality and Features

This section provides the following information:

- Functional Description
- Features
- Supported Operating Systems

Functional Description

The 8200 Series Adapters are Converged Network Adapters, and the 3200 Series Adapters are Intelligent Ethernet Adapters (IEA). The 8200 Series Adapters support 2×10GbE Enhanced Ethernet, FCoE, and iSCSI. The 3200 Series Adapters support 2×10GbE Enhanced Ethernet ports.

Features

The 8200 and 3200 Series Adapters provide the following features:

- NIC partitioning
- Message signaled interrupts (MSI-X)
- Device management for LAN and SAN

- Multi-boot capability including:
  - Preboot-eXecution environment (PXE)
  - iSCSI (8200 Series Adapters only)
  - Fibre Channel (8200 Series Adapters only)
  - Unified extensible firmware interface (UEFI)
- PCIe 2.0 ×8
- User diagnostics that can be run from the CLI and the GUI
- Ethernet functions include:
  - 2×10 GbE
  - Priority and virtual LAN (VLAN) tagging
  - Jumbo frames up to 9618 bytes
  - Advanced teaming
  - VLAN configuration and management
  - Preservation of teaming and VLAN configuration information during driver upgrade
- Enhanced Ethernet functions include:
  - Priority-based flow control (802.1Qbb)
  - Enhanced transmission selection (802.1Qaz)
  - Data center bridging exchange protocol (802.1Qaz)
- Advanced stateless offload features include:
  - IP, TCP, and UDP checksums
  - Large send offload (LSO)
  - Large receive offload (LRO)
- Stateful offload features (8200 Series Adapters only) include:
  - iSCSI
  - FCoE
- Several advanced management features for iSCSI and Fibre Channel adapters, including QConvergeConsole (QCC) (GUI and CLI) and NIC partitioning (NPAR)
- 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 8200 and 3200 Series adapters support commonly used operating systems: Windows, Linux, Solaris® and ESX®.

For a detailed list of the currently supported operating systems, refer to the adapter's Readme and Release Notes on the QLogic Web site:

http://driverdownloads.qlogic.com

Adapter Specifications

Physical Characteristics

The 8200 and 3200 Series Adapters are implemented as low-profile PCIe cards. The adapters ship with a full-height bracket for use in a standard PCIe slot or an optional spare low-profile bracket for use in a low-profile PCIe slot. Low-profile slots are typically found in compact servers.

Standards Specifications

The 8200 and 3200 Series adapters support the following standards specifications:

- IEEE: 802.3ae (10 Gb 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)
- iSCSI (RFC3720)

Environmental Specifications

Table 1-1. Environmental Specifications

Condition	Operating	Non-Operating
Temperature Ranges (for Altitude=900 m or 2952.75 ft)	10°C to 55°C [a] (50°F to 131°F)	–40°C to 65°C (–40°C to 149°F)
Temperature Ranges (for Altitude >900 m or 2952.75 ft)	10°C to *n*°C [b] (50°F to *n*°F [c])	–40°C to 65°C (–40°F to 149°F)
Temperature Gradient Maximum per 60 Min.	10°C	20°C
Humidity Percent Ranges—Non-condensing	20% to 80% (Max Wet bulb temperature= 29°C) [d]	5% to 95% (Max Wet bulb temperature= 38°C) [d]
Humidity Gradient Maximum per 60 Min.	10%	10%
Altitude Ranges—Low Limits	–15.2m (–50 ft)	–15.2m (–50 ft)
Altitude Ranges—High Limits	3,048m (10,000 ft)	10,668m (35,000 ft)
Airborne Contaminants—ISA-71 Level	G1 [d]	G1 [d]

[a] 200LFM is required to operate at this temperature.

[b] 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.

$$35 - \frac{\text{Maximum Altitude(in meters)} - 900}{300}\ ^\circ\text{C} \quad or \quad 35 - \frac{\text{Maximum Altitude(in ft)} - 2952.75}{984.25}\ ^\circ\text{C}$$

[c] 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.

$$95 - \left(\frac{(\text{Maximum Altitude(in meters)} - 900) \times 1.8}{300}\right)\ ^\circ\text{F} \quad or \quad 95 - \left(\frac{(\text{Maximum Altitude(in ft)} - 2952.75) \times 1.8}{984.25}\right)\ ^\circ\text{F}$$

[d] Maximum corrosive contaminant levels measured at =50% relative humidity; see Table 3 in ISA-71.04-1985.

Converged Network Adapter SuperInstaller Installation

QLogic Windows SuperInstaller

Refer to the *QLogic Windows SuperInstaller* Readme file for information.

QLogic Linux SuperInstaller

Refer to the *QLogic Linux SuperInstaller* Readme file for information.

Multi-boot Image for 10Gb CNA—CNA Function Configuration Utility

This section contains the following information:

For an alternate method of iSCSI boot from SAN (3200 Series Adapters only), use the fields in the iSCSI Boot Firmware Table (iBFT). 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 software-initiated iSCSI protocol. To view the iBFT specification, see the following Web page:

http://www.microsoft.com/whdc/system/platform/firmware/ibft.mspx

Operating System Support

This multi-boot code supports DOS, Windows Server® 2008; Solaris x86; and Linux on IA32, AMD64, and Intel® x64-based systems. Refer to the Readme file for operating system (OS) information.

Multi-boot Package Contents

The multi-boot package for 10Gb adapters is a compressed file that contains the 82*xx*/32*xx* BIOS, UEFI, PXE, and firmware. This package also includes the QLflash application.

Converged Network Adapter Function Configuration Package Contents

The following files are included for updating the adapter multi-boot code:

- `update.bat`—DOS batch file that calls the executable files to update the adapter multi-boot.
- `QLflash.exe`—Utility to update multi-boot code and firmware.
- `DOS4GW.exe`—This file is required to use the `QLflash.exe`.
- `p3pyyyyy.bin`—Combined binary file, which includes the binaries for the BIOS, UEFI, PXE, and firmware.

Using QLflash

`QLflash` is a native DOS utility. Refer to the Readme file for more information. To run this utility, boot to a DOS hard drive or USB removable drive.

Utility Version:

/`VER` = Display version of the QLflash utility

Help Options:

`/?` = Help menu

Updating the Multi-boot Code

To write the multi-boot code to Flash memory:

1. Insert the 10Gb adapter in the system.
2. Boot to DOS.
3. Run the update script at the command prompt:

   `C:\>update.bat`

   This script program updates the multi-boot image on the adapter.
4. Reboot the system.

NOTE:

You can also use QCC GUI/CLI to flash the multi-boot image. After the multi-boot code is updated, power cycle the server for the new changes to take effect.

QLflash Command Line Options

The executable file `QLflash.exe` is used by the `UPDATE.BAT` file to update your adapter multiboot-code. The application `QLflash.exe` may be used to read, write, or verify either the multi-boot image or the NVRAM on the adapter.

The following paragraphs describe the command line options available with this utility. Use of QLflash will modify the way your adapter operates, and it must be used with extreme caution.

Certain features (that is, the NVRAM options) may require additional data.

Files and passwords are not provided in this file.

QLflash Options

`QLflash.exe <1...N | ALL> [options]`

`/SLT`

`SLT`: Sets the application to silent mode

Use application return code for success or failure

`/SIL=filename.ext`

`SIL`: Load multiflash image from file

`/PRV | /VPP | /PRN`

`PRV`: Print firmware versions

`VPP`: Print VPD contents

`PRN`: Print MAC addresses and WWPN

`/CFU=filename.ext | /CFS=filename.ext | /CFC=filename.ext`

`CFU`: Flash board configuration from file

`CFS`: Save board configuration file (need template file `BRDCFG.DAT`)

`CFC`: Compare board configuration to a binary file

`/NVU=filename.ext | /NVS=filename.ext | /NVC=filename.ext`

`NVU`: Flash NVRAM from file

`NVS`: Save NVRAM to file (need template file `NVRAM.DAT`)

`NVC`: Compare Flash to NVRAM file

`/UIL=filename.ext | /UIS=filename.ext | /UIC=filename.ext`

`UIL`: Flash user info data from file

`UIS`: Save user info data to file

`UIC`: Compare user info data to a binary file

NOTE:

- If you used an `FC RAID` target in a cluster environment, you should enable the `Enable Target Reset = Enabled (Advanced Adapter Settings)`.
- Use the `/I` option if the update utility, QLflash, does not detect your adapter.
- QLogic recommends disabling the internal disk before installing the OS or booting to the FCoE disk.

Adapter Configuration (CNA/IEA Function Configuration Utility)

NOTE:

The CNA Function Configuration utility does not run on SUN™ SPARC® systems.

This section provides detailed configuration information for advanced users who want to customize the configuration of the 10Gb adapter and the connected devices. You can configure the adapter using the CNA Function Configuration utility.

To access the CNA Function Configuration utility, press ALT+ Q during the adapter initialization (it may take a few seconds for the menu to appear). If you have more than one adapter, the utility will ask you to select the adapter you want to configure. After changing the settings, the utility reboots your system to load the new parameters.

CAUTION!

If the configuration settings are incorrect, your adapter may not function properly.

Upon entering the CNA Function Configuration utility, the following selections are available from the Setup Menu/Function Configuration menu:

- Function
- Type
- MinBW%
- Protocol
- Restore to Non-NIC Partition Settings
- Exit the CNA Function Configuration Utility

Function

The Function parameter applies only to only function 0 and function 1. Press the ENTER key to configure either PXE or iBFT settings. (The Function option is predicated on the Protocol that is selected.)

Adapter Settings

Press the ENTER key on the Function number to access Adapter Settings, and to configure PXE Boot or iBFT.

PXE. The following are available when Protocol is set to PXE:

- `Setup Menu Wait Time` (0–15: the default is 5 seconds)

  Specifies the time in seconds the menu will wait.

- `Enable PXE Boot` (Enabled/Disabled: Default—`Disabled`)

  Enter to toggle between `Enabled` and `Disabled`. Use the `Enable PXE Boot` option to attempt a PXE boot on the selected function.

NOTE:

For an embedded LAN on motherboard (LOM), the **Enable PXE Boot** parameter is controlled by the system BIOS and is non-configurable in the CNA Function Configuration utility.

iBFT. The following parameters are available when **Protocol** is set to **iBFT**:

- `Enable ROM` (Enabled/Disabled: Default—`Enabled`)

  Enter to toggle between `Enabled` and `Disabled`. Use the **Enable ROM** option to attempt iBFT boot on the selected function.

NOTE:

For an embedded LOM, the **Enable PXE Boot** parameter is controlled by the system BIOS and is non-configurable in the CNA Function Configuration utility.

- `IP Format (`IPv4/IPv6: Default —`IPv4`)

  Enter to specify IPv4 or IPv6 address format.

- `VLAN Mode` (Enabled/Disabled: Default—`Disabled`)

  This setting enables VLAN.

- `VLAN ID` (2 – 4095: Default—`2`)

  Enter to specify the VLAN ID.

Initiator Settings

- `Enable DHCP` (Enabled/Disabled: Default—`Enabled`)

  When `Enable DHCP` is set to `Enabled`, the adapter utilizes DHCP to obtain its IP address, subnet mask, and gateway IP address.

- `Target Param From DHCP` (Enabled/Disabled: Default—`Enabled`)

  When `Target Param From DHCP` is set to Enabled, the adapter uses DHCP to obtain the target IP address, subnet mask, and gateway IP address.

- `Initiator Name` (Programmed during manufacturing)

  Enter to modify the iSCSI name of the initiator.

- `Initiator IP Address`

  When `Enable DHCP` is set to `Disabled`, this field must contain a valid IP address.

- `Initiator Subnet Mask`

  When `Enable DHCP` is set to `Disabled`, this field must contain a valid subnet mask.

- `Gateway`

  When `Enable DHCP` is set to `Disabled`, this field must contain a valid gateway address.

- `Reverse CHAP Name`

  Enter to configure bidirectional CHAP name.

- `Reverse CHAP Secret`

  Enter to configure Bidirectional CHAP secret.

Primary/Secondary Target Settings

When `Target Param From DHCP` is set to `Disabled`, enter to manually configure the target parameters.

- `Target Name`

  Enter to specify the iSCSI name of the intended boot target. (Depending on your configuration, it may not be necessary to configure this option).

- `Target IP Address`

  This field must contain a valid IP address of the intended iSCSI boot target.

- `Target Port`

  Enter to specify the port used by the target. (By default, iSCSI uses port 3260—do not modify this unless it is necessary.)

- `Boot LUN`

  Enter to specify the LUN of the boot device. (Depending on your configuration, it may not be necessary to change this option.)

- `CHAP Type`

  - `None`—CHAP is not used.
  - `One-Way`—Unidirectional CHAP is used.
  - `Mutual`—Bidirectional CHAP is used.

- `Target CHAP Name`

  Enter to specify the target CHAP name.

- `Target CHAP Secret`

  Enter to specify the target CHAP secret.

Type

Enter to set a function to a particular type.

- `NIC`—NIC protocols will be supported.
- `None`—The function will be disabled.
- `iSCSI`—The function will support iSCSI protocol.
- `FCoE`—The function will support FCoE protocol.

NOTE:

Functions 0 through 3 support only NIC or None. Functions 4 and 5 support NIC, iSCSI, or None. Functions 6 and 7 support NIC, FCoE, or None.

MinBW%

Enter to specify the percentage of bandwidth allocated to the function. The combined MinBW% values for a port's functions cannot exceed 100 percent.

Protocol

The Protocol parameter applies only to only function 0 and function 1.

- PXE—PXE will be supported on the NIC.
- iBFT—iBFT will be supported on the NIC.

NOTE:

For an embedded LOM, the Protocol parameter is controlled by the system BIOS and is non-configurable in the CNA Function Configuration utility.

Restore to Non-NIC Partition Settings

Press the ENTER key to restore the NIC partition settings.

Exit the CNA Function Configuration Utility

Press the ENTER key to select from the following:

- Reboot System
- Return to Fast!UTIL

Installing the QLogic Adapter vCenter Plug-in for ESX/ESXi

This section on installing the QLogic Adapter vCenter Plug-in contains the following:

Introduction

The QLogic Adapter vCenter Plug-in is a user-interface extension to the vSphere™ Client that allows you to manage and configure QLogic adapters, including the NIC partitioning feature, on ESX and ESXi servers. When the you select an ESX or ESXi host that has QLogic adapters installed, the extension appears in the client as an additional tab named QLogic Adapter, as shown in Figure 1-1.

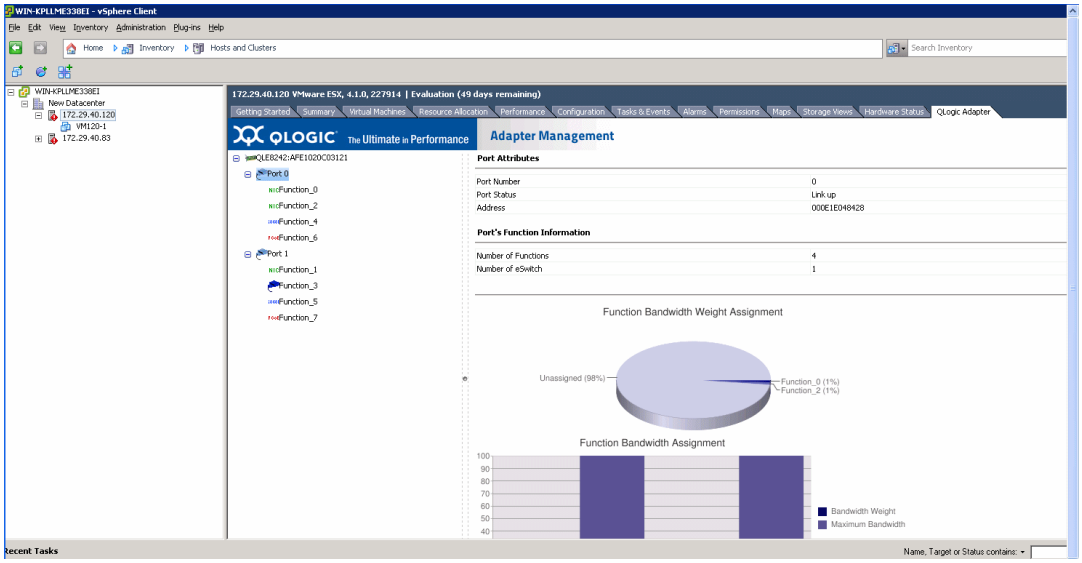

Figure 1-1. vSphere Client Showing QLogic Adapter Tab

Requirements

The QLogic Adapter vCenter Plug-in requires the components in Figure 1-2 to be installed and running.

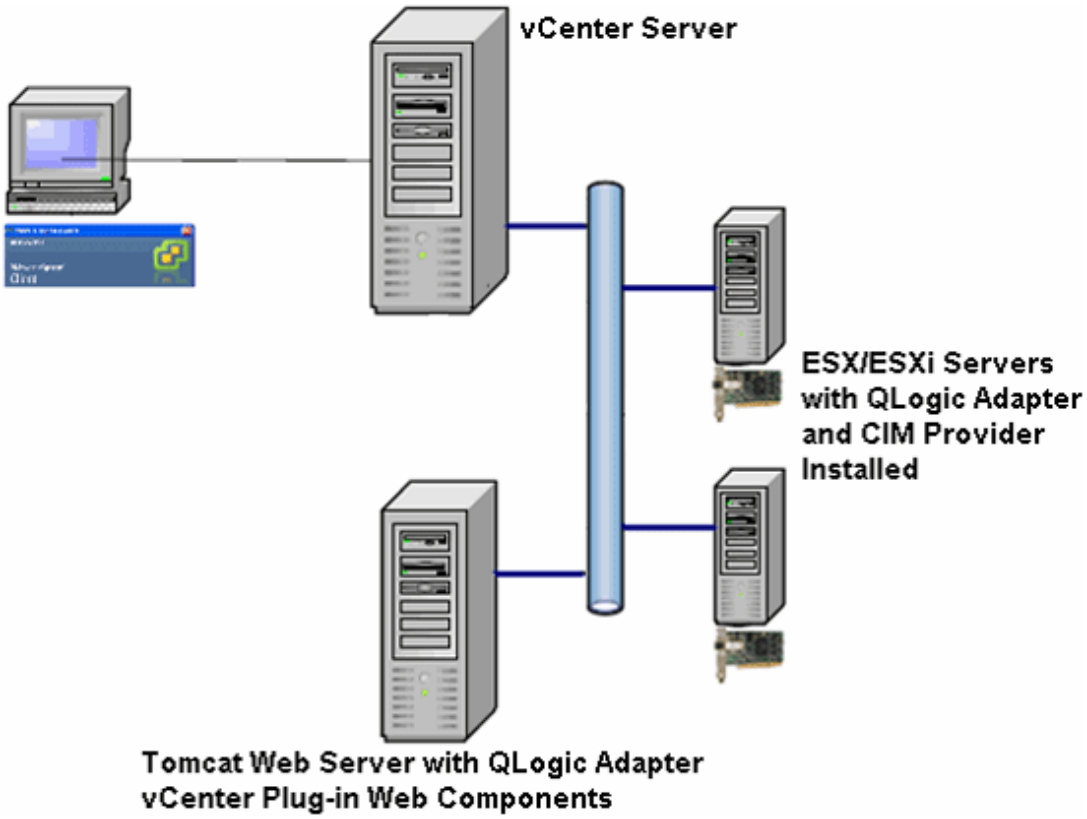

Figure 1-2. vCenter Plug-in Requirements

QLogic provides the following components that must be installed on the ESX or ESXi Server, vCenter Server, and Tomcat Web Server.

ESX/ESXi Server

- QLogic adapters with firmware and driver.
- QLogic Adapter CIM Provider

NOTE:

For the latest versions of QLogic firmware, drivers, and QLogic CIM Provider, please refer to the Release Notes on the QLogic Web site:

http:\\driverdownloads.qlogic.com

---

vCenter Server

- QLogic XML configuration file to register the plug-in to the vCenter Server

Tomcat Web Server

- QLogic Web-based extension to the vSphere Client

Installing the QLogic Adapter CIM Provider

The QLogic Adapter CIM Provider for VMware ESX was generated as a vSphere Installation Bundle (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.

This section provides the following installation procedures for the CIM Provider:

- Initial Installation
- Subsequent Update Installation

Initial Installation

To install the QLogic Adapter CIM Provider using the esxupdate command (ESX systems only):

1. Copy the offline-bundle.zip file into the root directory (/) of the ESX system.
2. Issue the esxupdate command as follows:

```
# cd /
# esxupdate --bundle offline-bundle.zip --nodeps –nosigcheck
--maintenancemode update
```

To install the QLogic Adapter CIM Provider from a remote host using the vSphere CLI vihostupdate command (ESXi systems only):

NOTE:

To update the provider from a remote host using the vihostupdate command, make sure that the ESXi system is in Maintenance Mode. To put the ESXi host in Maintenance Mode using vSphere Client, select **Inventory,** select **Host**, and then select **Enter Maintenance Mode**.

---

1. Copy the offline-bundle.zip file to any location on the host where either the vSphere CLI package is installed or vSphere Management Assistant (vMA) is hosted.
2. Navigate to the location of the offline-bundle.zip file.

3. Issue the `vihostupdate` command to install the offline bundle as follows:

   ```
   # vihostupdate.pl <conn_options> --install --bundle
   offline-bundle.zip --nosigcheck
   ```

   For available options, refer to the `vihostupdate` page.

4. Follow the on-screen instructions to complete the installation. You may need to reboot the ESXi system.

NOTE:

For more details on `vihostupdate`, refer to the documents on the VMware vSphere Command-Line Interface Documentation page, located here:

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

Subsequent Update Installation

To update the QLogic Adapter CIM Provider after a prior VIB installation:

1. Follow the instructions in the Removing the QLogic Adapter CIM Provider section to remove the existing VIB.
2. Follow the instructions in the Initial Installation section 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 should start automatically and load the QLogic Adapter CIM Provider when necessary.

For ESX systems, you can also manually stop, start, or restart the SFCB CIMOM by issuing the following commands.

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

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

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

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

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

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

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

Removing the QLogic Adapter CIM Provider

To remove the QLogic Adapter CIM Provider, use either the `esxupdate` or `vihostupdate` command.

To uninstall the QLogic Adapter CIM Provider using the `esxupdate` command:

1. Query and find the existing VIB you are updating as follows:

   `# esxupdate query --vib-view | grep qlogic`

   For example, the VIB ID may have a format similar to the following:

   ESX/ESXi 4.0: cross_qlogic-nic-provider_400.*x.x.x*-000000
   ESX/ESXi 4.1: cross_qlogic-nic-provider_410.*x.x.x*-000000

   where *x.x.x* is the version number of the existing provider.

2. Remove the existing VIB as follows:

   `# esxupdate remove -b <vibID>`

To uninstall from a remote host using the `vihostupdate` command:

NOTE:

To uninstall the provider from a remote host using the `vihostupdate` command, make sure that the ESX/ESXi system is in Maintenance Mode. To put the ESX/ESXi host in Maintenance Mode using 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 to find the Bulletin ID of the existing provider:

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

   For example, the Bulletin ID may have a format similar to the following:

   ESX/ESXi 4.0: QLGC_NIC_PROVIDER-ESX-4.0.0-qlogic-nic-provider-*x.x.x*
   ESX/ESXi 4.1: QLGC_NIC_PROVIDER-ESX-4.1.0-qlogic-nic-provider-*x.x.x*

   where *x.x.x* is the version number of the existing provider.

2. Remove the existing VIB as follows:

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

NOTE:

For more details on `vihostupdate`, refer to the documents on the VMware vSphere Command-Line Interface Documentation page, located here:

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

Starting the QLogic Adapter vCenter Plug-in

To start the QLogic Adapter vCenter Plug-in:

1. Start the VMware vSphere Client and connect to the vCenter Server by entering the IP address or name, user name, and password.
2. Click **Login.**
3. If the Security Warning dialog box appears, click **Ignore** to use the current SSL certificate.
   - If you start and connect the vSphere Client directly to an ESX/ESXi server, the vCenter plug-in does not open.
   - If you have not already done so, create a data center and add the ESX server.
4. In the left pane, select the IP address of the VMware server.
5. In the right pane, click the **QLogic Adapter** tab to view the Web page.

   The vCenter plug-in retrieves the adapter information from the server.

NOTE:

If the server does not have the QLogic Adapter CIM Provider and adapters installed, or if the vCenter Plug-in installation and registration was not successful, the **QLogic Adapter** tab is not shown.

Removing the QLogic Adapter vCenter Plug-in

To remove the QLogic Adapter vCenter Plug-in:

1. In the Windows **Control Panel**, select **Add or Remove Programs**.
2. In the **Add or Remove Programs** dialog box, select the **QLogic Adapter vCenter Plug-in**, and then click **Change/Remove**.
3. Follow the instructions in the QLogic Adapter vCenter Plug-in installer to remove the plug-in.

Installing the Flash Utility

For ESX, before you perform a Flash update on QLogic Adapters using the QLogic Adapter vCenter Plug-in, ensure that the QLflash utility is also installed on the ESX system.

To install the QLflash utility on ESX:

1. Unzip the `esx_qlflash.zip` file, which contains the `qlflash userworld` executable.
2. Copy the QLflash binary file to the `/usr/lib/vmware/bin/` directory.
3. Add executable permission for QLflash by issuing the following command:

   ```
   # chmod +x qlflash
   ```

4. If it does not already exist, edit and append the **`/usr/lib/vmware/bin/qlflash`** entry in the `/etc/vmware/UserWorldBinaries.txt` file.

For ESXi, the QLflash utility is not required, so no additional steps are necessary.

Updating the Flash

NOTE:

For ESX systems, before you perform a Flash update on QLogic Adapters using the QLogic Adapter vCenter Plug-in, ensure that the QLflash utility is also installed. Follow the instructions in the Installing the Flash Utility section.

To update the Flash from the QLogic Adapter vCenter Plug-in:

1. Follow the instructions in the Starting the QLogic Adapter vCenter Plug-in section to start the vCenter plug-in.
2. In the left pane of the **QLogic Adapter** page, select the adapter, and then click the **Update Adapter Flash Image** link.
3. In the **Select Flash File for Update** dialog box, click **Browse**.
4. In the **Choose File to Upload** dialog box**,** select the `.bin` Flash file from the extracted Flash kit package that is compatible with your adapter, and then click **Open**.
5. In the **Select Flash File for Update** dialog box, click **Send**.
6. Verify the current Flash version and file version, and then click **OK** to continue the update.

7. When asked “Do you want to reset the adapter to activate the firmware immediately after successful update?,” click **OK** if you want the new firmware to take effect immediately. Or, click **Cancel** to have the new firmware take effect after the next system reboot.
8. Wait for the Flash update process to complete. Processing time depends on the network connection and the system configuration. Do not interrupt the update process.

CAUTION!

The Flash update may take up to ten minutes to complete. Do not cancel the task or reboot the server during this time. Doing so may corrupt the firmware on the adapter.

9. In the Flash update successful completion message box, click **OK**.
10. If you clicked **OK** in Step 7, you can click **Refresh** to verify the new firmware version. Otherwise, you must reboot the system for the new firmware to take effect.

Using the vCenter Plug-in on a Tomcat Server

VMware requires that all vCenter plug-ins are web-based applications hosted on a Tomcat Server, which can be downloaded and installed on Linux and Windows systems. Here is a link to Tomcat:

http://tomcat.apache.org/

The QLogic vCenter Plug-in supports Tomcat 5, 6, and 7. The installer dynamically detects the Tomcat version and installs accordingly.

This section provides the following procedures for using the vCenter Plug-in on a Tomcat Server:

- Installing Tomcat on Linux
- Starting and Stopping Tomcat on Linux
- Installing Tomcat on Windows
- Starting and Stopping Tomcat on Windows

Installing Tomcat on Linux

To install Tomcat on a Linux OS:

1. Go to the following URL:

   http://tomcat.apache.org/download-55.cgi

2. Locate the following directory and file:

   Binary Distribution > Core > tar.gz (apache-tomcat-5.5.28.tar.gz)

3. Unzip the `tar.gz` file by issuing the following command:

   root # tar zxf file.tar.gz

4. Create a symbolic link to a Tomcat directory by issuing the following command:

   # ln -s apache-tomcat-5.5.28 tomcat

   Where setting variables include the following:

   export JAVA_HOME=/root/gwt/jdk1.6.0_17/
   export CATALINA_HOME=/root/gwt/tomcat/

Starting and Stopping Tomcat on Linux

To start Tomcat, issue the following command:

 # $CATALINA_HOME/bin/startup.sh

To stop Tomcat, issue the following command:

 # $CATALINA_HOME/bin/shutdown.sh

Installing Tomcat on Windows

To install Tomcat on a Windows OS:

1. Go to either of the following URLs:

   http://tomcat.apache.org/download-60.cgi

   http://tomcat.apache.org/download-70.cgi

2. Locate the following directory and file:

   Binary Distribution > Core > 32-bit/64-bit Windows Service Installer

3. To install the Tomcat service, get the 32-bit/64-bit Windows Service Installer, save it, and run it.

Starting and Stopping Tomcat on Windows

To start and stop Tomcat on a Windows OS:

1. Go to **Computer Management** > **Services and Application** > **Services** > **Apache Tomcat <*version number*>**.
2. Right-click, and then select **Start /Stop** to initiate Tomcat service, or select **Stop** to halt the Tomcat service.

Installing the vCenter Plug-in on Tomcat

1. Download the QLogic Adapter vCenter Plug-in installer EXE; for example, `QLogic Adapter VI Plugin 1.0.4.exe`.
2. Run the installer by double-clicking the `EXE` file or by typing the name of the `EXE` on a command prompt.

   The *InstallAnywhere* installer prepares to install the QLogic Adapter vCenter Plug-in.
3. On the Introduction window, click **Next**.
4. On the **Please Wait** window, wait while the QLogic Adapter vCenter Plug-in is configured for your system.
5. On the **Choose Install Folder** window, either accept the default installation directory, or click **Choose** to specify a different folder.
6. Click **Install** to install files to the installation directory specified in the previous step.

   A progress window shows the status of the installation.
7. On the **User Input** window, type your vCenter Server IP address, username, and password, as well as the IP address of your local Tomcat Server. Then click **Next** to continue.
8. On the **Please Wait** window, wait while the QLogic Adapter vCenter Plug-in is configured for your system and registers the plug-in with the vCenter Server.
9. On the **Registration Result** window, click **Finish** to complete the plug-in installation.

Plug-in Unregistration from a Manual Installation

If you have performed a manual installation of the vCenter plug-in, you must perform a manual uninstall before running the vCenter installation wizard.

VMware provides two type of scripts for vCenter plug-in registration:

For Windows PowerShell® scripting:

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

For Perl:

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

1. Before you can use the script, download the appropriate VI SDK from VMware.

   For PowerShell, download **vSphere PowerCLI**: http://communities.vmware.com/community/vmtn/vsphere/automationtools/powercli

   For Perl VI SDK, download **vSphere SDK for Perl**: http://www.vmware.com/support/developer/viperltoolkit/

2. After you download and install the SDK and the registration script, follow VMware instructions to register the vCenter 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"
   ```

3. Replace the username and password with the correct information to log into the vCenter Server.

Launching the Plug-in from vSphere Client

To launch the plug-in from vSphere client:

1. Start the vSphere Client and connect to the vCenter Server. (If you start and connect the vSphere Client directly to an ESX or ESXi server, the vCenter plug-in does not appear.)
2. If you have not already done so, create a data center and add the ESX server.
3. In the left pane, select the server.

   A row of tabs appear in the right pane. If the server has QLogic Adapter CIM Provider and adapters installed, and if the vCenter plug-in installation and registration were successful, the **QLogic Adapter** tab appears in the tab row.

4. Click the **QLogic Adapter** tab to view a web page.

2 Configuring NIC

This section describes the driver and QConvergeConsole (QCC) agent installation, configuration, operation, and troubleshooting of the NIC function of the 8200 and 3200 Series Adapters. Refer to the *QCC User's Guide* for information on QCC Agents.

Installing NIC in Linux

Refer to the *Intelligent Ethernet Adapter and Converged Network Adapter Networking Inbox Driver Update for Linux* Readme file for information.

Packaging Content

Refer to the *Intelligent Ethernet Adapter and Converged Network Adapter Networking Inbox Driver Update for Linux* Readme file for information.

Linux OS Support

Refer to the *Intelligent Ethernet Adapter and Converged Network Adapter Networking Inbox Driver Update for Linux* Readme file for information.

Supported Features

Refer to the *Intelligent Ethernet Adapter and Converged Network Adapter Networking Inbox Driver Update for Linux* Readme file for information.

Using the Driver

Refer to the *Intelligent Ethernet Adapter and Converged Network Adapter Networking Inbox Driver Update for Linux* Readme file for information on installing the driver, removing the driver, and driver system parameters.

Installing NIC in ESX

Refer to the *Intelligent Ethernet Adapter and Converged Network Adapter Networking Inbox Driver Update for ESX* Readme file for information.

Packaging Content

Refer to the *Intelligent Ethernet Adapter and Converged Network Adapter Networking Inbox Driver Update for ESX* Readme file for information.

ESX OS Support

Refer to the *Intelligent Ethernet Adapter and Converged Network Adapter Networking Inbox Driver Update for ESX* Readme file for information.

Using the Driver

Refer to the *Intelligent Ethernet Adapter and Converged Network Adapter Networking Inbox Driver Update for ESX* Readme file for information on installing and removing the driver.

Installing NIC in Windows

Refer to the *Intelligent Ethernet Adapter and Converged Network Adapter NDIS Miniport Driver for Windows* Readme file for information.

Packaging Content

Refer to the *Intelligent Ethernet Adapter and Converged Network Adapter NDIS Miniport Driver for Windows* Readme file for information.

Windows OS Support

Refer to the *Intelligent Ethernet Adapter and Converged Network Adapter NDIS Miniport Driver for Windows* Readme file for information.

Using the Driver

Refer to the *Intelligent Ethernet Adapter and Converged Network Adapter NDIS Miniport Driver for Windows* Readme file for information on installing and removing the driver.

Configuring PXE Boot

This section provides procedures for configuring the 8200 and 3200 Series Adapters to perform PXE boot. The example uses function 1 and NIC 1.

To configure PXE boot:

1. Enter the system BIOS by pressing the F2 key.
2. On the BIOS window (Figure 2-1), select **Integrated Devices**, and then press the ENTER key.

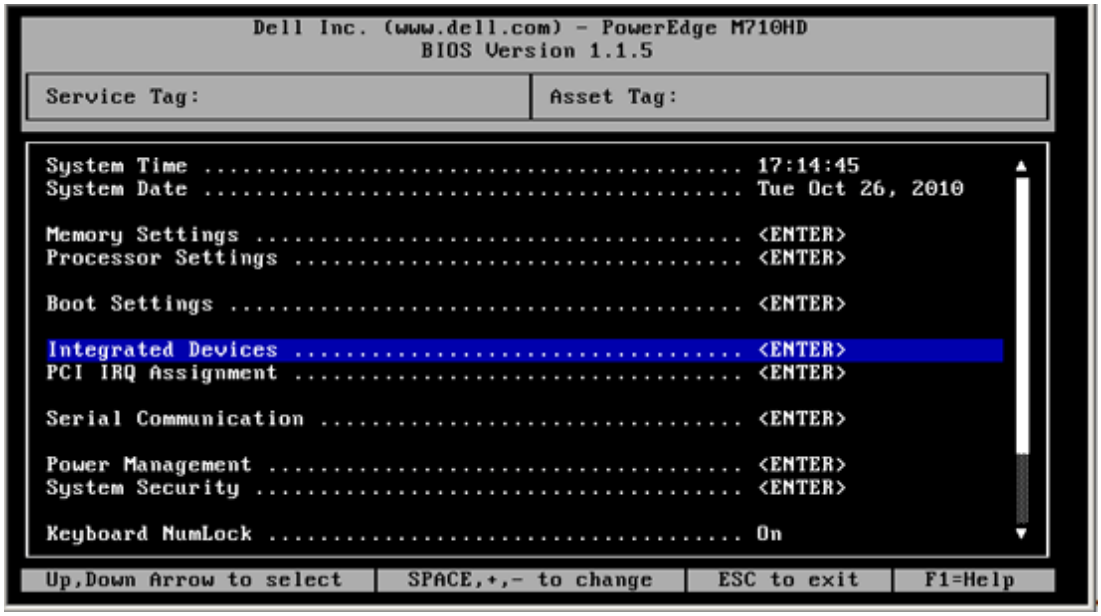

Figure 2-1. Dell BIOS: Integrated Devices

3. Set the **Embedded NIC1 and NIC2** option to **Enabled**.
4. Set the **Embedded NIC1** option to **Enabled with PXE**.
5. Press the ESCAPE key twice, and then select **Save changes and exit**. The system reboots.
6. (Optional) During POST, press the CTRL+Q keys to enter the **QLogic 8200 Series / 3200 Series CNA Function Configuration** window.

7. On the **CNA Function Configuration** main window, ensure that **Protocol** is set to **PXE** (Figure 2-2).

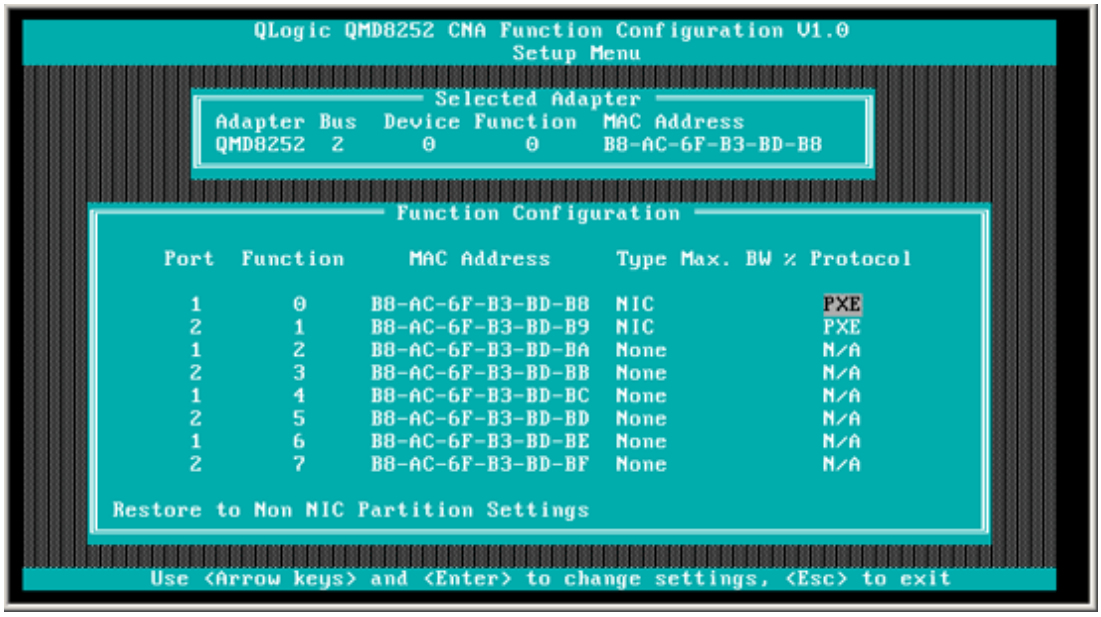

Figure 2-2. QLogic 8200 CNA Function Configuration

8. Press the ESC key to exit.
9. Select **Save changes** to save your edits, exit, and reboot the system.
10. During POST, press the F2 key to enter the BIOS system.

11. Select **Boot Settings**, and then press the ENTER key (Figure 2-3).

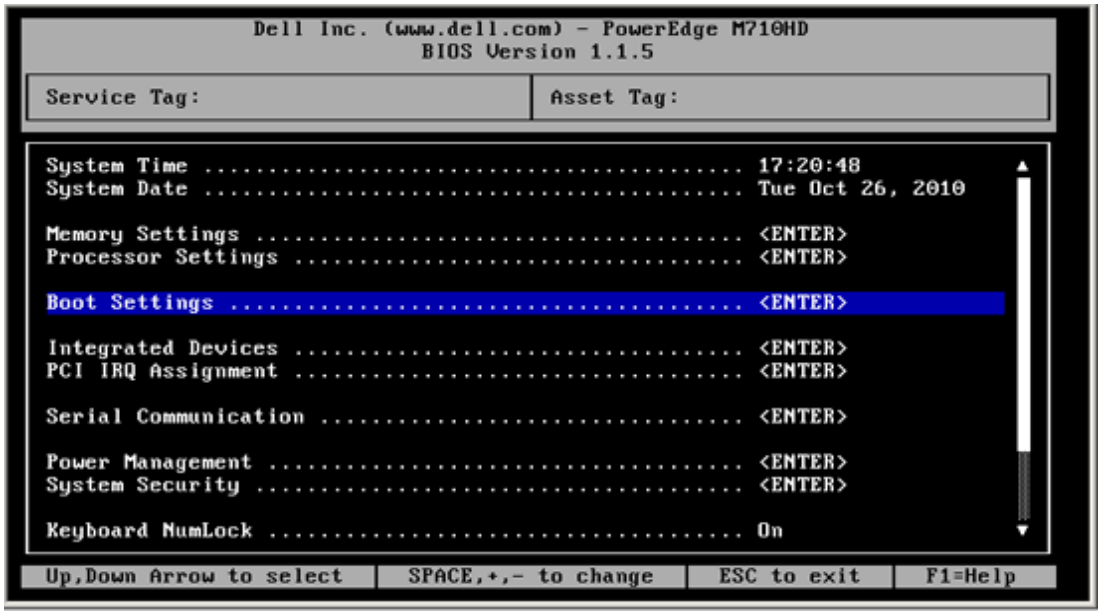

Figure 2-3. BIOS: Boot Settings

12. Select the **Boot Sequence** option, and then press the ENTER key.

13. Select the **Embedded NIC 1 QLogic PXE...** entry, and then press the UP ARROW key to move this entry to the first position (Figure 2-4).

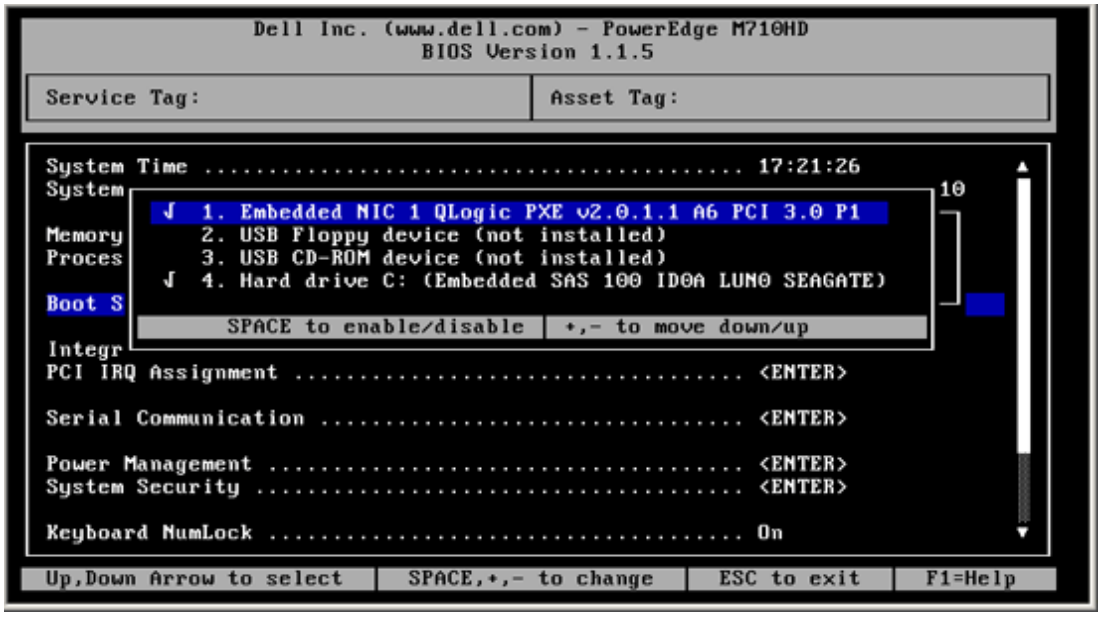

Figure 2-4. Embedded NIC 1 QLogic PXE

14. Press the ESCAPE key, and then select **Save changes and exit**.

    The system reboots.

15. After the system reboot, follow the screen prompt for PXE boot server for the installation of OS of your choice.

Configuring Driver Software Parameters

Linux NIC Driver Management Applications

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:

- qaucli Utility
- ethtool Utility
- phantomcore Utility

qaucli Utility

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

- QConvergeConsoleCLI-<version>_linux_<arch>.install.tar.gz—Package file
- QConvergeConsoleCLI-<version>_<arch>.rpm—RPM installer package file

Example RPM package installation:

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

```
rpm -qa | grep QConvergeConsoleCLI
```

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

   ```
   rpm -qa QConvergeConsoleCLI
   ```

2. If an older version is found, erase that version by issuing the following command:

   ```
   rpm -e QConvergeConsoleCLI
   ```

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

   ```
   rpm -ihv QConvergeConsoleCLI-1.0.00-04.i386.rpm
   ```

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

ethtool Utility

Use the ethtool utility to view adapter statistics and configure interface options. For additional details, refer to *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
```

phantomcore Utility

Use the phantomcore utility to collect a firmware dump that is useful for forwarding to QLogic Technical Support.

To collect a firmware dump:

1. To turn off auto recovery in the driver, issue the following command:

   ```
   insmod qlcnic.ko auto_fw_reset=0
   ```

2. To collect the firmware dump, issue the following command:

   ```
   phantomcore -i <interface name>
   ```

The utility generates a core file that contains the firmware dump.

Windows Property Pages

To access the property pages in Windows:

1. Access the Device Manager as follows:
   a. On the Windows desktop, click **Start**, **Control Panel**, click **Administrative Tools**, and then click **Computer Management**.
   b. In the **Computer Management** dialog box, click **Device Manager**.
2. In the left pane of the **Device Manager** dialog box, right-click **QLogic 10Gb Ethernet Adapter**, and then click **Properties**.
3. On the adapter properties dialog box, click the **Advanced** tab to bring that page to the front.
4. On the **Advanced** page, configure the parameters specified in the Table 2-1.
5. (Optional) Click other tabs to bring those pages to the front and view or change settings, for example:
   - Click the **Information** tab to view the configuration.
   - Click the **Details** tab to view NIC driver details. On the Details page, click items under **Property** to view the selected item's value.
   - Click the **Driver** tab to update, rollback, disable, or uninstall the NIC driver.
   - Click the **Resource** tab to view resource settings.
   - Click the **Statistics** tab to view the transmit and receive (general and QLogic-customized) statistics gathered while the NIC is operational.
6. When you are through viewing and changing driver properties, click **OK** to close the dialog box, or click **Cancel** to revert to the previous driver configuration.

Table 2-1. Windows Driver Configurable Parameters

Property	Description
Completion Queue Size	Specifies the size for the ring where command completion and incoming receive indication status will be posted. Registry Key: `CompletionQueueSize` Default: `16384` Values: `1024, 2048, 4096, 8192, 16384, 32768`
Flow Control	Sets the hardware flow control parameters. Registry Key: `*FlowControl` Default: `Rx and Tx Enabled`. Values: `Disabled, Rx Enabled, Tx Enabled, Rx and Tx Enabled`
Health Monitoring	Checks health of firmware, monitors normal operation, and ensures recovery of the firmware.
Interrupt Moderation	Allows interrupt coalescing during receive and transmit operation. Registry Key: `*InterruptModeration` Default: `Disabled` Values: `Enabled, Disabled`
IPv4 Checksum Offload	Enables and disables IPv4 checksum offload. Registry Key: `*IPChecksumOffloadIPv4` Default: `Rx and Tx Enabled`. Values: `Disabled, Rx and Tx Enabled, Rx Enabled, Tx Enabled`
Large Receive Offload	Enables and disables TCP large receive offload. Enables collapsing of multiple MTU size TCP packets into bigger segments before handing these over to the host. Registry Key: `LRO` Default: `Enabled` Values: `Enabled, Disabled`

Table 2-1. Windows Driver Configurable Parameters (Continued)

Property	Description
Large Send Offload V1 IP4	Enables and disables TCP large send offload. Allows host TCP stack to give bigger than maximum segment size (MSS) packets to the driver, and then with the help of hardware, splits the larger segments into MTU size packets before sending on wire. Registry Key: `*LsoV1IPv4` Default: `Enabled`. Values: `Enabled, Disabled` Size of LSO v1: `64K`
Large Send Offload V2 IP4	Enables and disables TCP large send offload. Allows host TCP stack to give bigger than MSS packets to the driver, and then with the help of hardware, splits the larger segments into MTU size packets before sending on wire. Registry Key: `*LsoV2IPv4` Default: `Enabled`. Values: `Enabled, Disabled` Size of LSO v2: `128K`
Large Send Offload V2 IP6	Enables and disables TCP large send offload. Allows host TCP stack to give bigger than MSS packets to the driver, and then with the help of hardware, splits the larger segments into MTU size packets before sending on wire. Registry Key: `*LsoV2IPv6` Default: `Enabled`. Values: `Enabled, Disabled` Size of LSO v2: `128K`
Locally Administered Address	Defines the locally administered address (LAA) that users and administrators can set on this interface. This address overrides the permanent address of the adapter (that may have been Flashed into the hardware). Registry Key: `NetworkAddress` Format: Hexadecimal Options: Value (a value is required, enter the 12 hex bytes of the MAC address to be used), or not present.

Table 2-1. Windows Driver Configurable Parameters (Continued)

Property	Description
Max Ethernet Frame Size	Specifies the Ethernet frame size for packet transmission and receive (includes the MAC header). Reg Key: `MaxFrameSize` Default: `1514` (corresponds to 1514 bytes on the wire + 4 bytes of CRC) Max: `9614` Min: `142` (Windows 2003), `590` (Windows 2008)
Max Jumbo Buffers	Specifies the number of jumbo-sized receive buffers allocated for the jumbo ring (not frame). Reg Key: `RxJumboRingSize` Default: `4096` Values: `1024, 2048, 4096, 8192`
Number of Receive Buffers	Specifies the number of buffers allocated for receiving standard MTU size (1514 byte) packets. Registry Key: `*ReceiveBuffers` Values: `1024, 2048,4096,8192,16384,32768.` Default: `16384`
Number of Transmit buffers	Specifies the number of stage buffers used by the driver during transmit. Registry Key: `*TransmitBuffers` Values: `1024, 2048, 4096` Default: `1024`
Priority & VLAN Tag	Enables and disables support for 802.1pQ priority tagging. This property must be enabled to set the VLAN ID. Registry Key: `PQTagging` Default: `Priority and VLAN Enabled` Values: `Priority and VLAN Enabled, Priority and VLAN Disabled, Priority Enabled and VLAN Enabled, Priority Disabled` and `VLAN Disabled`
Receive Side Scaling	Enables and disables the RSS feature. Registry Key: `RSS` Default: `Enabled` Values: `Enabled, Disabled`

Table 2-1. Windows Driver Configurable Parameters (Continued)

Property	Description
Receive Side Scaling Rings	Specifies the number of RSS rings used. Registry Key: `MaxStatusRings` Default: `2` Values: `1-4`
TCP Checksum Offload IPv4	Enables and disables TCP transmit and receive checksum offload. Registry Key: `*TCPChecksumOffloadIPv4` Default: `Rx and Tx Enabled` Values: `Disabled, Rx Enabled, Tx Enabled, Rx and Tx Enabled`
TCP Checksum Offload IPv6	Enables and disables the TCP transmit and receive checksum offload. Registry Key: `*TCPChecksumOffloadIPv6` Default: `Rx and Tx Enabled` Values: `Disabled, Rx Enabled, Tx Enabled, Rx and Tx Enabled`
UDP Checksum Offload IPv4	Enables and the disables user datagram protocol (UDP) transmit and receive checksum offload. Registry Key: `*UDPChecksumOffloadIPv4` Default: `Rx and Tx Enabled` Values: `Disabled, Rx Enabled, Tx Enabled, Rx and Tx Enabled`
UDP Checksum Offload IPv6	Enables and disables the UDP transmit and receive checksum offload. Registry Key: `*UDPChecksumOffloadIPv6` Default: `Rx and Tx Enabled` Values: `Disabled, Rx Enabled, Tx Enabled, Rx and Tx Enabled`
VLAN ID for setting	If the Priority and VLAN Tag is enabled, this parameter specifies a VLAN ID for this interface (also exposed through the standard object identifier [OID]). Registry Key: `VlanId` Default: `0 (no VLAN)` Range: `0-4094`

Configuring NIC Driver Parameters with QCC GUI

Refer to the *QConvergeConsole Help System: Setting General NIC Port Parameters and Setting Advanced NIC Port Parameters* for information.

Configuring NIC Driver Parameters with QCC Interactive CLI

Refer to the *QConvergeConsole CLI User's Guide:* section 7, “NIC Interactive Commands,” for information.

Configuring NIC Driver Parameters with QCC Non-Interactive CLI

Refer to the *QConvergeConsole CLI User's Guide,* section 4, “NIC Noninteractive commands,” for information.

VLAN Configuration

VLAN Configuration with QCC GUI

Refer to the *QConvergeConsole Help System: Setting Advanced NIC Port Parameters* for information.

VLAN Configuration with the QCC Interactive CLI

Refer to the appropriate section in the *QConvergeConsole CLI User's Guide* for information.

VLAN Configuration with the QCC Non-Interactive CLI

Refer to the appropriate section in the *QConvergeConsole CLI User's Guide* for information.

Teaming/Bonding

Windows Teaming

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 locally administered address (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 Modes** (802.3ad static, 802.3ad dynamic (active and passive Link Aggregation Control Protocol [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 2-2 shows that failsafe and transmit load-balancing modes are switch independent, which means they do not require switch configuration. LACP or 802.3ad require switch ports configured for LACP.

Table 2-2. Windows Teaming Modes

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

Table 2-2. Windows Teaming Modes (Continued)

Mode	Failover Capability	Switch Dependency	System Fault Tolerance (SFT)	Load Balancing	Number of Ports per Team (Range [a])
Dynamic 802.3ad	Yes	Yes	Yes	Yes	2–16

[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. Failsafe 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 *admin* 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, and so on.

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

All the adapters in the team share a common team MAC address. This is either a locally administered MAC address or 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 ARPs.

Link Aggregation Mode

Link aggregation provides increased bandwidth and high reliability by combining several NICs into a single, logical, network interface called a link aggregation group (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 one to two 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 (SLA)

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 teaming 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 (DLA)

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 teaming 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 rebalancing 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 the QConvergeConsole utility.

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 the following 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>
```

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 open 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 10Gigabit Ethernet** adapter, and then select **Properties**.

3. Click the **Team Management** tab to bring that page to the front (Figure 2-5) and perform teaming-related management.

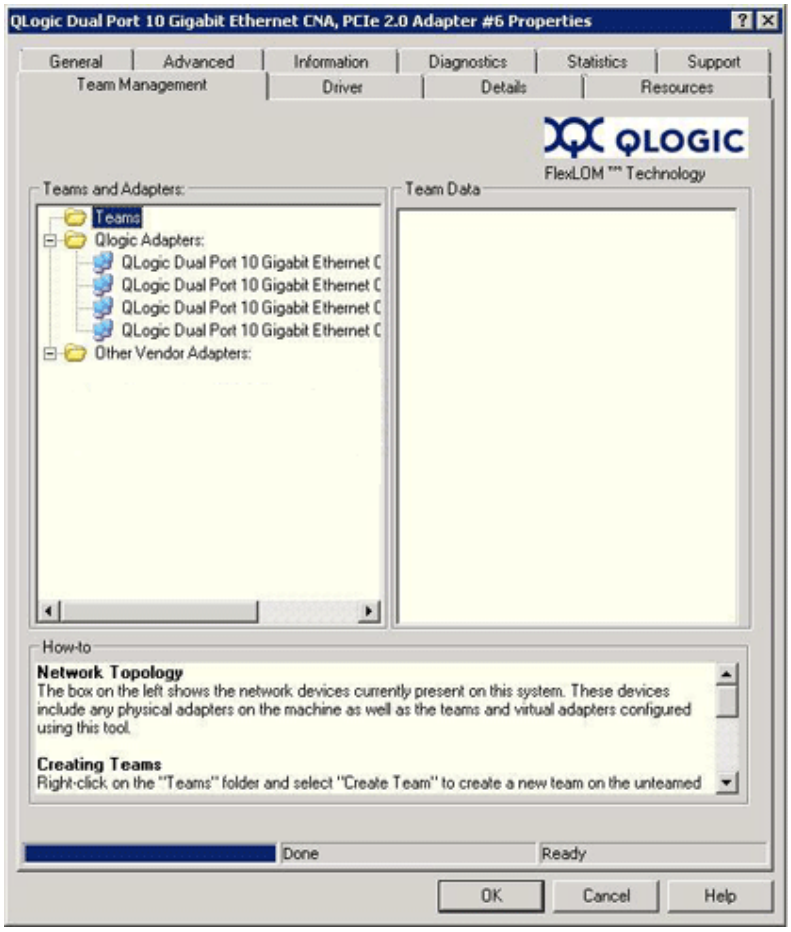

Figure 2-5. Team Management Property Page

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

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.

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:

1. To create a team, right-click the **Teams** folder icon, and then click **Create Team** (Figure 2-6).

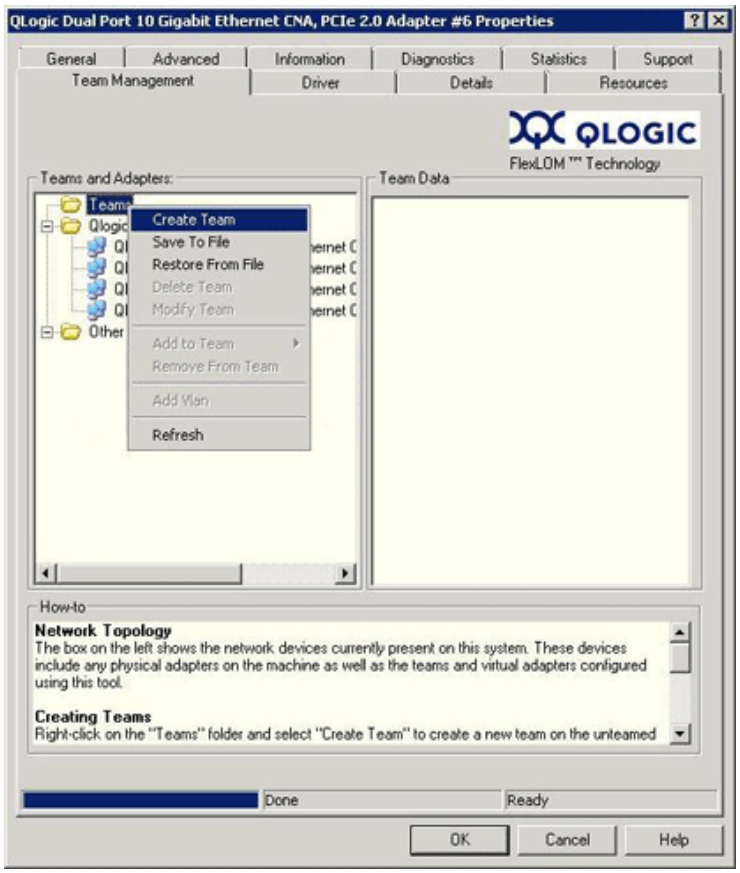

Figure 2-6. Creating a Team

2. The software automatically picks a unique team name, or you may choose to 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 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.
   - **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 configure the member adapters to avoid any conflict after a team has been created.

The following figures show the configuration of various teaming modes.

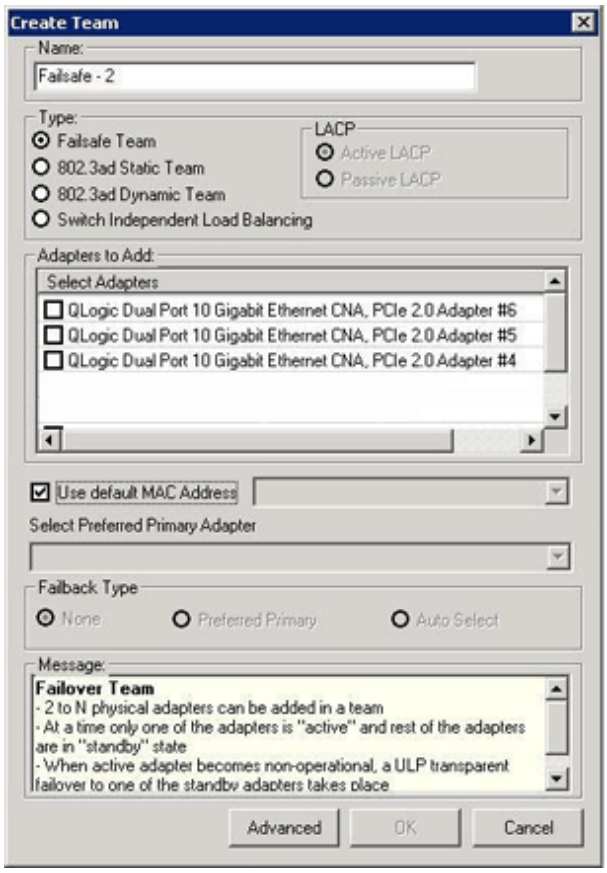

Figure 2-7. Creating a Failsafe Team

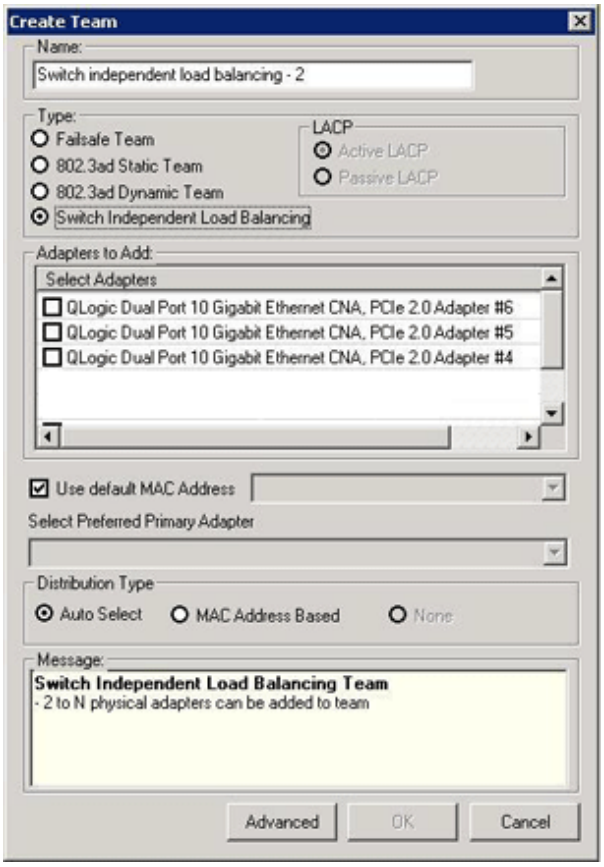

Figure 2-8. Creating a Switch-Independent Load Balancing Team

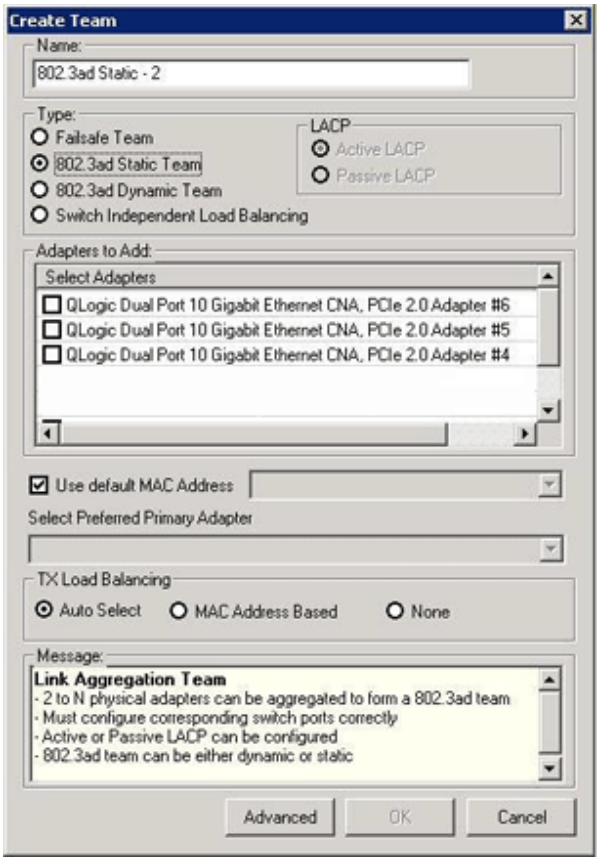

Figure 2-9. Creating an 802.3ad Static Team

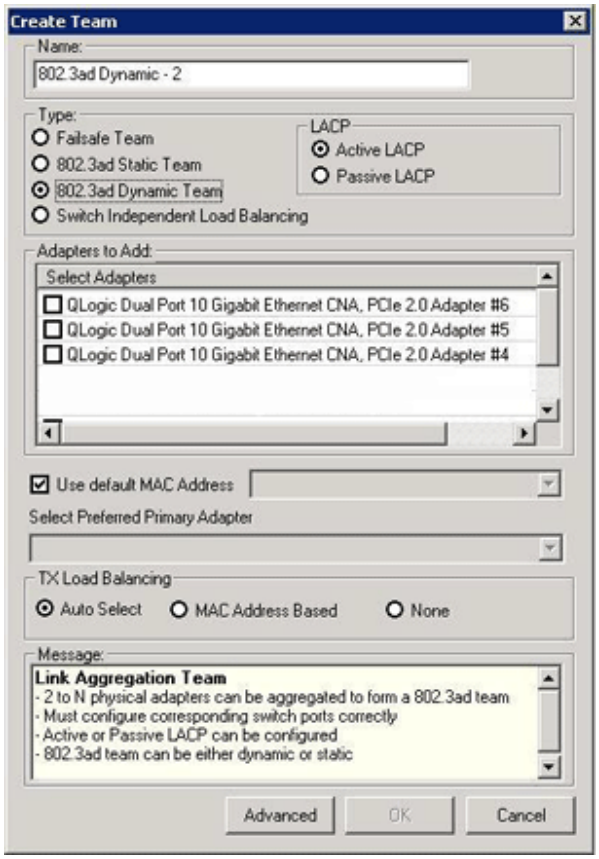

Figure 2-10. Creating an 802.3ad Dynamic Team

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

Figure 2-11 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.

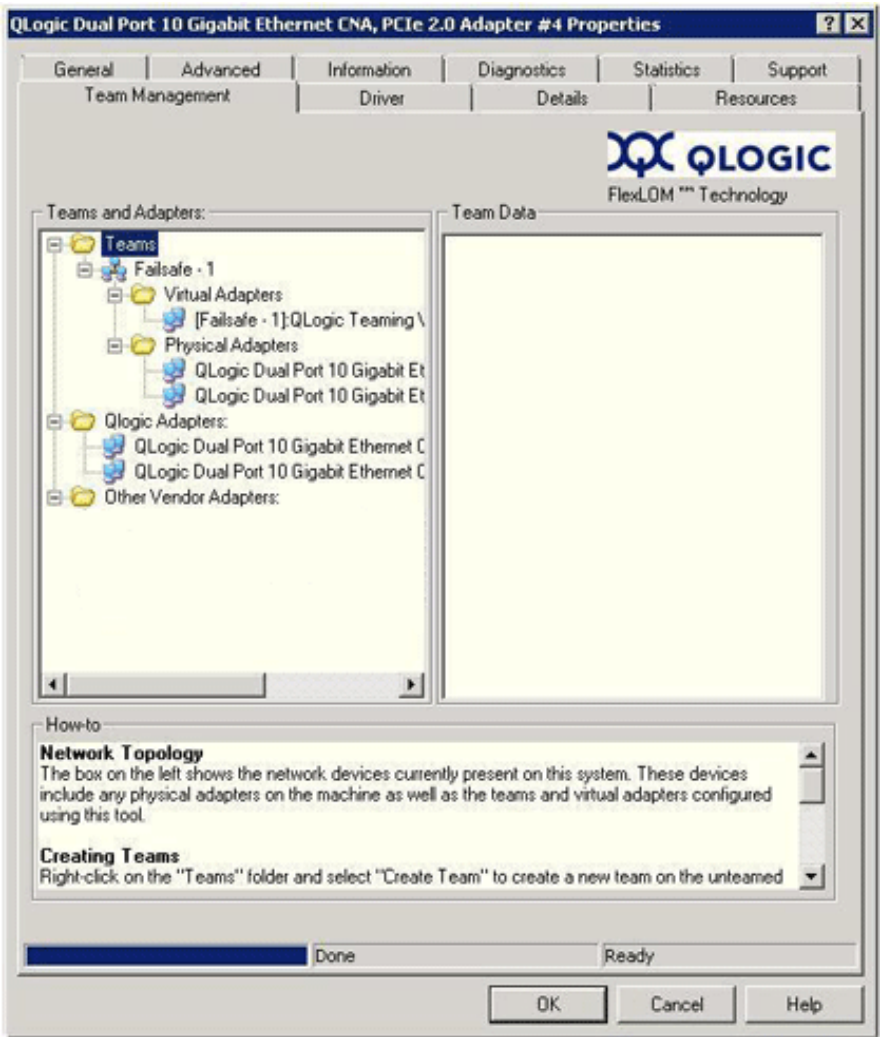

Figure 2-11. Confirming New Team Creation

Modifying a Team

A team can be modified by:

- 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 (Figure 2-12).

NOTE:

You cannot add an adapter to a team that is already a member of another team. Teaming of teams is not supported.

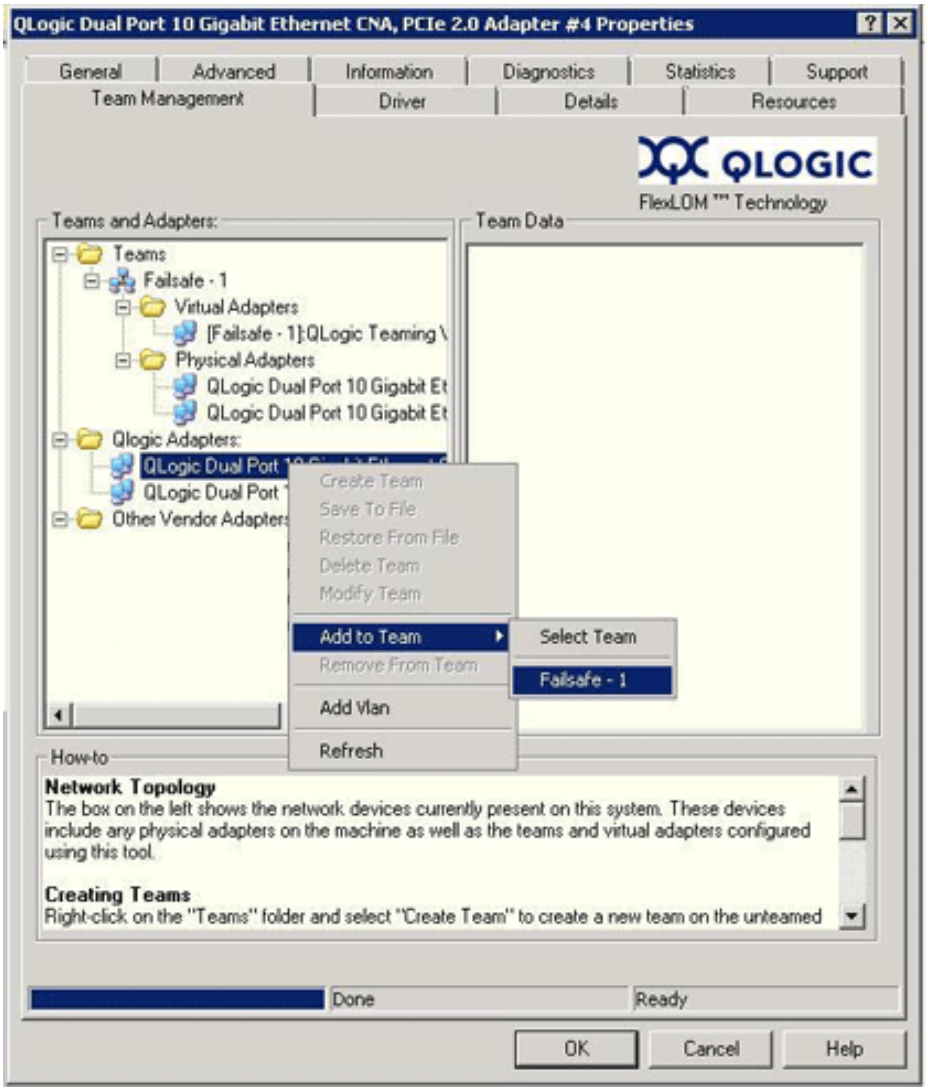

Figure 2-12. Adding a Team

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**.
3. 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 (Figure 2-13), specify a new property value, and then click **OK**.

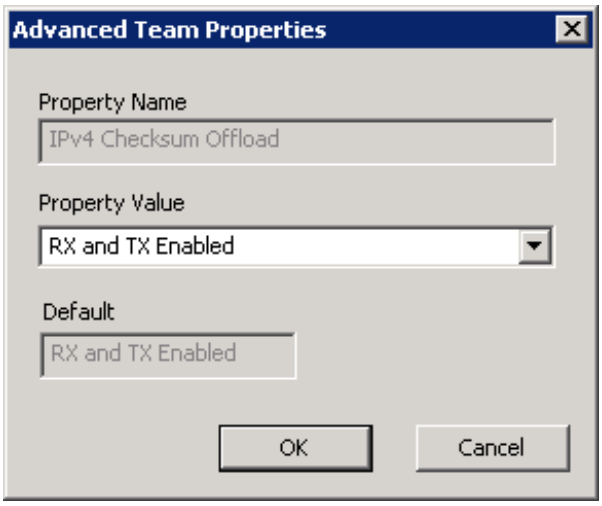

Figure 2-13. Modifying Advanced Team Properties

The team properties change takes effect immediately. Changing team properties causes the driver to reload, which could 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 out of sync 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** (Figure 2-14).

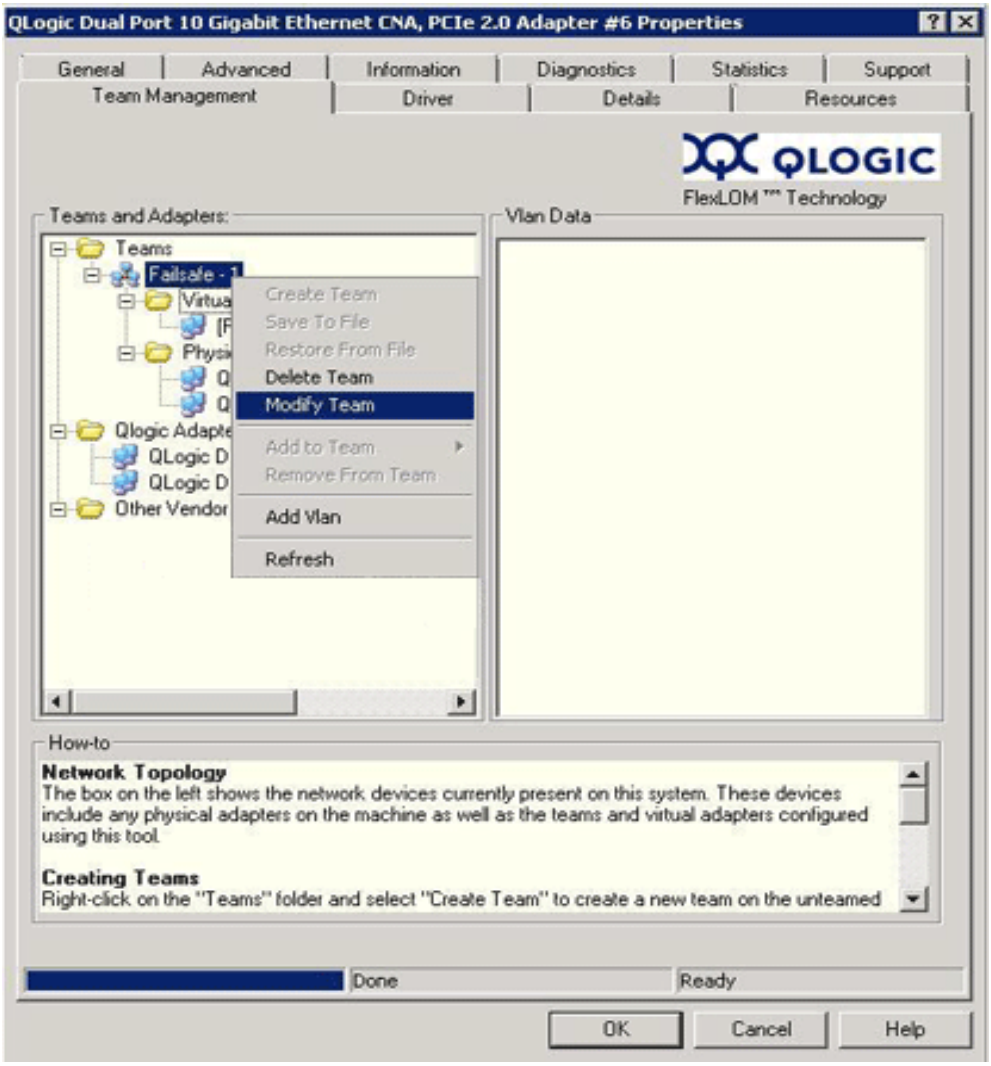

Figure 2-14. Modifying Team Properties

3. In 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 (Figure 2-15).

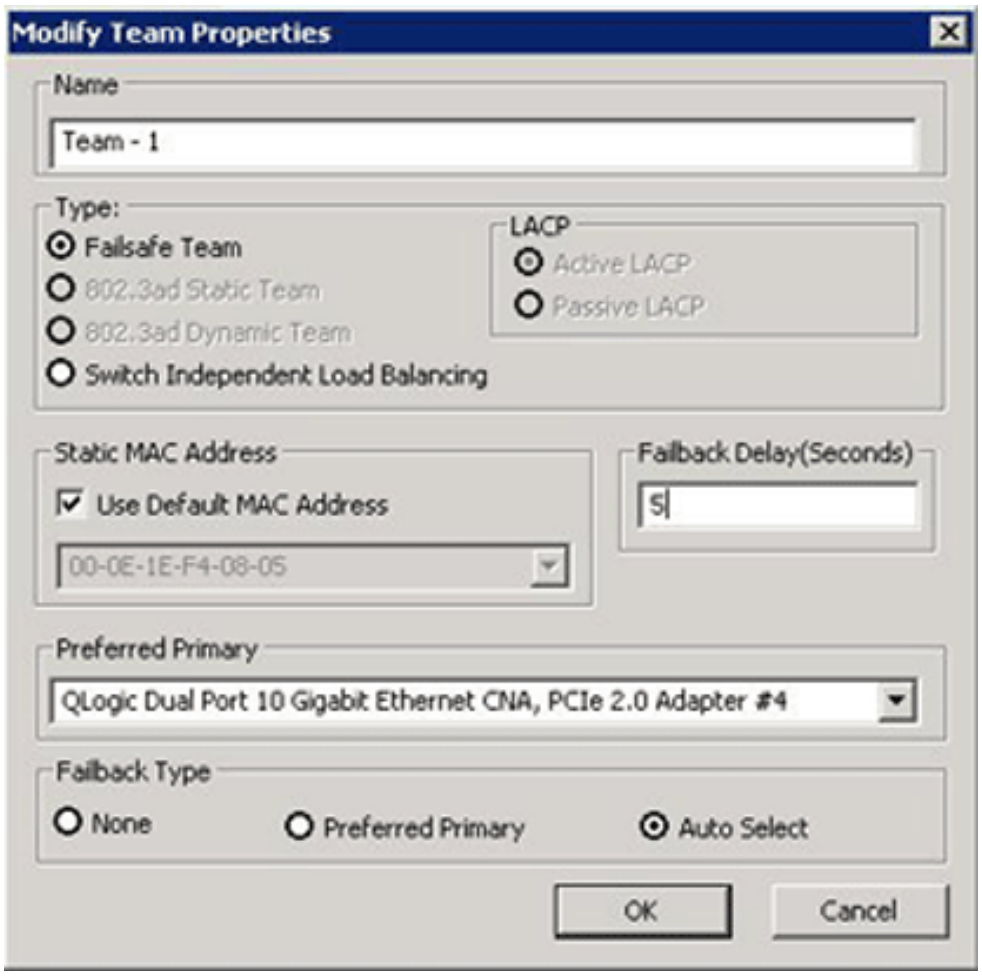

Figure 2-15. Modifying Failsafe Team Properties

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 2-16 shows a failsafe team modification, which shows the new team type and default values for team attributes. You can manually change attribute values.

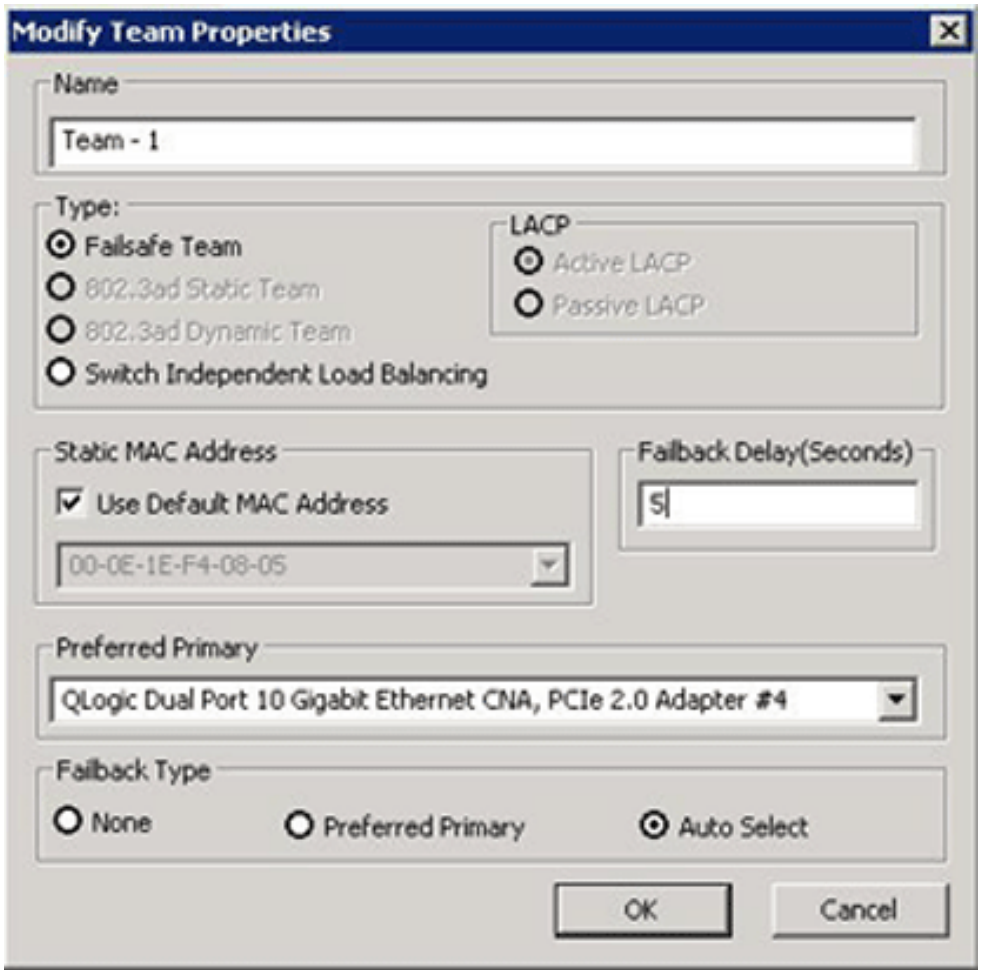

Figure 2-16. Modifying the Team Type

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

QLogic recommends 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 (QCC) CLI utility.

To view teaming statistics, 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]
```

Linux Bonding/Failover/Aggregation

The Linux `qlcnic` driver supports all the standard bonding modes supported by the Linux bonding driver for bonding, failover, and aggregation. Refer to the Linux bonding driver documentation for additional details on the bonding modes.

NIC Partition (NPAR)

The NPAR feature in QLogic adapters provides the ability to create multiple PCIe physical functions for each physical 10GbE port on the adapter. Each of the PCIe functions appears as an independent interface to the host operating system or hypervisor.

Dual-port, 10Gb adapters configured as Ethernet-only adapters contain eight Ethernet functions. 8200 Series and 3200 Series adapter configurations can contain up to four Ethernet functions corresponding to each physical port.

Configuration and Operation

By default, NPAR functionality is disabled on the adapters. Only two Ethernet functions are 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 functions.
- Functions 4 and 5 are iSCSI functions if iSCSI is enabled.
- Functions 6 and 7 are FCoE functions if FCoE is enabled.
- If iSCSI and FCoE functions are not enabled, they can be configured as NIC functions.
- Only one iSCSI function per physical port.
- Only one FCoE function per physical port.

The PCI function to port mapping is also fixed as follows:

- Functions 0, 2, 4, and 6 share one physical port.
- Functions 1, 3, 5, and 7 share the other physical port.

When performed, the following operations require a system reboot to take effect:

- Enabling or disabling NPAR mode
- Changing function type:
  - iSCSI to NIC
  - FCoE to NIC
  - NIC to FCoE

Perform NPAR configuration using any of the following methods:

- OptROM
- Lifecycle controller
- Managements tools under various OSs

After you have configured NPAR mode, you can modify the minimum and maximum bandwidth per NPAR. The changes take effect immediately; no reboot is required. The minimum and maximum bandwidths are specified as percentages of the link bandwidth, where:

- Minimum bandwidth is the minimum guaranteed bandwidth to a partition.
- Maximum bandwidth is the maximum value that a partition is permitted to use.

NPAR mode supports iSCSI boot, FCoE boot, and PXE boot in their respective partitions.

NOTE:

For additional information on configuring NPAR in the QConvergeConsole utility, refer to the *QConvergeConsole Help System*.

Using NPAR under ESX

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 command line interfaces (remote CLIs).

All NPAR 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 NIC Partition (NPAR).

You would typically create a virtual switch (vSwitch) for each NPAR 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 NPAR interfaces as uplinks, it is not useful to have multiple uplinks with NPARs belonging to the same physical port.

The vSwitch switches the packets when VMs are assigned to share the same NPAR. The embedded switch (eSwitch) on the adapter switches packets when packets are sent by a VM on an NPAR destined to another VM that is connected to a vSwitch on a different NPAR belonging to the same physical port. If the two NPARs 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 8200 and 3200 Series Adapters have 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.

Configuring NPAR in the BIOS

Refer to the Type section for information.

Configuring NPAR in QCC GUI

Refer to the *QConvergeConsole Help System: Partitioning NIC Ports* for information.

Configuring NPAR in QCC Interactive CLI

Refer to the *QConvergeConsole CLI User's Guide* for information.

Configuring NPAR in QCC Non-Interactive CLI

Refer to the *QConvergeConsole CLI User's Guide,* "NIC Noninteractive Commands" section for information.

Configuring NPAR in the QLogic Adapter vCenter Plug-in

The QLogic Adapter vCenter Plug-in provides the QLogic Adapter tab. Use this tab to manage the 8200 and 3200 Series Adapters in the VMware ESX/ESXi environment.

With the QLogic Adapter tab, you can:

- Enable/Disable NPAR
- Enable/Disable functions
- Configure function personalities
- Configure the relative bandwidth of NIC functions
- Configure eSwitch
- View the eSwitch Statistics

The default configuration of the 8200 Series Adapters is a NIC function, iSCSI function, and FCoE function. There is an additional NIC function that is defaulted as disabled, but is presented to the user to be optionally enabled.

Figure 2-17 shows the default view of QLogic Adapter tab in the vCenter Plug-in.

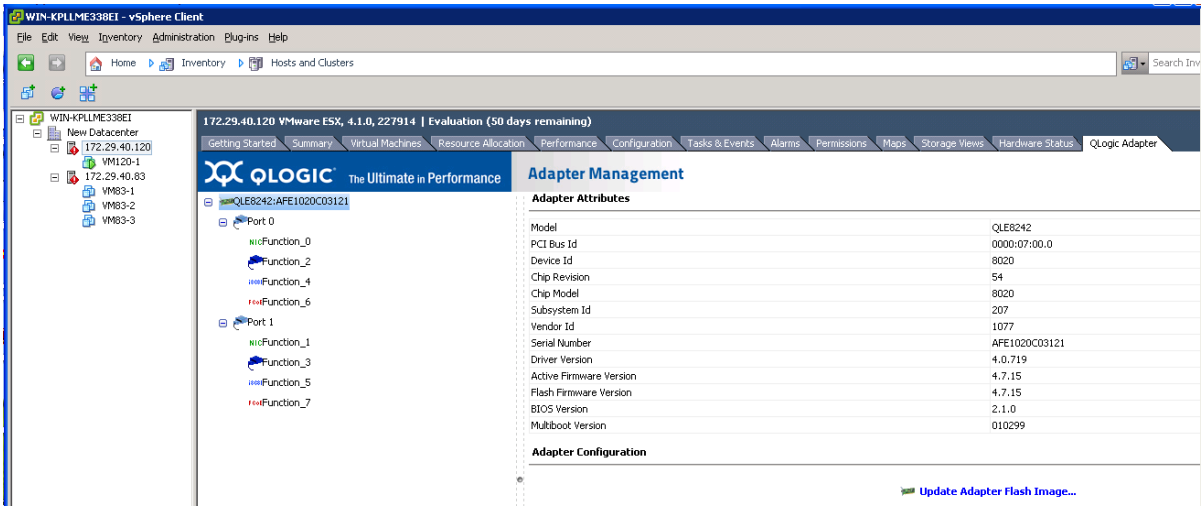

Figure 2-17. vCenter Plug-in QLogic Adapter Tab Default View

Figures 2-18 shows the configuration range of the 8200 Series Adapters without NPAR.

Figure 2-18. 8200 Series Configuration Range Without NPAR

Figures 2-19 shows the configuration range of the 8200 Series Adapters with NPAR.

PF0/1 NIC	PF2/3 NIC	PF4/5 iSCSI	PF6/7 FCoE
PF0/1 NIC	PF2/3 NIC	PF4/5 NIC	PF6/7 FCoE
PF0/1 NIC	PF2/3 NIC	PF4/5 iSCSI	PF6/7 NIC
PF0/1 NIC	PF2/3 NIC	PF4/5 NIC	PF6/7 NIC
PF0/1 NIC	PF2/3 NIC	PF4/5NIC	PF6/7 Disabled
PF0/1 NIC	PF2/3 NIC	PF4/5 Disabled	PF6/7 Disabled
PF0/1 NIC	PF2/3 Disabled	PF4/5 Disabled	PF6/7 Disabled

Figure 2-19. 8200 Series Configuration Range With NPAR

Figures 2-20 shows the Configuration Range of the 3200 Series Adapters with NPAR. Physical functions 2–7 can be independently disabled.

PF0/1 NIC	PF2/3 Disabled	PF4/5 NIC	PF6/7 NIC
PF0/1 NIC	PF2/3 Disabled	PF6/7 Disabled	PF6/7 Disabled
PF0/1 NIC	PF2/3 Disabled	PF4/5 NIC	PF6/7 Disabled
PF0/1 NIC	PF2/3 Disabled	PF4/5 NIC	PF6/7 NIC

Figure 2-20. 3200 Series Configuration Range With NPAR

To configure a physical function:

1. Select the function you want to change in the hierarchy on the left side of the screen.

   The Function Attributes on the right side of the screen reflect the selected physical function.

2. Make a selection from the **Function Type** menu.

   a. If the Function Type is NIC, iSCSI, or FCoE, select **Disabled** to disable the function.

   b. If the Function Type is Disabled, select the appropriate function type.

In the example in Figure 2-21, function 2 is disabled. To enable the function, NIC is selected from the Function Type menu.

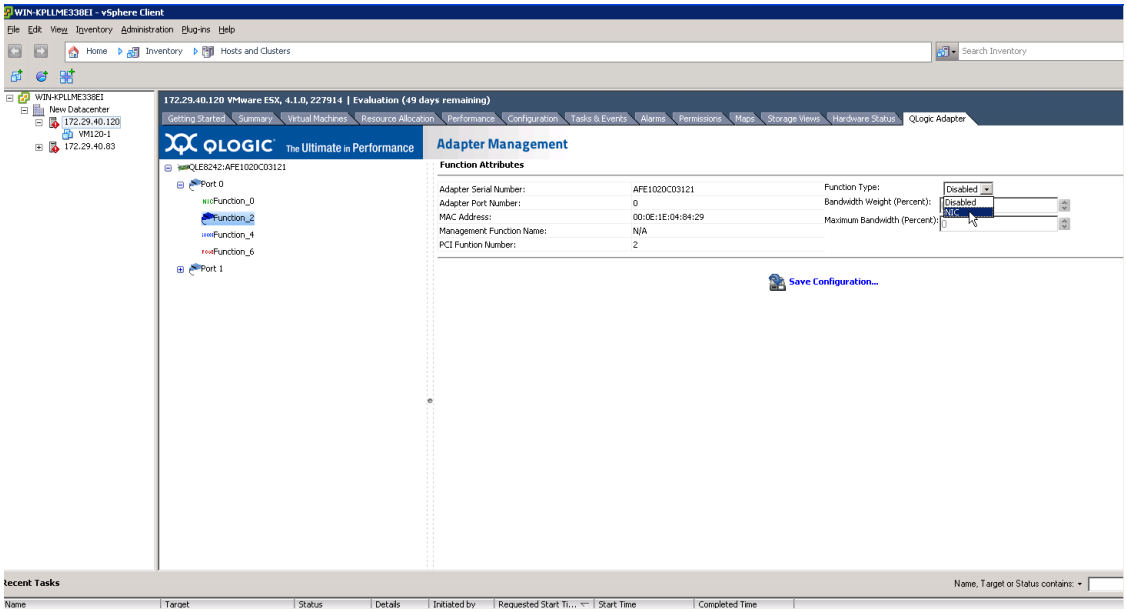

Figure 2-21. Enabling a Physical Function

3. Click **Save Configuration**.
4. Reboot your system for the change to take effect.
5. Open vCenter Plug-in.
6. Click the **QLogic Adapter** tab.
7. Right-click the adapter icon on the left side of the page, and then select **Refresh**.

The selected function now reflects the change made in Step 2. For example, the change made in Figure 2-21 is reflected in Figure 2-22 (function 2 changed from **Disabled** to **NIC**.

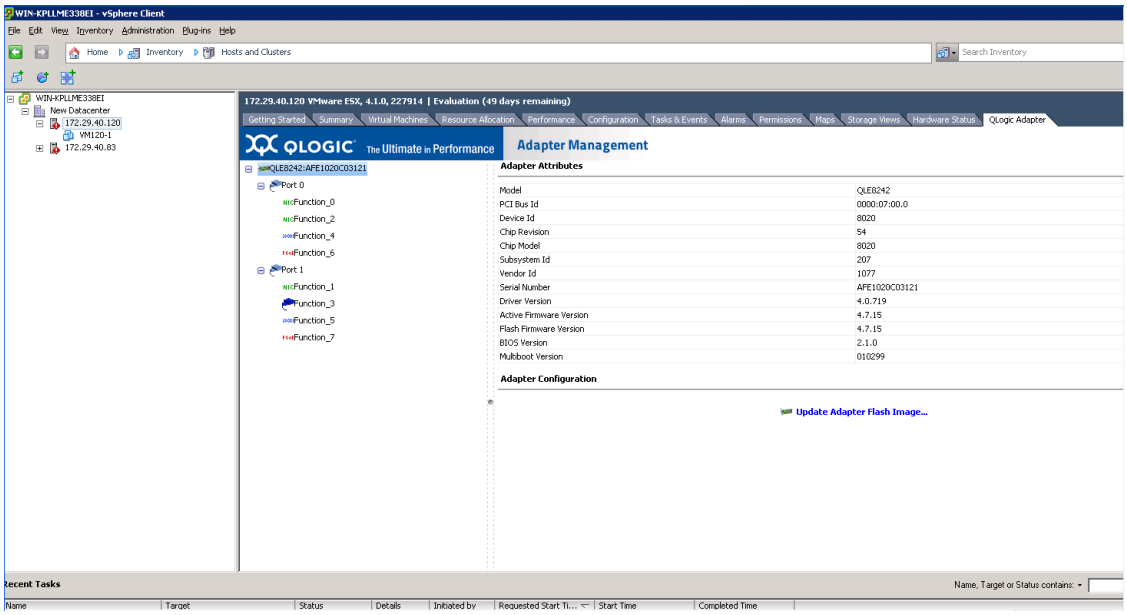

Figure 2-22. Port One Configured with NPAR Enabled

NPAR Personality Changes

- Changing a function's personality requires the ESX Host to be rebooted for the change to take effect.
- Each personality configuration change *must* be saved using the **Save Configuration** button.

  This means that to change two functions' personalities, each change must be separately saved using the **Save Configuration** button.
- Multiple personality changes may be made before rebooting the ESX server after saving the individual changes.
- The Function Type options for the FCoE and iSCSI capable functions do not reflect the total range of personalities that these ports may transition into.
- To transition from FCoE or iSCSI to NIC, the FCoE or iSCSI port must be first disabled, rebooted, configured as NIC, and then rebooted again.

Enabling NPAR

To enable NPAR for a port in the QLogic vCenter Plug-in, do the following:

1. Select function 2 for port 0 and function 3 for port 1.
2. Click **OK** to accept the *NPAR parameter(s) Successfully Changed!* confirmation message.
3. Reboot the ESX Host.

NOTE:

This is the Server where the Converged Network Adapter or Intelligent Ethernet Adapter is installed.

4. Wait until the ESX Host again becomes connected in the Data Center. The ESX Host's Name (IP Address) will be italicized (not responding) while the ESX Host reboots.
5. After the ESX Host connects again, on the shortcut menu on the QLogic Adapter tab, click **Refresh**.

Function 2 of port 0 is now configured for NIC; this enables NPAR.

- NPAR enabled on port 0—function 2 on port 0 is enabled.
- NPAR disabled on port 1—function 3 on port 1 is disabled.

NPAR is now enabled, and the configuration range of the FCoE and iSCSI storage functions is extended. The FCoE Function Type options are extended FCoE, NIC, and Disabled (see Figure 2-22).

Quality of Service

The QLogic vCenter Plug-in provides control of the NPAR quality of service (QoS) capability of the 8200 and 3200 Series Adapters. QoS is enabled whenever FCoE is not enabled. This is because FCoE requires ETS, which supersedes NPAR's QoS.

The vCenter QoS controls allocate bandwidth to NIC functions in increments of 1 percent of 10Gb. There is a Maximum Bandwidth control and a Bandwidth Weight. The Maximum Bandwidth control specifies the maximum bandwidth that the NIC function is allowed take in congested conditions. A NIC function can take more bandwidth than the Maximum Bandwidth setting if there is no other demand for that bandwidth. The Bandwidth control specifies the minimum band width that a NIC function is guaranteed in congested conditions.

See Figure 2-23 for an illustration.

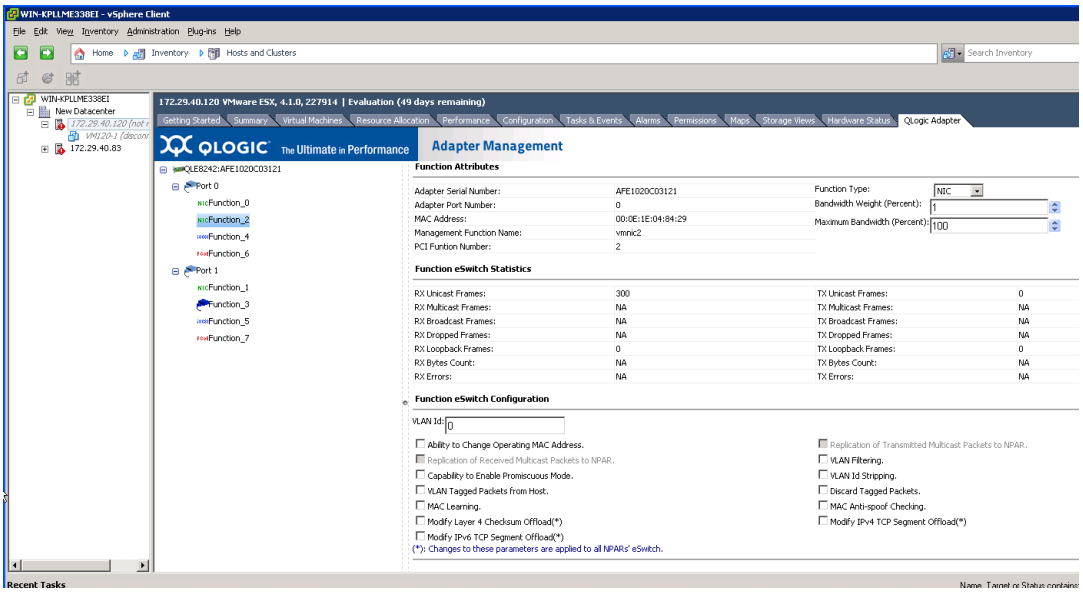

Figure 2-23. QLogic Adapter vCenter Plug-in with QoS Interface, Function eSwitch Configuration, and eSwitch Statistics

eSwitch

The 8200 and 3200 Series adapters connect the physical ports to the physical functions on the PCIe bus through eSwitches. An eSwitch is a VLAN aware MAC lookup virtual switch implemented in the 8200 and 3200 Series Adapters. Each port provides the uplink for an eSwitch with the enabled associated Physical PCI Functions connected through switch ports to the eSwitch. Functions 0, 2, 4, and 6 are connected through the port 0 eSwitch; and functions 1, 3, 5, and 7 are connected through the port 1 eSwitch. The eSwitch associated with a port will only conduct traffic between PCIe Ethernet functions if the physical port is linked.

Control of the eSwitch is through the QLogic vCenter Plug-in. The eSwitch management controls support VLAN configuration, promiscuous mode enable, locally administered MAC addresses, IPv4 checksum offload enable, IPv4 and IPv6 segmentation offloads, and MAC anti-spoof checking.

NIC TroubleShooting/Diagnostics

NIC Linux Diagnostics

Linux Diagnostics

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

NOTE:

The current release does *not* support external and internal loopback tests. Support is expected in a future release.

Running Linux User Diagnostics

Linux user diagnostics include QConvergeConsole diagnostics and ethtool diagnostics.

QConvergeConsole Diagnostics

NOTE:

Information on installing and starting the GUI version of QLogic's QConvergeConsole utility is provided in the *QConvergeConsole User's Guide*. All procedural information for that utility is covered in the *QConvergeConsole 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

- **Internal Loopback Test** performs internal packet loopback.
- **Flash Test** verifies the Flash read and write.
- **Hardware Test** verifies that the hardware is running.
- **Interrupt Test** enables and disables the interrupt and functional verification tests.
- **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.
- **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.

Windows Property Page Diagnostics

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

NOTE:

The current release does not support external and internal loopback tests. Support is expected in a future release.

Running Windows User Diagnostics

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

NOTE:

Information on installing and starting the GUI version of QLogic's QConvergeConsole utility is provided in the *QConvergeConsole User's Guide*. All procedural information for that utility is covered in the *QConvergeConsole 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 2-24 shows the **Diagnostics** page.

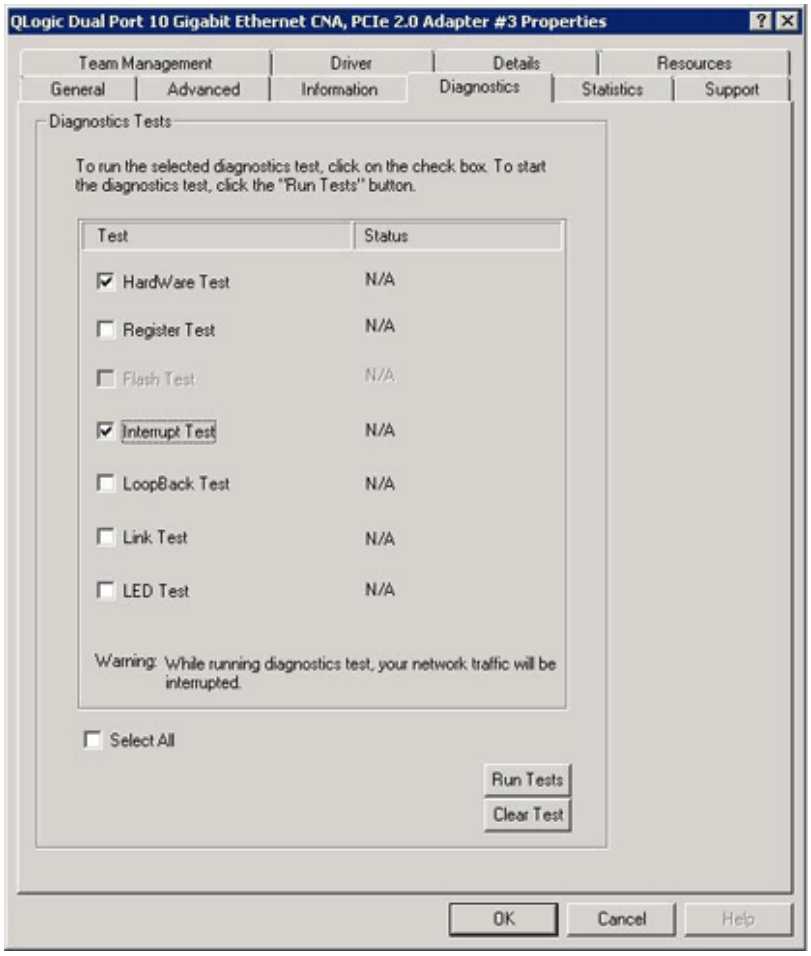

Figure 2-24. Diagnostics Tests on Windows

4. Under **Diagnostic Tests**, select one or more check boxes indicating the tests you want to run: **Hardware Test**, **Register Test**, **Interrupt Test**, **Loopback Test**, and **Link Test**. (Windows Diagnostic Test Descriptions 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: NIC, iSCSI, or Fibre Channel. If you do not specify a protocol, all protocols are enabled by default. Table 2-3 and Table 2-4 list the QConvergeConsole commands for selecting a protocol.

Table 2-3. Windows QConvergeConsole CLI—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 NPAR command line options

Table 2-4. Windows QConvergeConsole CLI—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

Table 2-4. Windows QConvergeConsole CLI—Selecting a Protocol in Legacy Mode (Continued)

Command	Description
`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 2-5.

Table 2-5. Windows QConvergeConsole CLI—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
`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 commands usage, and then exit

Table 2-6 lists miscellaneous Windows diagnostics commands.

Table 2-6. Windows QConvergeConsole CLI—Miscellaneous Commands

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

Table 2-7 lists the Windows CLI diagnostic test commands. Note that while running these tests, network traffic is interrupted.

Table 2-7. Windows QConvergeConsole CLI—Diagnostic Test Commands

Command	Description
`-i` `--interface`	Specifies the interface type (NX_NIC, NX_NIC1, and so on)

Table 2-7. Windows QConvergeConsole CLI—Diagnostic Test Commands (Continued)

Command	Description
`-a` `--all`	Perform all test, regardless of default
`-D` `--default`	Perform only the 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`)
`-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

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 2-8.

Table 2-8. 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]`
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, timeout_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
- Register Test
- Interrupt Test
- Loopback Test
- Link Test

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.

Loopback Test

The loopback test is a diagnostic tool that routes transmit data through a loopback connector back to the same adapter.

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 2-9. Note that this table does not list error messages for the interrupt and link tests.

Table 2-9. Windows Diagnostic Test Messages

Test	Error Message	Description
Loopback	LB_TEST_OK	Loopback test has passed
	LB_SEND_WAIT_QUEUE_ERR	Send queue blocked
	LB_NORCV_ERR	Receive packet not received
	LB_NOMEM_ERR	No memory error
	LB_TX_QUEUE_ERR	Transmit queue error
	LB_SHORT_DATA_ERR	Looped data short error
	LB_SEQUENCE_ERR	Looped data out of sequence
	LB_DATA_ERR	Looped data corrupted
	LB_ERR_CNT	Looped error count
Register	CR_TEST_OK	Control register test passed
	CR_NIU_MODE	Network interface unit (NIU) error
	CR_PHY	Physical layer (PHY) error
	CR_ERRCNT	Control register error count
Hardware	HW_TEST_OK	Hardware test has passed
	HW_DMA_BZ_0	DMA channel 0 is busy
	HW_DMA_BZ_1	DMA channel 1 is busy
	HW_DMA_BZ_2	DMA channel 2 is busy
	HW_DMA_BZ_3	DMA channel 3 is busy

Table 2-9. Windows Diagnostic Test Messages (Continued)

Test	Error Message	Description
	HW_SRE_PBI_HALT	Segmentation and reassembly engine currently halted
	HW_SRE_L1IPQ	Segmentation and reassembly engine currently paused due to L1 IPQ discard failure
	HW_SRE_L2IPQ	Segmentation and reassembly engine currently paused due to L2 IPQ discard failure
	HW_SRE_FREEBUF	Segmentation and reassembly engine free buffer list is currently empty
	HW_IPQ	IPQ is currently not empty
	HW_PQ_W_PAUSE	PQ write pause previously detected
	HW_PQ_W_FULL	PQ write full previously detected
	HW_IFQ_W_PAUSE	IFQ write pause previously detected
	HW_IFQ_W_FULL	IFQ write full previously detected
	HW_MEN_BP_TOUT	Memory backpressure timeout previously detected
	HW_DOWN_BP_TOUT	Downstream backpressure timeout previously detected
	HW_FBUFF_POOL_WM	Free buffer pool low watermark previously detected
	HW_PBUF_ERR	Packet buffer error previously detected
	HW_PBUF_ERR	Packet buffer error previously detected
	HW_FM_MSG_HDR	FM message header error previously detected
	HW_FM_MSG	FM message error previously detected
	HW_EPG_CTRL_Q	Egress packet generator (EPG) control queue is backed up
	HW_EPG_MSG_BUF	EPG message buffer error

Table 2-9. Windows Diagnostic Test Messages (Continued)

Test	Error Message	Description
	HW_EPG_QREAD_TOUT	EPG read queue timeout
	HW_EPG_QWRITE_TOUT	EPG write queue timeout
	HW_EPG_CQ_W_FULL	EPG completion queue write full
	HW_EPG_MSG_CHKSM	Egress packet generator (EPG) message checksum error
	HW_EPG_MTLQ_TOUT	EPG MTL queue fetch timeout
	HW_PEG0	PEG 0 is not used
	HW_PEG1	PEG 1 is not used
	HW_PEG2	PEG 2 is not used
	HW_PEG3	PEG 3 is not used
	HW_ERRCNT	Hardware error count

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
```

QCC GUI Diagnostics

Refer to the *QConvergeConsole Help System: Performing NIC Port Diagnostics* for information.

QCC CLI Diagnostics QCC Interactive CLI

Refer to the *QConvergeConsole User's Guide* for information.

QCC CLI Diagnostics QCC Non-Interactive CLI

Refer to the *QConvergeConsole User's Guide*, section 4, “NIC Noninteractive Commands,” for information.

NIC Troubleshooting Guidelines

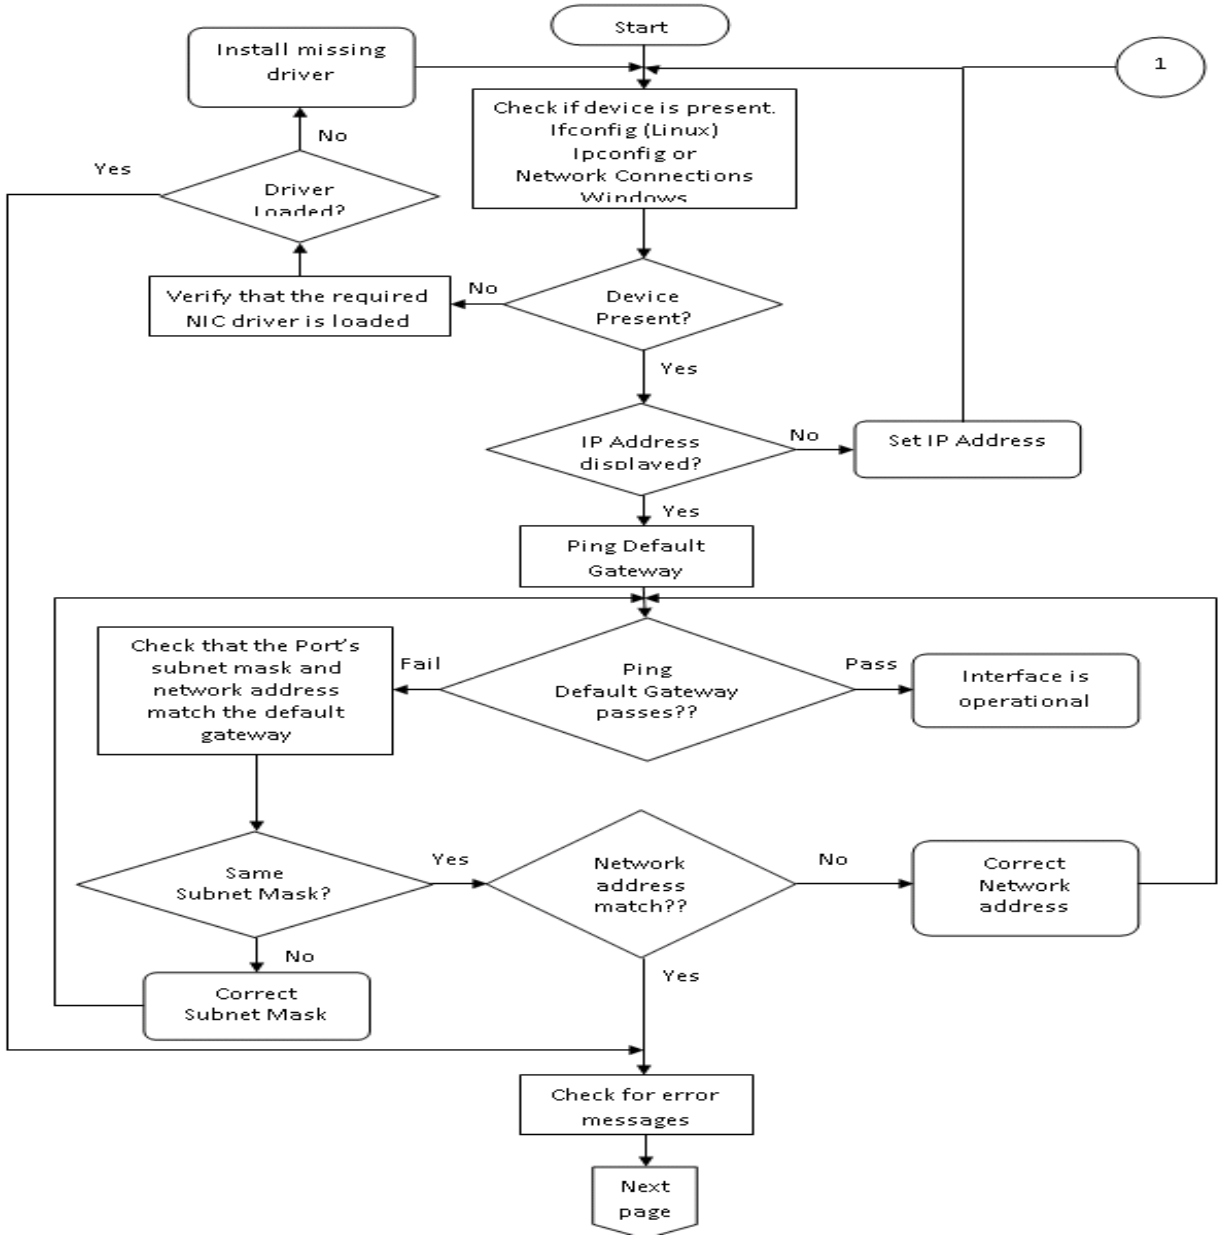

Figure 2-25. NIC Troubleshooting Diagram 1

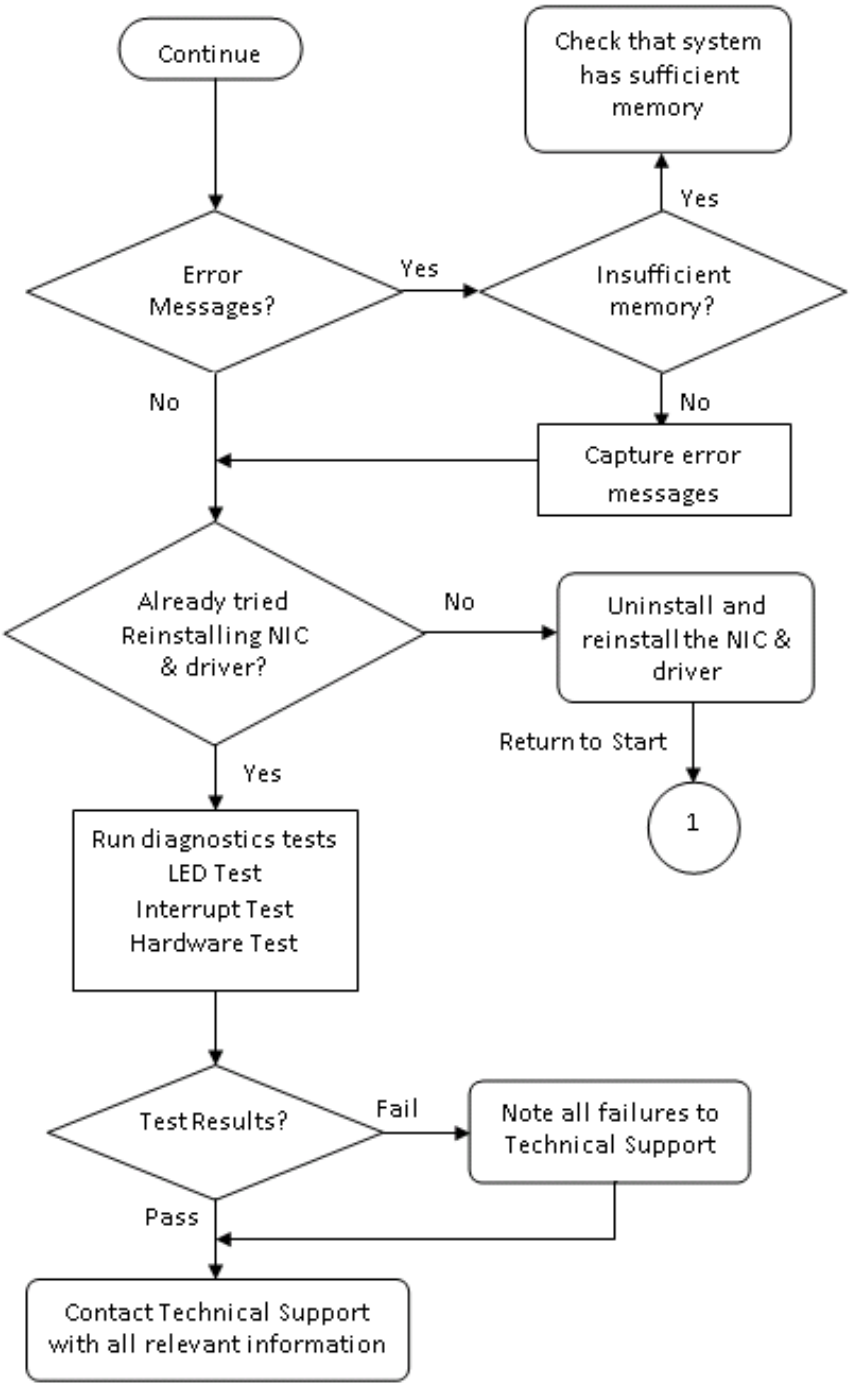

Figure 2-26. NIC Troubleshooting Diagram 2

3 Configuring iSCSI

iSCSI Overview

This section describes the driver and agent installation, configuration, operation, and troubleshooting of the iSCSI function of the 8200 Series Converged Network Adapter.

Installing iSCSI in Linux

Refer to the *iSCSI Adapter Inbox Driver Update for Linux Kernel 2.6.x* Readme file for information.

Packaging Content

Refer to the *iSCSI Adapter Inbox Driver Update for Linux Kernel 2.6.x* Readme file for information.

Linux OS Support

Refer to the *iSCSI Adapter Inbox Driver Update for Linux Kernel 2.6.x* Readme file for information.

Supported Features

Refer to the *iSCSI Adapter Inbox Driver Update for Linux Kernel 2.6.x* Readme file for information.

Deploying the Driver

Refer to the *iSCSI Adapter Inbox Driver Update for Linux Kernel 2.6.x* Readme file for information.

Driver System Parameters

Refer to the *iSCSI Adapter Inbox Driver Update for Linux Kernel 2.6.x* Readme file for information.

Installing iSCSI in ESX

Refer to the *iSCSI Driver for VMware ESX/ESXi 4.1* Readme file for information.

Packaging Content

Refer to the *iSCSI Driver for VMware ESX/ESXi 4.1* Readme file for information.

ESX OS Support

Refer to the *iSCSI Driver for VMware ESX/ESXi 4.1* Readme file for information.

Using the Driver

Refer to the *iSCSI Driver for VMware ESX/ESXi 4.1* Readme file for information.

Installing the Driver

Refer to the *iSCSI Driver for VMware ESX/ESXi 4.1* Readme file for information.

Installing iSCSI in Windows

Refer to the *iSCSI Adapter STOR Miniport Driver for Windows* Readme file for information.

Packaging Content

Refer to the *iSCSI Adapter STOR Miniport Driver for Windows* Readme file for information.

Windows OS Support

Refer to the *iSCSI Adapter STOR Miniport Driver for Windows* Readme file for information.

Using the Driver

Refer to the *iSCSI Adapter STOR Miniport Driver for Windows* Readme file for information on how to create the driver disk, install the driver, and remove the driver.

iSCSI Configuration

iSCSI Configuration with QCC GUI

Refer to the *QConvergeConsole Help System* for information.

iSCSI Configuration with Interactive QCC CLI

Refer to the *QConvergeConsole CLI User's Guide* for information.

iSCSI Configuration with Non-Interactive CLI

Refer to the *QConvergeConsole CLI User's Guide*, “iSCSI Noninteractive Commands” section for information.

iSCSI Initiator

Configuring an iSCSI Initiator in Linux

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

To configure a Linux iSCSI initiator:

1. Access the QConvergeConsole CLI by entering **`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 the ENTER key to accept the default.
   b. `DHCP to obtain IPv4 Network Information: [off]:` Press the ENTER key to accept the default.
   c. `IP_Address []:` Type the IP address of the initiator system, and then press the ENTER key.
   d. `IP_Subnet_Mask [255.255.255.0]:` Type the appropriate subnet mask, and then press the ENTER key.

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

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

9. On the options menu that opens, 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 the ENTER key, and then type **p** and press the ENTER key 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 the ENTER key to accept the default.

    b. `TGT_iSCSI_Name []`: Type the iSCSI qualified name (IQN) of the iSCSI target to connect to, and then press the ENTER key.

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

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

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

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

Configuring an iSCSI Initiator in Windows

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 **QCC 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 **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 the ENTER key to accept the default.
   b. `DHCP to obtain IPv4 Network Information: [off]:` Press the ENTER key to accept the default.
   c. `IP_Address []:` Type the IP address of the initiator system, and then press the ENTER key.
   d. `IP_Subnet_Mask [255.255.255.0]:` Type the appropriate subnet mask, and then press the ENTER key.
   e. `IP_Gateway [0.0.0.0]:` Press the ENTER key to accept the default.
   f. `Enable IPv6 [off]:` Press the ENTER key to accept the default.
9. On the options menu that opens, 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 the ENTER key, and then type **`p`** and press the ENTER key 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 the ENTER key to accept the default.
   b. `TGT_iSCSI_Name []:` Type the iSCSI qualified name (IQN) of the iSCSI target to connect to, and then press the ENTER key.
   c. `TGT_Port [3260]:` Press the ENTER key to accept the default.
   d. `TGT_TargetIPAddress [0.0.0.0]:` Type the IP address of the target, and then press the ENTER key.
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 an iSCSI Initiator in VMware

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.

iSCSI Name Server iSNS

Refer to the *Configuring iSNS to Automatically Discover Targets* section in the *QConvergeConsole Help System* for information.

iSCSI Boot

iSCSI Boot Setup Using Fast!UTIL

Fast!UTIL is the QLogic iSCSI BIOS Configuration utility used to configure the iSCSI TCP/IP Offload Engine (TOE).

Accessing Fast!UTIL

To access Fast!UTIL for PXE, iSCSI, and FCoE, respectively, press the CTRL+Q keys during the QLogic 8200 and 3200 Series Adapters BIOS initialization when the PXE, iSCSI, or FCoE QLogic banner displays.

It may 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 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 the C key to clear the boot device information from the primary and alternate boot locations.

Configure the following iSCSI boot setup information:

- Boot Device Primary and Alternate
- Adapter Boot Mode
- Primary and Alternate Boot Device Settings
- Configuring the iSCSI Boot Parameters
- Configuring the iSCSI Boot Settings
- Preparing the Image on the Local Hard Drive
- Transferring the OS Image to the iSCSI Target
- Booting

DHCP Boot Setup for iSCSI Boot

DHCP boot provides a method for the iSCSI initiator to acquire target parameters from a DHCP server. To configure the DHCP server to support iSCSI boot, first ensure that your DHCP server is set up.

NOTE:

This release does not support DHCP iSCSI boot for IPv6. Refer to future Readme and Release Notes for IPv6 support notification.

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 utilizes software-based initiators.

iSCSi Boot Setup Using QCC CLI

Refer to the *QConvergeConsole CLI User's Guide* for information.

Configuring iSCSI Boot using the QCC CLI

Refer to the *QConvergeConsole CLI User's Guide* for information.

Configuring iSCSI Boot using Interactive QCC CLI

Refer to the *QConvergeConsole CLI User's Guide* for information.

Configuring iSCSI Boot using Non-interactive QCC CLI

Refer to the *QConvergeConsole CLI User's Guide*, "iSCSI Noninteractive Commands" section for information.

Configuring iSCSI DHCP Boot Options using QCC CLI

Refer to the *QConvergeConsole CLI User's Guide*, "iSCSI Noninteractive Commands" section for information.

Configuring iSCSI DHCP Boot Options using Interactive QCC CLI

Refer to the *QConvergeConsole CLI User's Guide* for information.

Configuring iSCSI DHCP Boot Options using Non-interactive QCC CLI

Refer to the *QConvergeConsole CLI User's Guide*, "iSCSI Noninteractive Commands" section for information.

iSCSI Boot Setup Using QCC GUI

Refer to the *Configuring Boot Settings* section in the *QConvergeConsole Help System* for information.

Configuring iSCSI Boot using the QCC GUI

Refer to the *Configuring Boot Settings* section in the *QConvergeConsole Help System* for information.

Configuring iSCSI DHCP Boot Options

Refer to the *Configuring Boot Settings* section in the *QConvergeConsole Help System* for information.

iSCSI TroubleShooting

iSCSI Diagnostics

iSCSI Diagnostics using QCC GUI

Refer to *QConvergeConsole Help System: Pinging an iSCSI Port Target* for information.

iSCSI Diagnostics using Interactive QCC CLI

Refer to the *Interactive CLI User's Guide* for information.

iSCSI Troubleshooting Diagram

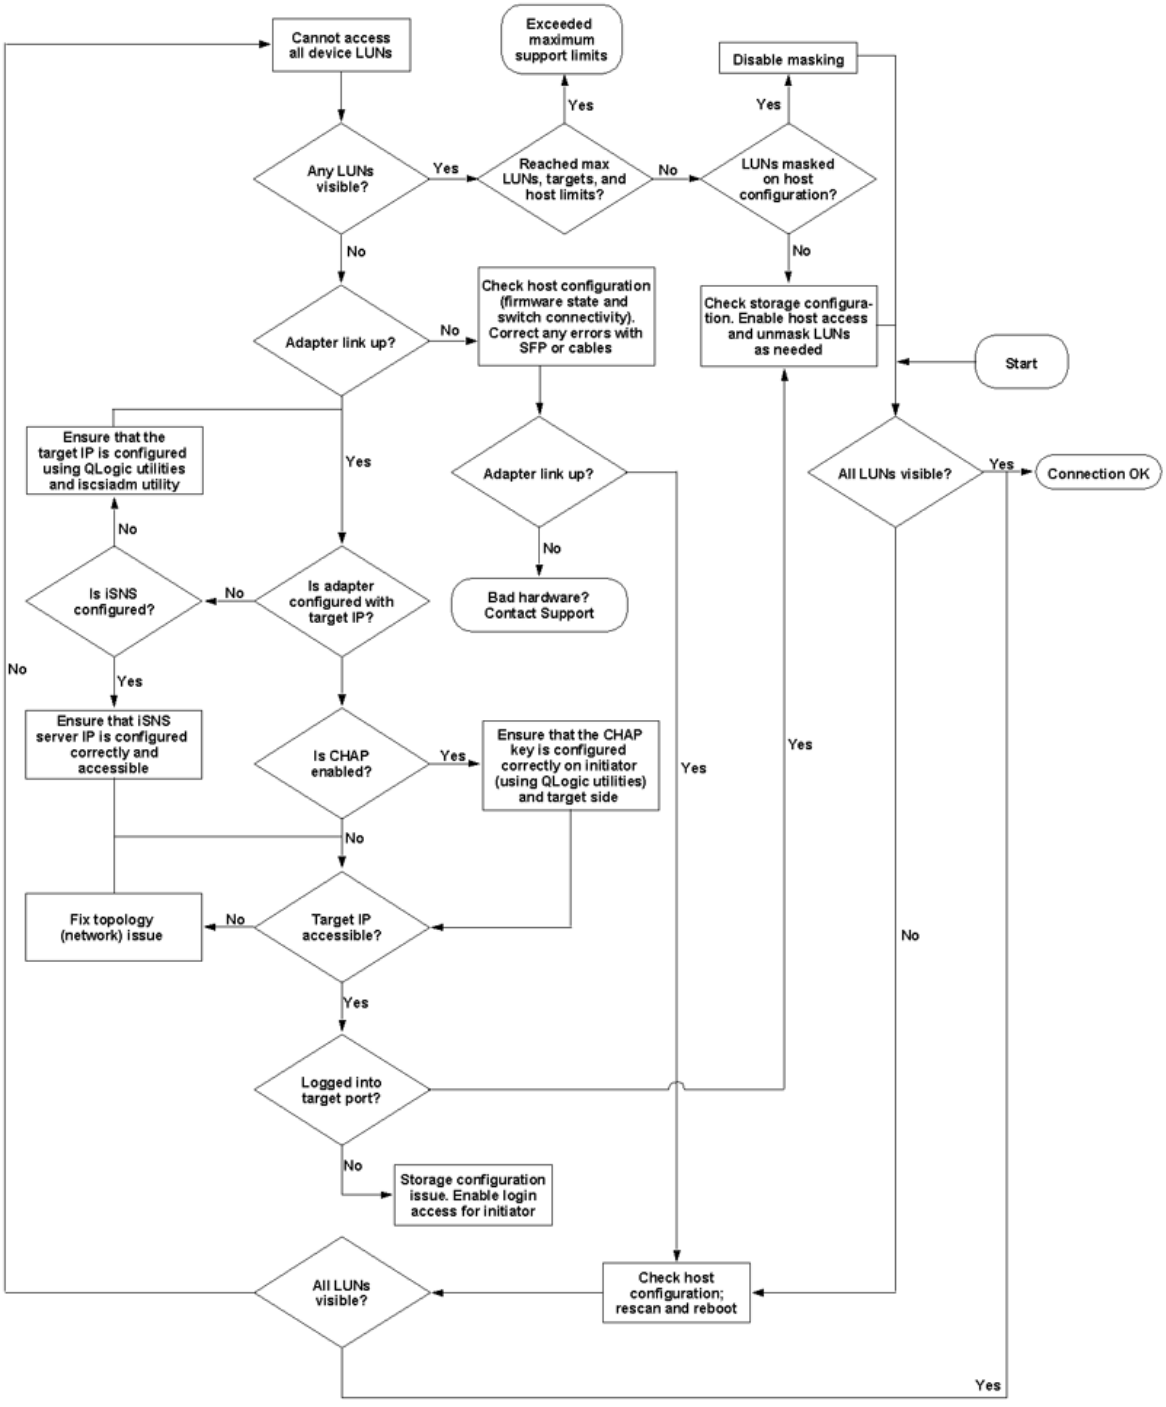

Figure 3-1. iSCSI Troubleshooting Diagram

4 Configuring FCoE

This section describes the driver and agent installation, configuration, operation and troubleshooting of the FCoE function of the 8200 Series Converged Network Adapters.

Installing FCoE

Installing FCoE in Linux

Refer to the *Fibre Channel Adapter and Converged Network Adapter Inbox Driver Update for Linux Kernel 2.6.x* Readme file for information.

Packaging Content

Refer to the *Fibre Channel Adapter and Converged Network Adapter Inbox Driver Update for Linux Kernel 2.6.x* Readme file for information.

Linux OS Support

Refer to the *Fibre Channel Adapter and Converged Network Adapter Inbox Driver Update for Linux Kernel 2.6.x* Readme file for information.

Supported Features

Refer to the *Fibre Channel Adapter and Converged Network Adapter Inbox Driver Update for Linux Kernel 2.6.x* Readme file for information.

Deploying the Driver

Refer to the *Fibre Channel Adapter and Converged Network Adapter Inbox Driver Update for Linux Kernel 2.6.x* Readme file for information.

Driver System Parameters

Refer to the *Fibre Channel Adapter and Converged Network Adapter Inbox Driver Update for Linux Kernel 2.6.x* Readme file for information.

Boot From SAN

Refer to the *Fibre Channel Adapter and Converged Network Adapter Inbox Driver Update for Linux Kernel 2.6.x* Readme file for information.

Firmware Installation

Refer to the *Fibre Channel Adapter and Converged Network Adapter Inbox Driver Update for Linux Kernel 2.6.x* Readme file for information.

Installing FCoE in ESX

Refer to the *Fibre Channel and Fibre Channel over Ethernet (FCoE) Driver for VMware ESX/ESXi 4.0.x and 4.1* Readme file for information.

Packaging Content

Refer to the *Fibre Channel and Fibre Channel over Ethernet (FCoE) Driver for VMware ESX/ESXi 4.0.x and 4.1* Readme file for information.

ESX OS Support

Refer to the *Fibre Channel and Fibre Channel over Ethernet (FCoE) Driver for VMware ESX/ESXi 4.0.x and 4.1* Readme file for information.

Using the Driver

Refer to the *Fibre Channel and Fibre Channel over Ethernet (FCoE) Driver for VMware ESX/ESXi 4.0.x and 4.1* Readme file for information on how to create the driver disk, install the driver, and remove the driver.

Installing FCoE in Windows

Refer to the *Converged Network Adapter FCoE STOR Miniport Driver for Windows* Readme file for information.

Packaging Content

Refer to the *Converged Network Adapter FCoE STOR Miniport Driver for Windows* Readme file for information.

Windows OS Support

Refer to the *Converged Network Adapter FCoE STOR Miniport Driver for Windows* Readme file for information.

Using the Driver

Refer to the *Converged Network Adapter FCoE STOR Miniport Driver for Windows* Readme file for information on how to create the driver disk, install the driver, and remove the driver.

Driver Parameters

Refer to the *Converged Network Adapter FCoE STOR Miniport Driver for Windows* Readme file for information about Storport and driver registry parameters.

HBA Parameters

Setting HBA Parameters with the QCC GUI

Refer to the *QConvergeConsole Help System: Managing Fibre Channel and FCoE Adapters and Ports* for information.

Setting HBA Parameters with the Interactive QCC CLI

Refer to the *QConvergeConsole User's Guide* for information.

Setting HBA Parameters with the Non-Interactive QCC CLI

Refer to the *QConvergeConsole CLI User's Guide,* "Fibre Channel Noninteractive Commands" section for information.

Target Persistent Binding

Configuring Persistent Binding with the QCC GUI

Refer to *QConvergeConsole Help System* for information.

Configuring Persistent Binding with the Interactive QCC CLI

Refer to the *QConvergeConsole CLI User's Guide* for information.

Configuring Persistent Binding with the Non-Interactive QCC CLI

Refer to *QConvergeConsole CLI User's Guide,* "Fibre Channel Noninteractive Commands" section for information.

Boot Devices Configuration

Configuring Boot Devices with the QCC GUI

Refer to *QConvergeConsole Help System: HBA Parameters-Viewing or Modifying Port Parameters - Boot Device Selection* for information.

Configuring Boot Devices with the Interactive QCC CLI

Refer to the *QConvergeConsole CLI User's Guide* for information.

Configuring Boot Devices with the Non-Interactive QCC CLI

Refer to the *QConvergeConsole CLI User's Guide*, "Fibre Channel Noninteractive Commands" section for information.

Configuring Boot Devices with the BIOS

Refer to Using QLflash for information.

Virtual Ports (NPIV)

Configuring NPIV with the QCC GUI

Refer to the *QConvergeConsole Help System: Using Virtualization (NPIV)* for information.

Configuring NPIV with the Interactive QCC CLI

Refer to the *QConvergeConsole CLI User's Guide* for information.

Configuring NPIV with the Non-Interactive QCC CLI

Refer to the *QConvergeConsole CLI User's Guide,* “Fibre Channel Noninteractive Commands” section for information

Driver Parameters

Configuring FCoE Driver Parameters with the QCC GUI

Refer to the *QConvergeConsole Help System: HBA Parameters-Viewing or Modifying Port Parameters* for information.

Configuring FCoE Driver Parameters with the Interactive QCC CLI

Refer to the *QConvergeConsole CLI User's Guide* for information.

Configuring FCoE Driver Parameters with the Non-Interactive QCC CLI

Refer to the *QConvergeConsole CLI User's Guide,* “Fibre Channel Noninteractive Commands” section for information.

Configuring FCoE Driver Parameters with the vCenter Plugin

Refer to the Configuring NPAR in the QLogic Adapter vCenter Plug-in for more information.

Selective LUNS

Configuring Selective LUNS with the QCC GUI

Refer to the *QConvergeConsole Help System* for information.

Configuring Selective LUNS with the Interactive QCC CLI

Refer to the *QConvergeConsole CLI User's Guide* for information.

Configuring Selective LUNS with the Non-Interactive QCC CLI

Refer to the *QConvergeConsole CLI User's Guide*, “Fibre Channel Noninteractive Commands” section for information.

Troubleshooting

FCoE Diagnostics

FCoE Diagnostics using QCC GUI

Refer to the *QConvergeConsole Help System: Performing Diagnostics* section for information.

FCoE Diagnostics using Interactive QCC CLI

Refer to the *QConvergeConsole CLI User's Guide* for information.

FCoE Troubleshooting Diagram

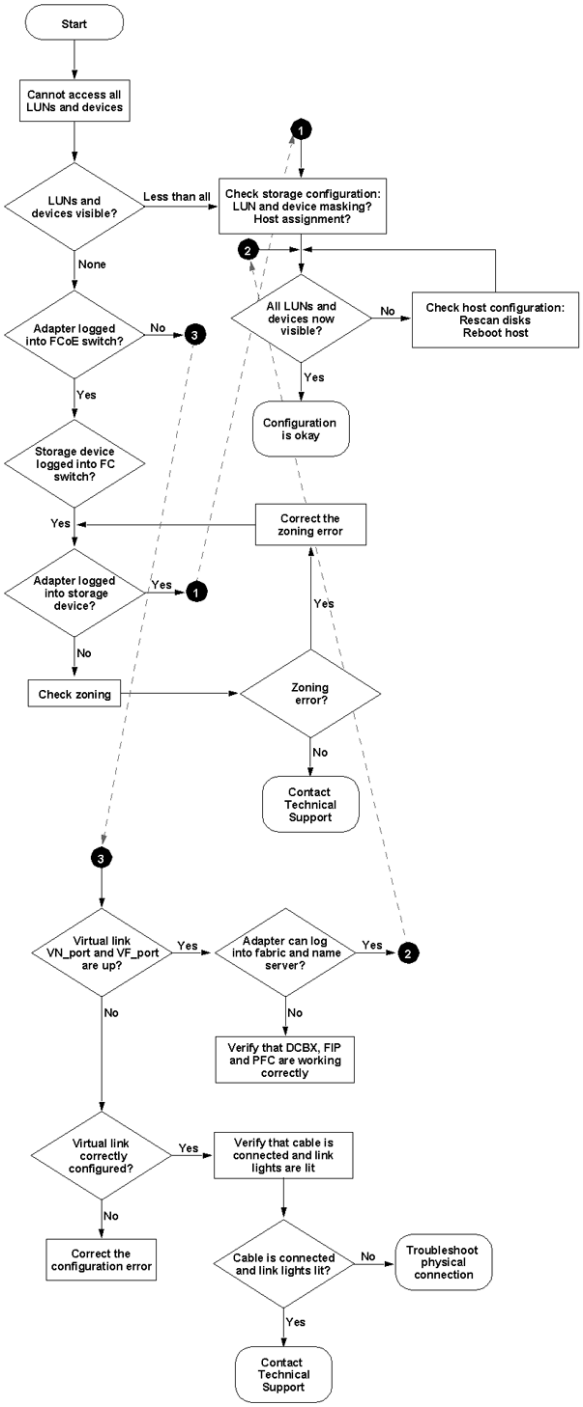

Figure 4-1. FCoE Troubleshooting Diagram

A Adapter LEDs

The LED behavior for the 8200 and 3200 Series Adapters is defined in Table A-1.

Table A-1. QLE824x and QLE324x LED Scheme

Link LED	Activity LED	Activity
Off	Off	Power off
Off	Off	Power on (no link)
On	On	Power on/link established, no activity
On	Flashing	Power on/link established, LAN and/or SAN activity

Glossary

adapter

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

adapter port

A port on the adapter board.

adapter port beacon

An LED on the adapter. Flashing it enables you to locate the adapter.

alias

A user-defined name for an adapter, adapter port, logical disk, or subsystem.

Basic input output system (BIOS)

The program (or utility) that serves as an interface between the hardware and the operating system and allows booting from the adapter at startup.

boot device

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

boot from SAN

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

challenge-handshake authentication protocol (CHAP)

CHAP is used for remote logon, usually between a client and server or a Web browser and Web server. A challenge/response is a security mechanism for verifying the identity of a person or process without revealing a secret password that is shared by the two entities. Also referred to as a" three-way handshake."

Converged Network Adapter

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

cyclic redundancy check (CRC)

A type of check value designed to catch most transmission errors.

device

A computer subsystem, such an adapter card, which mediates data in a computer network. The term "device" is used interchangeably with "target" and "target device".

driver

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

dynamic host configuration protocol (DCHP)

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

Enhanced Ethernet

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

Ethernet

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

failover path

Software feature that ensures data availability and system reliability by assigning alternate path and automatic adapter failover for device resources.

***Fast!*UTIL**

QLogic *Fast!*UTIL™ Fibre Channel Adapter BIOS utility.

Fibre Channel

High-speed serial interface technology that supports other higher layer protocols such as SCSI and IP, and is primarily used in SANs. Standardized under ANSI in 1994.

Fibre Channel over Ethernet (FCoE)

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

Flash

Nonvolatile memory where the boot code is saved. At times, Flash and boot code are used interchangeably.

Internet protocol (IP)

A method by which data is sent from one computer to another over the Internet. IP specifies the format of packets, also called datagrams, and the addressing scheme.

Internet small computer system interface (iSCSI)

Protocol that encapsulates data into IP packets to send over Ethernet connections. An alternative to FCIP.

load balancing

A software feature that improves system performance by balancing device access between multiple ports or adapters for maximum resource efficiency.

local area network (LAN)

Network typically with transmissions less than 5km.

logical unit number (LUN)

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

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

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

loopback

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

message signaled interrupts (MSI/MSIX)

One of two PCI-defined extensions to support message signaled interrupts (MSI), in PCI 2.2 and later and PCIe. MSIs are an alternative way of generating an interrupt through special messages that allow emulation of a pin assertion or deassertion.

MSI-X (defined in PCI 3.0) allows a device to allocate any number of interrupts between 1 and 2048 and gives each interrupt separate data and address registers. Optional features in MSI (64-bit addressing and interrupt masking) are mandatory with MSI-X.

network interface controller/card (NIC)

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

node port (N_Port)

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

node port ID virtualization (NPIV)

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

Nonvolatile random access memory (NVRAM)

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

path

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

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

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

PCIe (PCI Express)

A third-generation input/output (I/O) standard that allows enhanced Ethernet network performance beyond that of the older peripheral component interconnect (PCI) and PCI extended (PCI-x) desktop and server slots.

port

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

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

port instance

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

quality of service (QoS)

Methods used to prevent bottlenecks and ensure business continuity when transmitting data over virtual ports by setting priorities and allocating bandwidth.

redundant array of independent/inexpensive disks (RAID)

Fault-tolerant disks that look like either single or multiple volumes to the server.

small computer systems interface (SCSI)

The original SCSI specification was a hardware bus specification and a packet-oriented protocol specification for communicating on that bus. SCSI over Fibre Channel uses the packet-oriented protocol to communicate with storage devices on the Fibre Channel.

storage area network (SAN)

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

target

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

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

transmission control protocol/Internet protocol (TCP/IP)

Basic communication language of the Internet.

UEFI

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

worldwide port name (WWPN)

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

Corporate Headquarters QLogic Corporation 26650 Aliso Viejo Parkway Aliso Viejo, CA 92656 949.389.6000 **www.qlogic.com**

International Offices UK | Ireland | Germany | France | India | Japan | China | Hong Kong | Singapore | Taiwan