



Marvell® QLogic® Fibre Channel Adapters

2800 Series

User's Guide



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

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

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

Copyright 2026 Marvell or its affiliates. All rights reserved. Marvell and the Marvell logo are trademarks of Marvell or its affiliates. Please visit www.marvell.com for a complete list of Marvell trademarks. All third-party names, certification marks, trademarks (including logos and product names) and brands are the property of their respective owners. Use of such marks is for identification purposes only and does not imply endorsement or affiliation.

Table of Contents

Preface

What Is in This Guide	x
Intended Audience	xi
Related Materials	xii
Documentation Conventions	xii
Technical Support.	xv
Downloading Updates and Documentation	xv
Legal Notices	xvi
Laser Safety—FDA Notice.	xvi
Agency Certification.	xvi
EMI and EMC Requirements	xvi
Product Safety Compliance.	xviii

Quick Start

Installation Instructions.	xix
Step 1. Install the Adapter Hardware.	xix
Step 2. Install the Adapter Drivers.	xx
Step 3. Install an Adapter Management Tool.	xxi
Additional Resources	xxii

1 Product Overview

Functional Description	1
Features	1
Supported Operating Systems	3
Physical Characteristics	3
Standards Specifications	3

2 Software Installation and Initial Setup

Fibre Channel Adapter Installation with QLogic SuperInstallers	5
Multiboot Image for Fibre Channel Adapter Function Configuration Utility. .	5
Multiboot Package Contents	5
Fibre Channel Adapter Function Configuration Utility Package Contents.	5
Updating Multiboot Code and Firmware using QCC CLI.	6

Updating Multiboot Code and Firmware using UEFI	7
On-The-Fly Firmware Updates	7
Fibre Channel Adapter Function Configuration Utility	8
Using UEFI for Custom Configuration	8

3 Fibre Channel Configuration

Installing Fibre Channel	10
Installing Fibre Channel in Linux	10
Linux OS Support	10
Supported Features	10
Deploying the Driver	10
Driver System Parameters	10
Importing a Public Key for Secure Boot	11
Boot from SAN	13
Installing Fibre Channel Boot from SAN in Citrix Hypervisor Environments	13
Installing Fibre Channel in VMware ESXi	15
ESXi OS Support	15
Using the Driver	15
Installing Fibre Channel in Windows	16
Windows OS Support	16
Using the Driver	16
Setting Fibre Channel Adapter Parameters	16
Setting Fibre Channel Adapter Parameters with Interactive QConvergeConsole CLI	16
Setting Fibre Channel Adapter Parameters with Noninteractive QConvergeConsole CLI	17
Configuring Target Persistent Binding	17
Configuring Persistent Binding with Interactive QConvergeConsole CLI	17
Configuring Persistent Binding with Noninteractive QConvergeConsole CLI	17
Configuring Boot Devices.	18
Configuring Boot Devices with Interactive QConvergeConsole CLI . .	18
Configuring Boot Devices with Noninteractive QConvergeConsole CLI	18
Configuring Virtual Ports (NPIV).	19
Configuring NPIV with Interactive QConvergeConsole CLI.	19
Configuring NPIV with Noninteractive QConvergeConsole CLI	20
NPIV Quality of Service.	21
Setting QoS by Priority	22

Setting QoS by Bandwidth	23
Configuring Driver Parameters.	24
Configuring Fibre Channel Driver Parameters with Interactive QConvergeConsole CLI	24
Configuring Fibre Channel Driver Parameters with Noninteractive QConvergeConsole CLI	24
Configuring Selective LUNs	25
Setting Queue Depth	28
Enabling Fabric-Assigned World Wide Port Name.	28
Enabling FA-WWPN with Interactive QConvergeConsole CLI	28
Enabling FA-WWPN with Noninteractive QConvergeConsole CLI . . .	29
Using Fabric-Assigned Boot LUN.	29
Using a Fabric-Assigned Boot LUN with Interactive QConvergeConsole CLI	29
Using a Fabric-Assigned Boot LUN with Noninteractive QConvergeConsole CLI	30
FDMI Enhancements Support	30
Configuring End-to-End (CS_CTL) QoS	33
Configuring CS_CTL QoS on the Switch.	33
Configuring CS_CTL QoS on the Storage Device.	34
Configuring QoS Priority Levels	34
Changing the Priority Levels on Windows	34
Enabling Forward Error Correction.	35
FEC Process Overview	36
Enabling and Disabling FEC	36
Enabling and Disabling FEC on VMware ESXi.	36
Enabling and Disabling FEC on Linux and Windows	37
Using Buffer-to-Buffer Credit Recovery	37
Resetting Driver and Firmware Counters.	38
Configuring VMID.	38
Universal SAN Congestion Mitigation (USCM)	39
Prerequisites	40
SAN Congestion Mitigation	40
USCM I/O Throttling Profiles	41
USCM Virtual Lanes	42
FPIN Link Integrity Marginal Path Detection Support	43
Multipath Configuration File Changes for RHEL 9 and Later; SLES15 SP7, SLES16	43
Multipath Configuration File Changes for ESXi 8.0 U3 and Later	43
Using Fibre Channel Diagnostics.	43

	Device Ping Testing	44
	Fibre Channel Ping Testing (Echo ELS)	44
	CT Ping Testing	44
	CT Fibre Channel Trace Route Testing	45
	Link Cable Beacon Extended Link Service Command	45
	Diagnostic Port Testing	46
	Read Diagnostic Parameters Extended Link Service Command	46
	Running RDP Test from a Brocade 16/32/64G Switch	46
	Running RDP Test from a Cisco 16/32G Switch.	48
	Management Component Transport Protocol Support.	49
	MCTP Support for QLE2874 Adapters and SFP Polling	50
4	NVM Express over Fibre Channel	
	Hardware and Software Requirements	52
	Adapters	53
	Operating Systems	53
	Host Setup	53
	Installing the Drivers	53
	Finding Host NQN	54
	Linux	54
	Windows	55
	ESXi	55
	Programming the Host NQN onto the NVMe Storage Array	55
	Registering the Host NQN into the Fibre Channel Fabric	55
	Configuring System for NVMe Auto Discovery	55
	Verifying the Availability of FC-NVMe Devices	57
	Verifying in Linux	57
	Verifying in Windows	57
	Verifying in ESXi	59
	Unloading and Reloading the Linux Driver	60
	Limitations and Known Issues.	60
	Boot From SAN Setup	61
	Setup for NVME Boot-from-SAN Using ESXi 8.0 U3	61
	Setup for NVME Boot-from-SAN Using ESXi 9.0 and Later	64
A	Adapter LEDs	
B	Troubleshooting	
	Troubleshooting with Fibre Channel Diagnostics	67
	Fibre Channel Diagnostics Using Interactive QConvergeConsole CLI	68

Fibre Channel Diagnostics Using Noninteractive QConvergeConsole CLI	68
Fibre Channel Troubleshooting Diagram	69
Troubleshooting in Diagnostics Port Mode	70
Configuring, Running, and Viewing Results of a Diagnostics Port on a Brocade Fibre Channel Switch	73
Configuring, Running, and Viewing Results of a Diagnostics Port on a Cisco Fibre Channel Switch	75

C **Revision History**

Glossary

List of Figures

Figure		Page
	Server Motherboard and Slots	xx
3-1	Unloading the qla2xxx driver	14
3-2	Confirm Driver Update	14
3-3	Loading the OOB driver	15
3-4	Priority- and Bandwidth-based QoS	21
4-1	FC-NVMe Devices in Windows Device Manager	58
4-2	Listed NVMe Controllers	59
4-3	NVMe Namespace List	59
4-4	UUID NVME BIOS SETTING	64
A-1	QLE2800 Adapter LEDs	66
B-1	Fibre Channel Troubleshooting Diagram (Part 1 of 2)	69
B-2	Fibre Channel Troubleshooting Diagram (Part 2 of 2)	70

List of Tables

Table		Page
	Supported Adapter Models	x
	Adapter Management Tools	xxi
A-1	QLE2800, 64Gb Adapter LED Scheme	66

Preface

This guide provides detailed instructions on installing, configuring, and troubleshooting the 2800 Series Adapters listed in [Table](#) below. It also provides details on the use of adapter features to enhance the value of server virtualization. Such features include virtual adapter configuration using N_Port ID virtualization (NPIV) and boot-from-SAN configuration.

Supported Adapter Models

Model Number	Port Quantity	Speed
QLE2870	1	64/32/16 GFC
QLE2872	2	64/32/16 GFC
QLE2874	4	64/32/16 GFC

The 2800 Series Adapters run on the following operating systems:

- Microsoft® Windows® (including Hyper-V)
- Linux®
- VMware® ESXi
- Citrix® XenServer®

NOTE

For a detailed list of the currently supported operating systems, refer to the 2800 Series Adapter's *Read Me* and *Release Notes* at www.marvell.com.

What Is in This Guide

This preface specifies the intended audience, lists related documents, explains the typographic conventions used in this guide, and provides technical support, and regulatory information.

The remainder of this guide is organized into the following chapters and appendices:

- [Quick Start](#) provides high-level hardware and software installation instructions for advanced users, as well as additional resources.
- [Chapter 1 Product Overview](#) provides a functional description of the 2800 Series Adapters and lists the features, supported operating systems, physical characteristics, and standards specifications.
- [Chapter 2 Software Installation and Initial Setup](#) provides information about the SuperInstaller and the multi-boot image (Fibre Channel Adapter Function Configuration Utility).
- [Chapter 3 Fibre Channel Configuration](#) describes installing the Fibre Channel driver and agent across OSs; setting 2800 Series Adapter parameters; persistent binding; configuring the boot device, NPIV, and driver parameters; enabling the fabric to assign a world-wide port name or boot LUN; enabling end-to-end quality of service (QoS), diagnostics, configuring VMID, and Universal SAN Congestion Mitigation.
- [Chapter 4 NVM Express over Fibre Channel](#) describes the installation and configuration of NVMe over Fibre Channel (FC-NVMe).
- [Appendix A Adapter LEDs](#) describes the 2800 Series Adapter LEDs and defines the LED behavior.
- [Appendix B Troubleshooting](#) helps solve issues by running Fibre Channel diagnostics using QConvergeConsole® to identify and isolate link failures. It also provides a flowchart that walks you through troubleshooting problems that may occur with devices and LUNs on the 2800 Series Adapter.
- [Appendix C Revision History](#) contains a list of changes made to this guide since the last revision.

Following the appendices is a glossary of terms and acronyms used.

Intended Audience

This guide is intended for those responsible for deploying 2800 Series Adapters on Windows, Linux, VMware ESXi, VMware vCenter™ Server, and Citrix XenServer: from end users (such as data center managers and system administrators) to the test and development community.

Related Materials

For additional information, refer to following documents that are available from the Marvell Web site, www.marvell.com:

- *2800 Series Adapter's Read Me and Release Notes* documents contain information about installation, OS support, and known issues.
- *User's Guide—UEFI Human Interface Infrastructure, 2690 Series Enhanced 16GFC, 2740/2760 Series 32GFC, 2770 Series Enhanced 32GFC, 2800 Series 64GFC Marvell QLogic Fibre Channel Adapters* (part number BK3254602-00) describes how to use the HII application to configure the adapter parameters and boot-from-SAN settings.
- *User's Guide—QConvergeConsole CLI* (part number SN0054667-00) provides specific command line use in both interactive and noninteractive modes.¹
- *User's Guide—Marvell® QLogic® Fibre Channel QConvergeConsole Extension for Windows® Admin Center* (part number TD-001728) provides information on installing and using the Marvell QLogic Fibre Channel (FC) QConvergeConsole Extension for Windows Admin Center (WAC) to manage the Marvell QLogic 2600/2700/2800 Series Adapters.¹
- *User's Guide—Marvell QLogic PowerKit: Windows PowerShell® cmdlets for 2600, 2700, 2800 Series Marvell QLogic Fibre Channel Adapters* (part number TD-001187) provides information on installing the Marvell QLogic Fibre Channel (MRVLFC) PowerKit (Windows PowerShell cmdlets) software to manage Marvell QLogic Fibre Channel (FC) adapters.¹
- *User's Guide—Marvell ESXCLI Plug-in for VMware vSphere* (part number BC0151101-00) provides information on installing and using the Marvell ESXCLI VMware® Plug-in. The ESXCLI plug-in extends the capabilities of the ESX® command line interface to manage the following adapters installed in VMware ESX/ESXi hosts.¹

For information about downloading documentation from the Marvell Web site, see [“Downloading Updates and Documentation” on page xv](#).

Documentation Conventions

This guide uses the following documentation conventions:

- The adapters described in this guide are collectively referred to as the 2800 Series Adapters.
- **NOTE** provides additional information.

- **CAUTION** without an alert symbol indicates the presence of a hazard that could cause damage to equipment or loss of data.
- Text in **blue** font indicates a hyperlink (jump) to a figure, table, or section in this guide, and links to Web sites are shown in underlined blue. For example:
 - ❑ **Table 9-2** lists problems related to the user interface and remote agent.
 - ❑ See **"Installation Checklist"** on page 3-6.
 - ❑ For more information, visit www.marvell.com.
- Text in **bold** font indicates user interface elements such as a menu items, buttons, check boxes, or column headings. For example:
 - ❑ Click **Start**, point to **Programs**, point to **Accessories**, and then click **Command Prompt**.
 - ❑ Under **Notification Options**, select the **Warning Alarms** check box.
- Text in **Courier** font indicates a file name, directory path, or screen output. For example:
 - ❑ To return to the root directory from anywhere in the file structure, type `cd /root`, and then press the ENTER key.

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

- ❑ Issue the following command: `# sh ./install.bin`
- Key names and key strokes are indicated with UPPERCASE:
 - ❑ Press 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 applicable *Software End User License Agreement*.
 - ❑ What are *shortcut keys*?
 - ❑ To enter the date type *mm/dd/yyyy* (where *mm* is the month, *dd* is the day, and *yyyy* is the year).
- Topic titles between quotation marks identify related topics either within this manual or in the online help, which is also referred to as *QConvergeConsole Help* throughout this document.

- QConvergeConsole CLI non-interactive command syntax conventions include the following:
 - Plain text indicates items that you must type as shown. For example:
 - `gauccli -pr fc -ei`
 - `< >` (angle brackets) indicate a variable whose value you must specify. For example:
 - `<hba instance>`

NOTE

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

- `[ ]` (square brackets) indicate an optional parameter. For example:
 - `[<file_name>]` means specify a file name, or omit it to select the default file name.
- `|` (vertical bar) indicates mutually exclusive options; select one option only. For example:
 - `on|off`
 - `1|2|3|4`
- `...` (ellipsis) indicates that the preceding item may be repeated. For example:
 - `x...` means *one* or more instances of `x`.
 - `[x...]` means *zero* or more instances of `x`.
- `:` (vertical ellipses) within command example output indicate where portions of repetitious output data have been intentionally omitted.
- `( )` (parentheses) and `{ }` (braces) are used to avoid logical ambiguity. For example:
 - `a|b c` is ambiguous
 - `{ (a|b) c }` means `a` or `b`, followed by `c`
 - `{ a | (b c) }` means either `a`, or `b c`

Technical Support

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

Downloading Updates and Documentation

The Marvell Web site provides periodic updates to product firmware, software, and documentation.

To download firmware, software, and documentation:

1. Go to www.marvell.com.
2. Click **Support**, and then under **Tools & Resources**, click **Driver Downloads**.
3. In the Marvell Drivers window:
 - a. (MUST) Under CATEGORY, select FIBRE CHANNEL ADAPTERS.
 - b. (optional) Under PLATFORM/OS, select the platform/OS that matches your system.
 - c. (optional) Under PART NUMBER, select the part number for your adapter.
 - d. (optional) Under KEYWORDS, type a keyword describing what you are looking for.
4. Click **Apply**.
5. Locate the firmware (boot code), software (drivers, management tools), or document (documentation for user's guides) you need, and then do one of the following:
 - a. Click the [blue](#) text in the DESCRIPTION column.
 - b. Click the arrow in the DOWNLOAD column.

NOTE

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

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

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

Legal Notices

This section provides laser safety (FDA notice), agency certification, and product safety compliance for the following products:

- QLE2870 (part number MA2810401-13, MA2810401-B3)
- QLE2872 (part number MA2810401-14, MA2810401-B4)
- QLE2874 (part number MA2810402-17, MA2810402-57)

Laser Safety—FDA Notice

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

CLASS I LASER

Class 1 Laser Product	Caution —Class 1 laser radiation when open Do not view directly with optical instruments
Appareil laser de classe 1	Attention —Radiation laser de classe 1 Ne pas regarder directement avec des instruments optiques
Produkt der Laser Klasse 1	Vorsicht —Laserstrahlung der Klasse 1 bei geöffneter Abdeckung Direktes Ansehen mit optischen Instrumenten vermeiden
Luokan 1 Laserlaite	Varoitus —Luokan 1 lasersäteilyä, kun laite on auki Älä katso suoraan laitteeseen käyttämällä optisia instrumenttejä

Agency Certification

The following sections contain a summary of EMC and EMI test specifications performed on models MA2810401-13, MA2810401-B3, MA2810401-14, MA2810401-B4, MA2810402-17, and MA2810402-57 to comply with radiated emission, radiated immunity, and product safety standards.

EMI and EMC Requirements

FCC Rules, CFR Title 47, Part 15, Subpart B: 2017 Class A

Compliance information statement: This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

Industry Canada, ICES-003:2016: Class A

This Class A digital apparatus complies with Canadian ICES-003.
Cet appareil numérique de la classe A est conforme à la norme NMB-003 du Canada.

CE Mark 2014/30/EU EMC Directive Compliance

EN55032:2012/CISPR 32:2015 Class A

EN55024:2010

EN61000-3-2:2014

EN61000-3-3:2013

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

AS/NZS; CISPR 32:2015 Class A

CNS 13438:2006 Class A

KCC: Class A

KC-RRA KN32 KN35 (2014 for BK series) Class A



Product Name/Model—MA2810401-13, MA2810401-B3,
MA2810401-14, MA2810401-B4, MA2810402-17, and
MA2810402-57 Adapters

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.

VCCI: CISPR 32: 2015 Class A

This is a Class A product based on the standard of the Voluntary Control Council for Interference (VCCI). If this equipment is used in a domestic environment, radio interference may occur, in which case the user may be required to take corrective actions.

この装置は、クラスA機器です。この装置を住宅環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

VCCI-A

BSMI: Class A

BSMI CNS 13438:2006 Class A for the BK3210407

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。

Product Safety Compliance

UL, cUL Product Safety

UL60950-1 (2nd Edition), A1 + A2 2014-10-14

CSA C22.2 60950-1-07 (2nd Edition) A1 +A2 2014-10

Use only with listed ITE or equivalent.

Complies with 21 CFR 1040.10 and 1040.11.

For use with UL Listed Servers or equivalent.

Power off the computer and all attached devices such as monitors, printers, and external components before installing the accessory card or equivalent.

2014/35/EU Low Voltage Directive

TUV EN60950-1:2006+A11+A1+A12 2nd Edition

TUV IEC 60950-1 :2005 2nd Edition Am 1 :2009 + Am2 :2013 CB

TUV IEC62368 2nd and 3rd Edition CB TUV EN62368-1 :2014+A11:2017

Quick Start

To help you quickly get started using the 2800 Series Adapters, this chapter provides installation steps and lists additional resources.

Installation Instructions

Install and configure your new 2800 Series Adapter in three simple steps:

- [Step 1. Install the Adapter Hardware](#)
- [Step 2. Install the Adapter Drivers](#)
- [Step 3. Install an Adapter Management Tool](#)

Step 1. Install the Adapter Hardware

CAUTION


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

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.

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](#) below).



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 optical cable into the adapter.
10. Plug in the power cable and turn on the computer.

Step 2. Install the Adapter Drivers

Follow these steps to download and install adapter drivers.

To install the adapter drivers:

1. Go to www.marvell.com.
2. Click **Support**, and then under **Tools & Resources**, click **Driver Downloads**.
3. In the Marvell Drivers window:
 - a. (MUST) Under CATEGORY, select FIBRE CHANNEL ADAPTERS.
 - b. (optional) Under PLATFORM/OS, select the platform/OS that matches your system.
 - c. (optional) Under PART NUMBER, select the part number for your adapter.
 - d. (optional) Under KEYWORDS, type **drivers**.

4. Click **Apply**.
5. Locate the drivers you need, and then do one of the following:
 - a. Click the [blue](#) text in the DESCRIPTION column.
 - b. Click the arrow in the DOWNLOAD column.

NOTE

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

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

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

Step 3. Install an Adapter Management Tool

Based on your system, download one of the Marvell tools in to configure and manage your 2800 Series Adapter.

Adapter Management Tools

If you are using ...	Install
Windows, Linux, Citrix XenServer	QConvergeConsole CLI
Windows, VMware ESXi, Linux	Marvell QLogic PowerKit
Windows	Marvell QLogic Fibre Channel QConvergeConsole Extension for Windows Admin Center

To download an application, follow the instructions in [“Downloading Updates and Documentation” on page xv](#).

For detailed information about using the adapter management tools, see the documents specified in [“Related Materials” on page xii](#).

Additional Resources

For more information on using the 2800 Series Adapter, Marvell provides these additional resources:

- To obtain the most current drivers, management tools, user instructions, and documentation, download the latest versions from the Downloads and Documentation page at www.marvell.com.
- For descriptions of the user instructions provided in this document, see [“What Is in This Guide” on page x](#).
- For important product information, including warranty, laser safety, and agency certification, see the [“Legal Notices” on page xvi](#).

1 Product Overview

This chapter describes the function and features of the 2800 Series Adapter including:

- [Functional Description](#)
- [Features](#)
- [“Supported Operating Systems” on page 3](#)
- [“Physical Characteristics” on page 3](#)
- [“Standards Specifications” on page 3](#)

Functional Description

The 2800 Series Adapters are Marvell Gen 7 Fibre Channel Adapters. They boast industry-leading native Fibre Channel (FC) performance—achieving quad-port (QLE2874), dual-port (QLE2872), or single-port (QLE2870) line-rate, 64GFC throughput—at extremely low CPU use with full hardware offloads.

Features

The 2800 Series Adapters provide the following features. For more information about these features, see the QConvergeConsole CLI online help.

- QLE2870, QLE2872, QLE2874 offer connectivity to 64Gb, 32Gb, and 16Gb Fibre Channel networks.
- Physical PCIe characteristics (with backward compatibility):
 - ❑ PCIeGen4 x 8 (QLE2870, QLE2872)
 - ❑ PCIeGen4 x16 (QLE2874)
- Device management for SAN
- Full hardware offload for the Fibre Channel protocol
- Message-signaled interrupts (MSI-X) and legacy interrupts (INT-X)
- N_Port ID virtualization (NPIV)
- Buffer-to-buffer credit recovery (BB-CR)

- Firmware and driver counter reset
- Boot from SAN
- Advanced management features for the 2800 Series Adapter include QConvergeConsole CLI management utility. This utility is available if you are running Windows, Linux, or Citrix XenServer.
- Diagnostics port
- FC-NVMe
- Host NVMe Qualified Name (NQN) registration (FOS version 9.2.2a and later, for 32G and 64G fabrics only)
- Universal SAN Congestion Mitigation (USCM)
 - FPIN Monitoring
 - I/O Throttling Profiles
 - Virtual Lanes
 - Marginal Paths (FPIN-LI/MPIO)
- Management Component Transport Protocol (MCTP)
- Secure Firmware Update
- Link Cable Beaconsing (LCB)
- FDMI Registrations
- CS_CTL-based QoS
- Device ping
- Fibre Channel ping (Echo ELS)
- Common transport (CT) ping
- CT Fibre Channel trace route (FTR)
- Loop back testing
- Read buffer/write buffer testing
- Send read diagnostic parameters (RDP)
- 16G Forward error correction (FEC) (enable/disable, counter display, and reset)
- Principal/adjacent fabric WWN
- Fabric-assigned WWPN (FA-WWPN) (FOS 9.2.0 or earlier)
- Fabric-assigned boot LUN (FA-BLD) (FOS 9.2.0 or earlier)
- VMID support on VMware ESXi 8.0 U3 and 9.x
- Support for long range 10km/5km optics

Supported Operating Systems

The 2800 Series Adapters support commonly used operating systems, including Windows Server, Red Hat® Enterprise Linux (RHEL®), SUSE® Linux Enterprise Server (SLES), VMware vSphere ESXi, openEuler®, Kylin®, Unity®, Citrix®, Oracle Linux Unbreakable Enterprise Kernel (UEK), and Ubuntu®.

For a detailed list of the currently supported operating systems, refer to the adapter's *Read Me* and *Release Notes* available at www.marvell.com (see "Downloading Updates and Documentation" on page xv).

Physical Characteristics

Physical characteristics for the 2800 Series Adapters include the following:

- The QLE2870, QLE2872 Adapters are implemented as low-profile PCIe cards. The adapter ships with a standard-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.
- The QLE2874 Adapter is implemented as standard-height PCIe card with a standard-height bracket installed.

Standards Specifications

The 2800 Series Adapter supports the following standards specifications:

- SCSI Fibre Channel Protocol 4 (FCP-4)
- Fibre Channel Tape (FC-TAPE) Profile
- Fibre Channel Generic Services 8 (FC-GS-8)
- Fibre Channel-Physical Interface-5 (FC-PI-5)
- Fibre Channel-Physical Interface-6 (FC-PI-6)
- Fibre Channel Link Services 4 (FC-LS-4)
- Fibre Channel Framing and Signaling 5 (FC-FS-5)
- Fibre Channel NVMe (FC-NVMe)

For more information on these specifications, refer to the Technical Committee T11 Web site:

<http://www.t11.org>

2

Software Installation and Initial Setup

This chapter provides the following information about installing and setting up the software:

- [Fibre Channel Adapter Installation with QLogic SuperInstallers](#)
- [“Multiboot Image for Fibre Channel Adapter Function Configuration Utility” on page 5](#)

NOTE

If you are using VMware ESXi, refer to—*User’s Guide Marvell® Fibre Channel Adapters for VMware® ESXi 8.0 U3 and 9.0* (part number SN0154529-00) for detailed instructions on installation and initial setup.

If you are using Marvell ESXCLI VMware® Plug-in, refer to—*User’s Guide—Marvell® ESXCLI Plug-in for VMware vSphere®* (part number BC0151101) for detailed instructions on installation and usage of Marvell ESXCLI VMware® Plug-in.

If you are using Marvell QLogic Fibre Channel (MRVLFC) PowerKit (Windows PowerShell cmdlets) software, refer to —*User’s Guide—Marvell® QLogic® PowerKit for VMware ESXi 8.0 U3 or 9.0* (part number TD-001187) for information on installation of the Marvell QLogic Fibre Channel (MRVLFC) PowerKit (Windows PowerShell cmdlets) software to manage Marvell QLogic Fibre Channel (FC) adapters.

Fibre Channel Adapter Installation with QLogic SuperInstallers

The QLogic SuperInstallers are easy-to-use tools that install the following components:

- Fibre Channel
- QConvergeConsole CLI management tool

For installation instructions for the QLogic Windows SuperInstaller, refer to the QLogic Windows SuperInstaller *Read Me* document.

Multiboot Image for Fibre Channel Adapter Function Configuration Utility

This section contains the following information:

- [Multiboot Package Contents](#)
- [Fibre Channel Adapter Function Configuration Utility Package Contents](#)
- [Updating Multiboot Code and Firmware using QCC CLI](#)
- [On-The-Fly Firmware Updates](#)
- [Fibre Channel Adapter Function Configuration Utility](#)

Multiboot Package Contents

The multiboot package for the 2800 Series Adapter is a compressed file that contains the Unified Extensible Firmware Interface (UEFI) code and firmware.

Fibre Channel Adapter Function Configuration Utility Package Contents

The following files are included in the boot code package:

- `update.nsh`—UEFI batch file run from a UEFI shell that calls the executable files to update the adapter multiboot
- `ql28xx.drv`—UEFI driver
- `EFlashX64.efi`—UEFI utility
- `Readme`—Boot Code folders contain Readme and Release notes
- `Read1st`—Component versions
- `mhxxxxxx.bin`—Combined binary file, which includes the binaries for the UEFI and firmware

NOTE

The files may be different for each respective OEM software release.

Updating Multiboot Code and Firmware using QCC CLI

To update the multiboot code and firmware using QCC CLI:

1. Boot to the OS with the latest driver and QConvergeConsole CLI installed.
2. Download the MBI kit for your 2800 Series Adapter from Marvell Support Website following the procedures in [“Step 2. Install the Adapter Drivers” on page xx](#), and selecting the relevant package in the Boot Code section.
3. Unzip the following MBI kit to a directory:
`qlogic_ql28xx_img10.07.04_b1_9.12.01` or later
4. Navigate to the directory from step 3.
5. Update the firmware using QCC CLI by doing one of the following:
 - ☐ In non-interactive mode, enter the following command:
`# qauccli -b all mhxxxxxx.bin`
 - ☐ In interactive mode:
 - a. From the **Main** menu, select **Adapter Updates**.
The **FC Adapter Update** menu appears.
 - b. Select **1 Flash Update**.
 - c. From the **Flash Update** menu, select the adapter port, and then type the Flash update file name.
The update confirmation dialog box appears.
 - d. Select **Yes** to continue the update process.

NOTE

The following steps apply when using Linux. Other operating systems may vary or not allow this option.

6. Verify that the new firmware becomes active after the system reboots by doing one of the following using QCC CLI:
 - ☐ In noninteractive mode, enter the following command:
`# qauccli -i`

Confirm that the **Running Firmware Version** matches the **Flash Firmware Version**.

- ❑ In interactive mode:
 - a. From the **Main** menu, select **Adapter Information**.
The **FC Adapter Information** menu appears.
 - b. Select **1 FC Adapter Information**.
 - c. From the **FC Adapter Information** menu, select the adapter for which to view adapter information.
Confirm that the **Running Firmware Version** matches the **Flash Firmware Version**.

Updating Multiboot Code and Firmware using UEFI

You can also update the MBI kit from the UEFI Shell.

1. Unzip MBI kit to a directory, for example:

```
qlogic_ql28xx_img10.07.04_b1_9.12.01
```

The prefix of the MBI zip file indicates whether the file is applicable to either a Marvell QLogic branded adapter or an OEM branded adapter.

2. Boot to the UEFI Shell.
3. Run `update.nsh` from the MBI folder.

On-The-Fly Firmware Updates

The following Marvell management tools perform additional checks to confirm that the new MBI image file to be flashed is consistent with the current running firmware on the adapter.

Marvell Management Tool	Supported Version
QConvergeConsole CLI	version 3.2 and later
Marvell QLogic QConvergeConsole Extension for Windows Admin Center	version 2.5.0 and later
Marvell QLogic PowerKit	version 2.6.1 and later
Marvell ESXCLI Plug-in for VMware vSphere	version 1.9.0 and later

Consistent flash components can be activated and reset on-the-fly.

Components that are consistent with the running firmware:

- Limit changes between the running version and the version to be flashed to Fibre Channel firmware, board configuration, or boot code.

- Have the same firmware attributes in the running firmware and the new firmware.

Components that are not consistent with the running firmware:

- Have changes to components in the flash that need a PCIe reset (PE) or a power-on reset to start (for example, firmware other than Fibre Channel)

If the new MBI image file is consistent with the currently running flash components, the application issues a Fibre Channel firmware reset at the end of the flash update process to activate the new image.

If the new MBI image file is not consistent with the current running flash components, the application does not issue a Fibre Channel firmware reset at the end of the flash update process; instead, the user is prompted to reboot the system.

Fibre Channel Adapter Function Configuration Utility

This section provides detailed configuration information for advanced users who want to customize the configuration of the 2800 Series Adapter and the connected devices. You can configure the 2800 Series Adapter using UEFI:

- Use UEFI for Cisco®, Dell®, HP®, IBM®, Lenovo®, and Oracle x64 OEM adapters where the platform supports UEFI (see [Using UEFI for Custom Configuration](#)).

Using UEFI for Custom Configuration

For details on how to use UEFI to perform custom configuration, refer to the relevant server vendor documentation. For additional information, see the latest UEFI files under **Boot Code** on the Downloads and Documentation Web pages for 2800 Series Adapter models. See also the *User's Guide—UEFI Human Interface Infrastructure*, 2690 Series Enhanced 16GFC, 2740/2760 32GFC, 2770 Series Enhanced 32GFC, 2800 Series 64GFC Marvell QLogic Fibre Channel Adapters (part number BK3254602-00), which describes how to use the HII application to configure adapter parameters and boot-from-SAN settings. For download assistance, see [“Downloading Updates and Documentation” on page xv](#).

3 Fibre Channel Configuration

This chapter describes the driver and agent installation, configuration, operation, and troubleshooting of the Fibre Channel function of the 2800 Series Adapters, including:

- [“Installing Fibre Channel” on page 10](#)
- [“Setting Fibre Channel Adapter Parameters” on page 16](#)
- [“Configuring Target Persistent Binding” on page 17](#)
- [“Configuring Boot Devices” on page 18](#)
- [“Configuring Virtual Ports \(NPIV\)” on page 19](#)
- [“Configuring Driver Parameters” on page 24](#)
- [“Configuring Selective LUNs” on page 25](#)
- [“Setting Queue Depth” on page 28](#)
- [“Enabling Fabric-Assigned World Wide Port Name” on page 28](#)
- [“Using Fabric-Assigned Boot LUN” on page 29](#)
- [“FDMI Enhancements Support” on page 30](#)
- [“Configuring End-to-End \(CS_CTL\) QoS” on page 33](#)
- [“Enabling Forward Error Correction” on page 35](#)
- [“Using Buffer-to-Buffer Credit Recovery” on page 37](#)
- [“Resetting Driver and Firmware Counters” on page 38](#)
- [“Configuring VMID” on page 38](#)
- [“Universal SAN Congestion Mitigation \(USCM\)” on page 39](#)
- [“Using Fibre Channel Diagnostics” on page 43](#)
- [“Management Component Transport Protocol Support” on page 49](#)

Installing Fibre Channel

This section provides information on installing Fibre Channel on the following platforms:

- [Installing Fibre Channel in Linux](#)
- [Installing Fibre Channel Boot from SAN in Citrix Hypervisor Environments](#)
- [Installing Fibre Channel in VMware ESXi](#)
- [Installing Fibre Channel in Windows](#)

Installing Fibre Channel in Linux

For more information, download the latest driver package and the associated *Read Me* and *Release Notes* documents, available at www.marvell.com (see “[Downloading Updates and Documentation](#)” on page xv).

Linux OS Support

To obtain the latest information about supported Linux operating systems, see the latest *Read Me* and *Release Notes*.

Supported Features

The Fibre Channel (inbox and/or out-of-box) Linux driver supports the following:

- Point-to-point, and fabric topologies
- Initiator mode only
- Fibre Channel tape
- MSI-X
- NPIV
- Device mapper
- Support for block SG (BSG) interface
- Support for sysfs interface
- Support for an adapter API

Deploying the Driver

For information on deploying the driver, see the latest *Read Me*.

Driver System Parameters

For information on setting driver system parameters, see the latest *Read Me*.

Importing a Public Key for Secure Boot

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

This section is for use with one of the following packages:

■ SLES:

```
qlgc-qla2xxx-kmp-default-<driver version>_<kernel  
version>.<OSV-version>.<architecture-type>.rpm
```

For example:

```
qlgc-qla2xxx-kmp-default-10.02.08.01_k_k5.14.21_150400.22-1.  
sles15sp4.x86_64.rpm
```

■ RHEL®:

```
kmod-qlgc-qla2xxx-<driver-version>.<OSV-version>.<architectu  
re-type>.rpm
```

For example:

```
kmod-qlgc-qla2xxx-10.02.08.01_k-1.rhel9u0.x86_64.rpm
```

Here, OSV means Operating System Vendor.

NOTE

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

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

Method 1: To import and enroll the Marvell QLogic public key:

1. Download the `qlgc_module-202x.der` or `qlgc_module-202x.zip` public key file from the following web page:

<https://www.marvell.com/support/ldriver.html>

2. To install the `qlgc_module-202x.der` public key, issue the following command:

```
# mokutil --root-pw --import qlgc_module-202x.der
```

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

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

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

4. Reboot the system.

5. When the shim launches MokManager, enter the root password to confirm the certificate importation to the Machine Owner Key (MOK) list.
6. To determine if the newly imported key was enrolled:

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

Method 2: To launch MOK manually and enroll the Marvell QLogic public key:

1. Download the `qlgc_module-202x.der` or `qlgc_module-202x.zip` public key file from the following web page:
<https://www.marvell.com/support/ldriver.html>
2. Issue the following command:

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

```
chainloader $efibootdir/MokManager.efi  
boot
```
5. Select **Enroll key from disk**.
6. Navigate to the `qlgc_module-202x.der` file and then press ENTER.
7. Follow the instructions to enroll the key. Generally, this includes pressing the 0 (zero) key and then pressing the Y key to confirm.

NOTE

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

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

<https://documentation.suse.com/sles/15-SP4/html/SLES-all/cha-uefi.html#sec-uefi-secboot>

Boot from SAN

Booting from SAN means booting the OS from a Fibre Channel target device. You should use the QLogic inbox driver to install the OS to the Fibre Channel target device that is attached to a 2800 Series Adapter. If there is no QLogic inbox driver that supports the 2800 Series Adapter, you should use a DD-Kit¹ to boot from SAN.

NOTE

The following procedure is only for RHEL 9.0 and later distributions.

To boot from SAN when installing to a LUN other than LUN 0 and LUN 0 is present:

1. On the Boot Loader Configuration window, select **Configure Advance Boot Loader**, and then click **Next**.
2. On the Advanced Boot Loader Configuration window, select **Change Drive Order**.
3. On the Edit Drive Order window, move the entry for the boot LUN.
4. To continue, click **OK**.
5. Continue with the OS installation as usual.

Installing Fibre Channel Boot from SAN in Citrix Hypervisor Environments

NOTE

The following procedure is only for Citrix Hypervisor 8.0 and later.

To Boot from SAN when installing to a LUN other than LUN ID 0, or LUN 0 follow the procedure below:

1. Insert the Citrix Hypervisor 8.0 installation CD and begin the installation in a root shell.
2. When the system is booted from XS8.0 install CD, choose the **Shell Mode Installation** option. This selection logs in the user with the required privileges.

¹ Marvell provides DD-Kits (driver update disks) for Linux platforms.

3. Once the system boots to shell, unload the inbox qla2xxx driver.

```
root@localhost /l# modprobe -rv qla2xxx  
rmmod qla2xxx  
rmmod scsi_transport_fc  
rmmod nune_fc  
rmmod nune_fabrics  
rmmod nune_core
```

Figure 3-1. Unloading the qla2xxx driver

4. Type “exit” to return to the GUI installer.
5. Insert the driver update disk (DUD) CD (ISO image) and load the driver using the “F9” option as shown in the following figure.

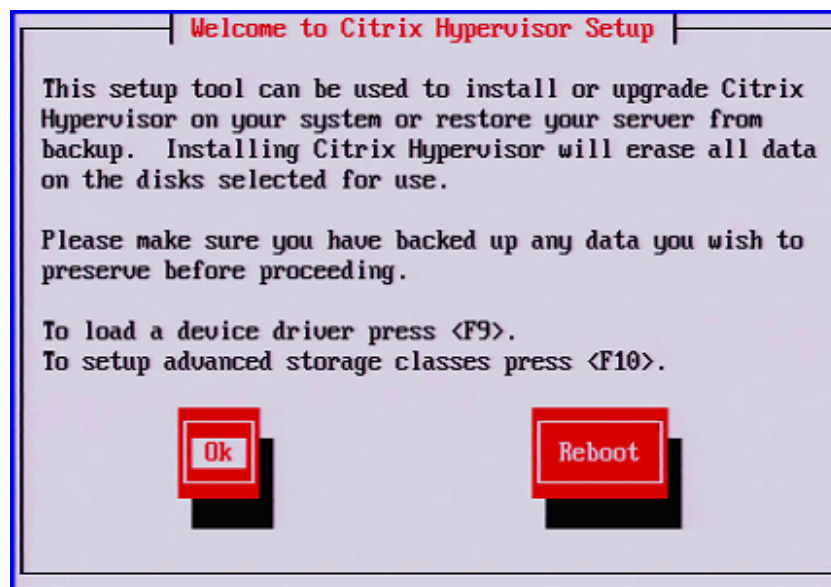


Figure 3-2. Confirm Driver Update

6. Go to the shell by pressing keys “ALT+F2” and load the OOB driver.

```
root@localhost ~]# modprobe -v qla2xxx  
insmod /lib/modules/4.19.0+1/kernel/drivers/scsi/scsi_transport_fc.ko  
insmod /lib/modules/4.19.0+1/kernel/drivers/nvme/host/nvme-core.ko  
insmod /lib/modules/4.19.0+1/kernel/drivers/nvme/host/nvme-fabrics.ko  
insmod /lib/modules/4.19.0+1/kernel/drivers/nvme/host/nvme-fc.ko  
insmod /lib/modules/4.19.0+1/updates/qla2xxx.ko
```

Figure 3-3. Loading the OOB driver

7. Return to the GUI installer by pressing “ALT+F1” and continue the OS installation without removing the Driver Update Disk CD.
8. Skip the supplemental package installation.
9. Reboot the system when prompted and remove the OS installer CD and driver update CD.
10. Hypervisor should boot with new driver installed.

NOTE

Both the OS install disk and driver update disk should be inserted during the installation process.

Installing Fibre Channel in VMware ESXi

For information on installing Fibre Channel in VMware ESXi, refer to the *Read Me* and *Release Notes* documents for VMware ESXi 8.0 U3 and 9.x.

NOTE

For specific service packs (SPs) and OS updates, go to the Downloads and Documentation page at www.marvell.com (see “Downloading Updates and Documentation” on page xv).

ESXi OS Support

To obtain the latest information about supported ESXi operating systems, download the associated *Read Me* and *Release Notes* documents available at www.marvell.com (see “Downloading Updates and Documentation” on page xv).

Using the Driver

For information on how to install and remove the driver, refer to the *Read Me* document for VMware ESXi 8.0 U3 and 9.x, available at www.marvell.com (see “Downloading Updates and Documentation” on page xv).

Installing Fibre Channel in Windows

For information on how to install Fibre Channel in Windows, refer to the *Fibre Channel Adapter STOR Miniport Driver Read Me* document.

NOTE

Marvell recommends performing Windows Boot from SAN using a single path regardless of FC or FC-NVMe protocol.

Windows OS Support

To obtain the latest information about supported Windows operating systems, download the associated *Read Me* and *Release Notes* documents available at www.marvell.com (see “Downloading Updates and Documentation” on page xv).

Using the Driver

For information on how to create the driver disk, install the driver, and remove the driver, refer to the *Fibre Channel Adapter STOR Miniport Driver Read Me* document.

Setting Fibre Channel Adapter Parameters

You can set parameters for Fibre Channel using QConvergeConsole CLI or QConvergeConsole Plug-ins.

Setting Fibre Channel Adapter Parameters with Interactive QConvergeConsole CLI

To set Fibre Channel Adapter parameters with interactive QConvergeConsole CLI:

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

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

Setting Fibre Channel Adapter Parameters with Noninteractive QConvergeConsole CLI

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

Configuring Target Persistent Binding

You can configure target persistent binding using either the QConvergeConsole CLI tool.

NOTE

Target persistent binding is supported only for Windows.

Configuring Persistent Binding with Interactive QConvergeConsole CLI

In the interactive mode of QConvergeConsole CLI, follow these steps to bind a selected target to a specific adapter.

To bind a target using interactive QConvergeConsole CLI:

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

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

Configuring Persistent Binding with Noninteractive QConvergeConsole CLI

In the noninteractive mode of QConvergeConsole CLI, bind a selected target to a specific adapter by issuing the following command:

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

Where:

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

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

Configuring Boot Devices

You can configure boot devices using QConvergeConsole CLI, QConvergeConsole Plug-in.

You can configure boot devices using QConvergeConsole CLI or QConvergeConsole Plug-in.

Configuring Boot Devices with Interactive QConvergeConsole CLI

In the interactive mode of QConvergeConsole CLI, open the **Adapter Configuration** menu, and then select **Configure Boot Devices**. For more information on configuring boot devices with the interactive mode of QConvergeConsole CLI, refer to the section about the Fibre Channel interactive commands in the QConvergeConsole.

Configuring Boot Devices with Noninteractive QConvergeConsole CLI

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

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

Where:

`hba instance` = Adapter number (issue the `-g` command to find)
`hba wwpn` = Worldwide port name of the adapter
`target wwnn` = Worldwide node name of the target

target wwpn = Worldwide port name of the target
target id = ID to which the target is bound
lun id = ID of the LUN

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

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

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

Configuring Virtual Ports (NPIV)

You can configure NPIV using the QConvergeConsole CLI tool. This section also covers NPIV quality of service (QoS).

NOTE

To configure NPIV on VMware ESXi, refer to the *User's Guide—Marvell Fibre Channel Adapters for VMware ESXi 8.0 U3 and 9.0* (part number SN0154529-00), “Deploying NPIV for VMware ESXi 8.0 U3 and 9.0” section, available at www.marvell.com (see “[Downloading Updates and Documentation](#)” on page xv).

NPIV is not supported for ESXi 9.0 and later.

Configuring NPIV with Interactive QConvergeConsole CLI

In the interactive menu of QConvergeConsole CLI, open the **Adapter Configuration** menu, and then select **N_Port ID Virtualization (NPIV)**. The **N_Port ID Virtualization (NPIV)** menu shows options for viewing and configuring virtual ports.

For more information about configuring NPIV with interactive QConvergeConsole CLI, refer to the *User's Guide—QConvergeConsole CLI* (part number SN0054667-00).

Configuring NPIV with Noninteractive QConvergeConsole CLI

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

To create a virtual port with an automatic WWPN:

- Issue the following command:

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

Where:

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

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

To create a virtual port with a specific WWPN:

1. Issue the following command:

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

Where:

hba instance = Adapter number (issue the -g command to find)
hba wwpn = Worldwide port name of the adapter
vport hex = World wide port name of the virtual port with the two hexadecimal digits in byte three supplied by the user

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

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

NPIV Quality of Service

The 2800 Series Adapter solution provides for standards-based QoS, ensuring high-quality performance for applications that require preferential delivery. The QConvergeConsole QoS solution assigns QoS levels to virtual ports (NPIV ports) in contrast to end-to-end QoS, which assigns QoS levels from initiator to target across the fabric. For information about end-to-end (CS_CTL) QoS, see [“Configuring End-to-End \(CS_CTL\) QoS” on page 33](#).

You can configure the NPIV QoS using the **priority method** or the **bandwidth method**, as shown in [Figure 3-4](#).




Figure 3-4. Priority- and Bandwidth-based QoS

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

- The **priority method** assigns priority levels (low, medium, or high) to virtual ports based on the QoS required for the applications or VMs using the virtual port. This method provides for various levels of latency requirements across virtual ports. See [“Setting QoS by Priority” on page 22](#).
- The **bandwidth method** assigns bandwidth to applications or VMs that are using virtual ports as a percentage of the total bandwidth available (64Gbps) or a specific value, such as 1Gbps. This method provides a minimum bandwidth guarantee. The bandwidth-based QoS configuration applies only when the total bandwidth requirement of the applications exceeds the available bandwidth. The bandwidth is not reserved for a specific virtual port, and unused bandwidth is shared among the other virtual ports. See [“Setting QoS by Bandwidth” on page 23](#).

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

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

- Server operating system: Microsoft Windows Server® 2022 or later
- Server support for MSI-X
- A 2800 Series Adapter
- Latest version of the QLogic STOR miniport driver, available from the Downloads and Documentation page at www.marvell.com (see [“Downloading Updates and Documentation” on page xv](#)).
- Fibre Channel switch that supports NPIV
- Physical port attached in a point-to-point connection

Setting QoS by Priority

To set the QoS by priority with interactive QConvergeConsole CLI:

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

Setting QoS by Bandwidth

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

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

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

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

NOTE

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

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

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

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

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

Configuring Driver Parameters

You can configure driver parameters for Fibre Channel using QConvergeConsole CLI.

Configuring Fibre Channel Driver Parameters with Interactive QConvergeConsole CLI

To configure Fibre Channel driver parameters using interactive QConvergeConsole CLI:

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

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

Configuring Fibre Channel Driver Parameters with Noninteractive QConvergeConsole CLI

To configure the driver settings in the noninteractive mode of QConvergeConsole CLI, issue the following command:

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

Where:

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

`hba wwpn` = Worldwide port name of the adapter

`param name` = Name of the parameter

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

PO Persistently bound targets only

PN Present persistently bound targets plus any new targets with driver default

NO Present targets with driver default

BW Bind devices by WWPN

BP Bind devices by port ID

`param value` = New value of the parameter

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

Configuring Selective LUNs

You can configure selective LUNs using QConvergeConsole CLI on Windows. To enable selective LUNs of a target, the selected target must be currently persistently bound. See [“Configuring Target Persistent Binding” on page 17](#) for more additional information.

To configure selective LUNs with interactive QConvergeConsole CLI:

1. On the QConvergeConsole CLI **Adapter Configuration** menu, select **Selective LUN Mapping**.
2. On the port menu, select a port to open the **Selective LUN Mapping** menu.

The **Selective LUN Mapping** menu contains options for displaying LUN configuration, manually configuring LUNs, and auto-configuring LUNs.

The following example shows the **Selective LUN Menu (Display)** and the general view of selective LUNs:

```
QConvergeConsoleCLI
Version 3.0.x (Build xx)

Selective LUN Menu (Display)

=====
HBA          : 0 Port: 1
SN           : RFE1305T342151
HBA Model    : QLE2872
HBA Descr.   : QLE2872 QLogic 2-port 64Gb Fibre Channel Adapter
FW Version   : 132.00.00
WWPN         : 21-00-00-0E-1E-13-49-A0
WWNN         : 20-00-00-0E-1E-13-49-A0
Host NQN     : nqn.2014-08.org.nvmexpress:uuid:
              37363836-3239-4d32-3233-313430315759
Host ID      : 363836373932324d3233313430315759
Link         : Online
=====

1: Display Selective LUN Mapping
2: Configure a specific LUN
3: Configure all LUNs
```

3–Fibre Channel Configuration

Configuring Selective LUNs

(p or 0: Previous Menu; m or 98: Main Menu; ex or 99: Quit)
Please Enter Selection: **1**

QConverge ConsoleCLI

Version 3.0.x (Build xx)

Selective LUN Mapping

```
=====
HBA          : 0 Port: 1
SN           : RFE1305T342151
HBA Model    : QLE2872
HBA Descr.   : QLE2872 QLogic 2-port 64Gb Fibre Channel Adapter
FW Version   : 132.00.00
WWPN         : 21-00-00-0E-1E-13-49-A0
WWNN         : 20-00-00-0E-1E-13-49-A0
Host NQN     : nqn.2014-08.org.nvmexpress:uuid:
              37363836-3239-4d32-3233-313430315759
Host ID      : 363836373932324d3233313430315759
Link         : Online
=====
```

- 1: Display current setting of a specific LUN
- 2: Display current setting of all LUNs

(p or 0: Previous Menu; m or 98: Main Menu; ex or 99: Quit)
Please Enter Selection: **2**

```
-----
HBA Instance 0: QLE2872 Port 1 WWPN 21:00:00:0e:1e:13:49:a0 Port ID 01:07:00
Link: Online
-----
```

Enable	Type	Target/LUN Info	Port Name	LUN ID
Yes	Disk	ST336707FC	22:00:00:11:c6:2d:5a:3e	0
Yes	Disk	ST336707FC	22:00:00:0c:51:68:1f:9a	0
Yes	Disk	ST336707FC	22:00:00:0c:51:68:35:d9	0

Press <Enter> to continue: _

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

To configure selective LUNs with non-interactive QConvergeConsole CLI:

- To enable (select) or disable (deselect) a LUN for a specific target on a specific adapter, issue the following command:

```
# qauccli -pr fc -m (<hba instance> | <hba wwpn>) {<target  
wwnn> <target wwpn>} <lun id> (select | deselect)
```

Where:

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

- To enable (select) or disable (deselect) all LUNs for a specific target on a specific adapter, issue the following command:

```
# qauccli -pr fc -m (<hba instance> | <hba wwpn>) (select |  
deselect) {<target wwnn> <target wwpn>}
```
- To enable (select) or disable (deselect) all LUNs of all targets on a specific adapter, issue the following command:

```
# qauccli -pr fc -m (<hba instance> | <hba wwpn>) (select |  
deselect) all
```

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

Setting Queue Depth

Values range from 0–65535. Queue depth refers to the maximum LUN queue depth to report for target devices. Values range from 0–65535. The default value for ESXi 8.0 U3 and 9.x is 64.

To change the queue depth, edit the `/etc/modprobe.d/scsi.conf` file and add the following line:

```
options qla2xxx ql2xmaxqdepth=<value>
```

where `<value>` is the queue depth value.

Enabling Fabric-Assigned World Wide Port Name

This section describes how to enable the fabric to assign the 2800 Series Adapters WWPN (FA-WWPN) on the following platforms:

- [Enabling FA-WWPN with Interactive QConvergeConsole CLI](#)
- [Enabling FA-WWPN with Noninteractive QConvergeConsole CLI](#)

When using the FA-WWPN feature:

- FA-WWPN must be enabled on both the initiator (2800 Series Adapter) and the Brocade switch for the FA-WWPN to be assigned to the 2800 Series Adapter.
- If FA-WWPN is disabled on the 2800 Series Adapter and enabled on the Brocade switch, the FA-WWPN cannot be assigned to the 2800 Series Adapter.
- If FA-WWPN is enabled on the 2800 Series Adapter and disabled on the Brocade switch, the FA-WWPN cannot be assigned to the 2800 Series Adapter.

Enabling FA-WWPN with Interactive QConvergeConsole CLI

To enable FA-WWPN with interactive QConvergeConsole CLI:

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

Enabling FA-WWPN with Noninteractive QConvergeConsole CLI

To enable FA-WWPN with noninteractive QConvergeConsole CLI:

- Issue the following command:

```
# qauccli -pr fc -n <hba instance>|<hba WWPN>
{<EnableFabricAssignWWN>|<FN> 1}
```

Where:

hba instance = Adapter number (issue the `-g` command to find)
hba wwpn = Worldwide port name of the adapter
EnableFabricAssignWWN = 1 enables the fabric to assign a WWN to this port
or FN 0 disables the fabric from assigning WWN to this port

To disable FA-WWPN with noninteractive QConvergeConsole CLI:

- Issue the following command:

```
# qauccli -pr fc -n <hba instance>|<hba WWPN>
{<EnableFabricAssignWWN>|<FN> 0}
```

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

Using Fabric-Assigned Boot LUN

This section describes how to enable a fabric-assigned boot LUN on the following platforms:

- [Using a Fabric-Assigned Boot LUN with Interactive QConvergeConsole CLI](#)
- [Using a Fabric-Assigned Boot LUN with Noninteractive QConvergeConsole CLI](#)

Using a Fabric-Assigned Boot LUN with Interactive QConvergeConsole CLI

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

1. From the main menu, select the adapter type (**Fibre Channel Adapter**).
2. Select an adapter port.
3. On the **Adapter Configuration** menu, select **Configure Boot Devices**.
4. On the **Boot from SAN BIOS Configuration** menu, select **Boot-from-SAN Drive Mapping Parameters**, and then press ENTER.

5. Select **Configure BIOS Boot Parameters**, and then press ENTER.
6. On the **Boot from SAN BIOS Parameters** menu, select **Fabric Assign Boot LUN**, and then press ENTER.
7. To enable this feature, type 1.

Using a Fabric-Assigned Boot LUN with Noninteractive QConvergeConsole CLI

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

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

Where:

```
hba instance = Adapter number (issue the -g command to find)  
hba wwpn     = Worldwide port name of the adapter  
FabricAssignBootLUN or fb = Configure enable fabric assign boot LUN setting
```

FDMI Enhancements Support

FDMI enables the management of devices such as host adapters. Enabling additional attributes using FDMI provides you with detailed information about the 2800 Series Adapter through a centralized management entity such as the Fibre Channel fabric switches.

To verify the additional 2800 Series Adapter attributes registered on the Brocade switch, issue the `fdmishow` command on Brocade Fibre Channel switches, and then confirm that all the requested attributes are displayed in the switch. For details, refer to your Brocade switch documentation.

The following examples show output of the `fdmishow` command.

Port attributes (Brocade switch):

```
FC4 Types: FCP FC-NVMe  
Supported Speed: 16 32 64 Gb/s  
Port Speed: 64 Gb/s  
Max Frame Size: 2048 bytes  
Device Name: qla2xxx:host1140  
Host Name: evt87.qlogic.org  
Node Name: 20:00:f4:c7:aa:08:2e:12  
Port Name: 21:00:f4:c7:aa:08:2e:12  
Port Type: Nx_Port (0x7f)
```

Port Symb Name: QLE2872 FW:v9.12.00 DVR:v10.02.09.00.a6-k1-debug
Class of Service: 3
Fabric Name: 10:00:d8:1f:cc:47:14:58
FC4 Active Type: FCP
Port State: 0x2
Discovered Ports: 0x1
Port Identifier: 0x012100

Port attributes (Cisco switch):

Port-id: 21:00:34:80:0d:73:05:59
Supported FC4 types:scsi-fcp NVMe
Supported Speed :8G 16G 32G
Current Speed :32G
Maximum Frame Size :2048
OS Device Name :qla2xxx:host12
Host Name :nightly28fcini650

HBA attributes (Brocade switch):

Node Name: 20:00:f4:c7:aa:08:2e:12
Manufacturer: Marvell Semiconductor, Inc.
Serial Number: RFD2213X41486
Model: QLE2872
Model Description: QLogic QLE2872 Dual Port 64/32/16GFC PCIe Gen4 x8 Adapter
Hardware Version: MA2810401-C5 00
Driver Version: 10.02.09.00.a6-k1-debug
Option ROM Version: 0.00
Firmware Version: 9.12.00 (d0d5)
OS Name and Version: Linux 3.10.0-957.el7.x86_64 x86_64
Max CT Payload Length: 512 words
Symbolic Name: QLE2872 FW:v9.12.00 DVR:v10.02.09.00.a6-k1-debug
Vendor Info: 0x00001077
Number of Ports: 1
Fabric Name: 10:00:d8:1f:cc:47:14:58
Bios Version: BIOS 0.00
Vendor Identifier: QLGC
Local Port database contains (port 32):
21:00:f4:c7:aa:08:2e:12

HBA attributes (Cisco switch):

```
CISCO9706FW1# show fDMI database detail hba-id 21:00:f4:c7:aa:06:5d:b9 vsan 1
Node Name           :20:00:f4:c7:aa:06:5d:b9
Manufacturer        :QLogic Corporation
Serial Num          :RFD2205W56894
Model               :QLE2872
Model Description:QLogic QLE2872 Dual Port 64/32/16GFC PCIe Gen4 x8 Adapter
Hardware Ver        :MA2810401-K2 00
Driver Ver          :10.02.07.00.a15-k-debug
ROM Ver             :0.00
Firmware Ver        :82.12.01 (f0d5)
OS Name/Ver         :Linux 3.10.0-957.el7.x86_64 x86_64
Port-id: 21:00:f4:c7:aa:06:5d:b9
Supported FC4 types:scsi-fcp NVMe
Supported Speed      :16G 32G 64G
Current Speed        :64G
Maximum Frame Size  :2048
OS Device Name       :qla2xxx:host54
Host Name           :evtbench3.fwlab.local
```

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

```
Cisco# sh fDMI database detail
```

```
-----
HBA-ID: 21:00:00:24:ff:8f:d8:88
-----
```

```
Node Name           :20:00:00:24:ff:8f:d8:88
Manufacturer        :Marvell Semiconductor Inc.
Serial Num          :AFD1536Y03339
Model               :QLE2872
Model Description:QLogic 64Gb LR 2-port FC to PCIe Gen7 x8 Adapter
Hardware Ver        :BK3210407-05 01
Driver Ver          :8.08.00.04.07.0-k
ROM Ver             :3.62
Firmware Ver        :8.08.03 (d0d5)
```

```
Port-id: 21:00:00:24:ff:8f:d8:88
```

```
Supported FC4 types:scsi-fcp
```

```
Supported Speed      :16G 32G 64G
Current Speed        :64G
Maximum Frame Size   :2048
OS Device Name       :
Host Name            :localhost.localdomain
```

Configuring End-to-End (CS_CTL) QoS

NOTE

End-to-end QoS requires the initiator, target and fabric to support CS_CTL QoS. QoS must also be supported by both the switch and target devices.

The class specific control (CS_CTL) QoS implementation is an end-to-end solution, encompassing the initiator (2800 Series Adapter), the switches, and the target (storage). The 2800 Series Adapter stamps the priority value (high, medium, low) in the CS_CTL field of the Fibre Channel frame. All three entities (initiator, switch, and target) honor this priority value at their respective levels.

This feature is enabled by default in the 2800 Series Adapter Windows and VMware drivers. Therefore, when all three entities support QoS (and it is enabled on the switch and the target), a default priority value is populated in all FCP_CMND, FCP_DATA, FCP_XFER_RDY, FCP_RSP, and FCP_CONFIRM frames.

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

- [Configuring CS_CTL QoS on the Switch](#)
- [Configuring CS_CTL QoS on the Storage Device](#)
- [Configuring QoS Priority Levels](#)

Configuring CS_CTL QoS on the Switch

To enable CS_CTL QoS on the Brocade SW6510 switch:

1. From the switch, issue the following command:

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

2. A list of attributes appear one at a time. After either accepting the default or entering a different value, the next attribute appears.
 - a. To accept the default for `cfgload` attributes, Custom attributes, and system attributes, press the ENTER key.
 - b. For `fos` attributes, type **y**.
 - c. For the CSCTL QoS Mode attribute, type **1**.

For example:

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

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

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

This command should return the following response:

```
fos.csctlMode:1
```

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

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

Configuring CS_CTL QoS on the Storage Device

For QoS configuration instructions, see your device's user guide.

Configuring QoS Priority Levels

The QoS priority levels are set in the 2800 Series Adapter virtual ports (vPorts). The Windows and VMware adapter drivers assign default vPort priority levels. The following sections describe how to change the default vPort priority levels in Windows and VMware ESXi systems.

Changing the Priority Levels on Windows

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

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

To create a vPort in interactive QConvergeConsole CLI:

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

To change the vPort priority level using interactive QConvergeConsole CLI:

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

Enabling Forward Error Correction

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

FEC is mandatory when operating at 32Gb or 64Gb line rates. 2800 Series Adapters automatically enable FEC when operating at 32G or 64G. If the adapter is operating at 16G, enabling FEC is optional.

NOTE

The FEC feature is supported only if the current adapter port is connected to a Fibre Channel switch with FEC support. For more information, refer to the section about setup tips in the *Application Note—Enabling Forward Error Correction* (part number 83830-514-03).

FEC Process Overview

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

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

NOTE

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

Enabling and Disabling FEC

FEC is mandatory when operating at 32Gb or 64Gb line rates. 2800 Series Adapters automatically enable FEC when operating at 32G or 64G. If the adapter is operating at 16G, enabling FEC is optional.

Use the appropriate procedure to enable or disable FEC:

FEC:

- [Enabling and Disabling FEC on VMware ESXi](#)
- [Enabling and Disabling FEC on Linux and Windows](#)

Enabling and Disabling FEC on VMware ESXi

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

Enabling and Disabling FEC on Linux and Windows

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

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

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

All guides are available on the Support pages at www.marvell.com (see “Downloading Updates and Documentation” on page xv).

Using Buffer-to-Buffer Credit Recovery

Buffer-to-buffer credit recovery (BB-CR) enables two FC peer ports (N_Port, F_Port, or E_Port) to periodically send and receive the quantity of receiver ready (R_RDY) signals transmitted. The BB-CR feature allows the peer port to recover from possible R_RDY signals lost over a lossy link. BB-CR enables two FC ports logged in with each other to recover lost buffer-to-buffer credits. These lost credits can impact throughput, cause link resets, and disrupt traffic flow.

NOTE

BB-CR takes effect only when the operational link speed is 16G or higher.

During the login process, the peer ports exchange a nonzero BB_SC_N value. During frame transmission/reception, one port maintains counters to track the R_RDY and frames received. Based on a periodic exchange of BB_SC_N values, the receiving port determines how many R_RDYs should have been received. BB-CR compares this value with the actual counters maintained at the receiving port, from which the receiver can compute the actual quantity of credits lost, and thus recover without incurring a link reset.

By default, BB-CR functionality is enabled on the 2800 Series Adapters that are operating at 16Gbps or higher, and that integrate with Brocade Fabric OS Manager, Brocade Network Advisor. You can disable this feature using the boot code (UEFI) or QConvergeConsole CLI.

For details about how to recover buffer credits, see the Marvell management tool user guide for your system; the guides are listed in “Related Materials” on page xii.

Resetting Driver and Firmware Counters

When 2800 Series Adapters are running I/Os, they collect and display many driver and firmware counters, which are an accumulation of normal traffic statistics or error conditions. These counters can help you assess the traffic flow, identify I/O bottlenecks, and isolate link issues that affect the SAN performance.

Some counters are accumulated in the adapter's hardware and persist throughout the life of the adapter. Other counters are transient in memory and are reset after a reload or re-initialization. The 2800 Series Adapter software provides the ability to reset counters as needed. For example, you may want to see if counters or rates change as a result of taking corrective actions in response to issues or after changing configuration parameters.

For details about how to reset counters, see the management tools user guide for your system; the guides are listed in [“Related Materials” on page xii](#).

Configuring VMID

VM storage I/O activity automatically generates VMID tagging in FC frames. VMID tags I/O frames with an ID that identifies the virtual machine that sent or received them. This ID allows management tools to monitor traffic flows down to the virtual machine level. In ESXi 8.0 U3, VMID is disabled by default for both inbox and async drivers. In ESXi 9.x, VMID is enabled by default for inbox drivers, but remains disabled for async drivers. VMID requires a Brocade or Cisco switch running a supported FOS or NX-OS version. For detailed VMID information, see the documentation for your switch.

- To enable VMID (and disable QoS), type the following command and then reboot the system:

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

- To disable VMID, type the following command and then reboot the system:

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

For details about configuring VMID, see—*User's Guide Marvell® Fibre Channel Adapters for VMware® ESXi 8.0 U3 and 9.0* (part number SN0154529-00).

NOTE

ESXi 8.0 U3 and ESXi 9.x now associate and display VC UUIDs with VMIDs. For more information, see the VMware vSphere 8.0 U3 or VMware vSphere 9.0 documentation.

Universal SAN Congestion Mitigation (USCM)

NOTE

Universal SAN congestion mitigation (USCM) is supported on the QLE2870, QLE2872, QLE2874 adapters. See [“Prerequisites” on page 40](#) for more information.

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

NOTE

Marvell USCM requires a Brocade or Cisco switch running FOS 9.0 or later, or NX-OS 9.2(1) or later, respectively.

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

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

NOTE

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

For example, to enable USCM in QConvergeConsole CLI:

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

For example, to disable USCM in QConvergeConsole CLI:

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

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

For more information about using USCM, see the appropriate documents in [“Related Materials” on page xii](#).

Prerequisites

USCM is supported on the following drivers, which you can download from the Marvell Web site.

- Linux: 10.02.04.00 or later
- Windows: 9.4.11.20 or later
- ESX: 3.1.64.0/4.1.34.0/5.3.1 or later

USCM is supported on the following Marvell management tools:

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

To use these tools for USCM, see the associated user’s guides in [“Related Materials” on page xii](#).

SAN Congestion Mitigation

USCM utilizes two mitigation techniques: throttling the bandwidth to and from the SAN, and quarantining congested flows to slower virtual lanes. In USCM, these techniques are referred to as USCM Profiles and USCM Virtual Lanes, respectively, and are described in the following sections.

These USCM features can be controlled:

- Through the Marvell QLogic applications listed in [“Prerequisites” on page 40](#)
- Using driver modules or registry parameters

USCM I/O Throttling Profiles

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

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

The USCM I/O Throttling Profiles feature provides controls by which an administrator can select an appropriate mitigation strategy using one of the Marvell applications listed in [“Prerequisites” on page 40](#) or using a driver module/registry parameter.

Using Driver Modules or Registry Parameters for USCM Profiles

You can control USCM I/O Throttling Profiles through either a driver module (Linux, VMware ESXi) or registry parameter (Windows). This parameter applies to all adapter ports on the server.

scmrprofile	Windows/driver registry parameter	Requests are throttled based on defined driver profiles. The default profile is Monitor Only. USCM Profiles are: 0 = Monitor Only (default) 1 = Conservative 2 = Moderate 3 = Aggressive
ql2x_scmr_profile	Linux/driver module parameter	
ql2x_scmr_driver_profile	VMware ESXi/driver module parameter	

USCM Virtual Lanes

Overview

NOTE

USCM Virtual Lanes are available only when the adapter is connected to a supported Cisco switch running a NX OS version 9.3.2, or later, that supports this feature in the fabric. This feature is also supported on Brocade switches running FOS 9.2.2a and later that supports this feature in the fabric.

The Marvell Fibre Channel Adapters support the *Virtual Lanes* feature. Traffic can be steered on each virtual lane; each lane can carry a different storage protocol or a different session within the same protocol. This configuration is useful in multi-protocol environments (for example, NVMe and FCP) where an administrator may not want congestion on one protocol to affect the other protocol.

Within USCM, the driver prioritizes the use of Virtual Lanes over I/O throttling based on profiles for peer congestion. For example, when there is congestion, the driver switches the session associated with the slow device to a *slow* Virtual Lane. If that peer continues to be congested, then the host may throttle requests to that device.

USCM Virtual Lanes can be enabled using one of the Marvell applications listed in [“Prerequisites” on page 40](#) or using a driver module/registry parameter.

Using a Driver Module or Registry Parameter for USCM Virtual Lanes

You can control Virtual Lanes through either a driver module (Linux, VMware ESX) or registry parameter (Windows). This parameter applies to all adapter ports on the server.

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

virtuallane	Windows/driver registry parameter	When enabled, the driver switches a session with a slow device to a slow Virtual Lane. Virtual Lane values are: 0 = Disable 1 = Enable
ql2xvirtuallane	Linux and VMware ESXi/driver module parameter	

FPIN Link Integrity Marginal Path Detection Support

The Device Mapper (DM) Multipath FPIN link integrity (FPIN-LI) marginal path detection is currently supported only on RHEL 9 and SLES16 or later, OSs with Brocade and Cisco fabrics. DM Multipath support for FC-NVMe connections is supported in RHEL 9.

RHEL 9 and SLES16 or later all support marginal path detection natively. Therefore, you do not have to install the Fibre Channel transport services daemon on these OSs. However, changes are needed in the `multipath.config` file for marginal path detection to work correctly. The steps are provided in [“Multipath Configuration File Changes RHEL9 and Later; SLES15 SP7, SLES16”](#) on page 43.

ESXi 8.0 U3 or later: the configuration information is provided in [“Multipath Configuration File Changes for ESXi 8.0 U3 and Later”](#) on page 43.

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

Multipath Configuration File Changes for RHEL 9 and Later; SLES15 SP7, SLES16

Edit the `multipath.conf` file as follows:

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

```
marginal_pathgroups fpin
```

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

NOTE

RHEL 9.4 and later support FPIN marginal pathgroups for both FC-NVMe and FCP protocols. For more information, see Red Hat RHEL 9.4 Release Notes and documentation.

Multipath Configuration File Changes for ESXi 8.0 U3 and Later

For more information about configuring FPINs for FPIN-LI/MPIO operation in Pluggable Storage Architecture (PSA), see the VMware vSphere ESXi 8.0 U3 and ESXi 9.x documentation.

Using Fibre Channel Diagnostics

The following sections show and describe the 2800 Series Adapter Fibre Channel diagnostic functions:

- [Device Ping Testing](#)
- [Fibre Channel Ping Testing \(Echo ELS\)](#)
- [CT Ping Testing](#)

- [CT Fibre Channel Trace Route Testing](#)
- [Link Cable Beacon Extended Link Service Command](#)
- [Diagnostic Port Testing](#)
- [Read Diagnostic Parameters Extended Link Service Command](#)

Device Ping Testing

QConvergeConsole supports device ping testing. The device ping test sends a SCSI Inquiry command to each selected target, measuring the time it takes for the target to respond.

NOTE

Before and after each SCSI command, the link statistics are collected to verify that no counters are incremented during each pass.

For more information on performing general diagnostic tests and viewing the diagnostic test results, refer to the help system, *QConvergeConsole Help*.

Fibre Channel Ping Testing (Echo ELS)

QConvergeConsole supports Fibre Channel (FC) ping testing. The FC ping test sends an ECHO extended link service (ELS) from the adapter port to the selected switch, fabric, or target. For more information on performing general diagnostic tests and viewing the diagnostic test results, refer to the appropriate QConvergeConsole user guide in [“Related Materials” on page xii](#).

CT Ping Testing

QConvergeConsole supports common transport (CT) ping testing. The CT ping test pings each device attached to the adapter port. For more information on performing general diagnostic tests and viewing the diagnostic test results, refer to the appropriate QConvergeConsole user guide in [“Related Materials” on page xii](#).

NOTE

Run the CT Ping test only from a switch with either FOS version 6.4.x or later or NX-OS version 8.4.1 or later.

CT Fibre Channel Trace Route Testing

QConvergeConsole supports CT Fibre Channel trace route (FTR) testing. The CT FTR test traces the route from the adapter port to the target device.

NOTE

All inter-link switches between the initiator and the target must have the Brocade switch firmware 7.1.1 or later and must support CT FTR for the test to work end-to-end. To confirm that your target is supported, contact Technical Support.

For more information on performing general diagnostic tests and viewing the diagnostic test results, refer to the appropriate QConvergeConsole user guide in [“Related Materials” on page xii](#).

Link Cable Beacon Extended Link Service Command

NOTE

Run the LCB test only from a switch with FOS version 7.4.x or later or NX-OS 8.3.1 or later.

The link cable beacon (LCB) extended link service command locates the corresponding ends of a cable by flashing the specific LEDs (see [“Adapter LEDs” on page 65](#) for more information) on the corresponding ports of a link.

The LCB command can be sent through a Brocade switch using the following commands, which show examples of enabling, disabling, or verifying the LCB. For details, refer to your switch documentation.

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

To execute the LCB command through a Cisco switch, use the following commands to issue and stop beacon. For details, refer to your switch documentation.

```
# beacon interface fc1/1 peer
# no beacon interface fc1/1 peer
```

Diagnostic Port Testing

QConvergeConsole CLI supports diagnostics port testing to identify cable connection problems from a Brocade switch (FOS 7.4.x or later) or a Cisco switch (NX-OS 8.2.1 or later). For information about diagnostics port testing from the host server, see the *QConvergeConsole Help* or the *User's Guide—QConvergeConsole CLI* (part number SN0054667-00). For information about performing diagnostics port testing from the FC switch, see [“Troubleshooting in Diagnostics Port Mode” on page 70](#).

Read Diagnostic Parameters Extended Link Service Command

The read diagnostic parameters (RDP) extended link service (ELS) command is a powerful diagnostic tool that can be used while a link is up and passing traffic. When the adapter port receives an RDP command, it responds with a comprehensive diagnostic report containing several blocks of data useful for monitoring the fabric or trouble-shooting. The report includes:

- SFP Diagnostic Parameters: temperature, V_{CC} , and Tx and Rx power
- Port Speed Descriptor: speed capabilities and operating speed
- Link Error Status Block Descriptor
- Port Names Descriptor: WWNN and WWPNN
- FEC Status Descriptor: corrected and uncorrectable blocks

NOTE

You must run the RDP test from a switch with firmware as follows:

- Brocade switches: FOS 9.1.1 or later
- Cisco switches: NX-OS 8.2.1 or later

All of Marvell's management tools support running the RDP test. For information about running the RDP test from the host server, see the user's guide for your Marvell adapter's management tool in [“Related Materials” on page xii](#).

Running RDP Test from a Brocade 16/32/64G Switch

The following example shows the RDP ELS command output for Port 13 on a Brocade 64G switch (9.1.1 or later):

```
G720:FID128:admin> sfpshow 16 -link -force
Identifier: 3      SFP
Connector:  7      LC
Transceiver: 0204406000000000 16,32,64_Gbps M5 sw Inter,Short_dist
Encoding:    8      PAM4
Baud Rate:   255    (units 100 megabaud)
Length 9u:   0      (units km)
```

3-Fibre Channel Configuration

Using Fibre Channel Diagnostics

Length 9u: 0 (units 100 meters)
Length 50u (OM2): 2 (units 10 meters)
Length 50u (OM3): 7 (units 10 meters)
Length 62.5u: 0 (units 10 meters)
Length 50u (OM4): 10 (units 10 meters)
Vendor Name: BROCADE
Vendor OUI: 00:05:1e
Vendor PN: 57-1000495-01
Vendor Rev: B
Wavelength: 850 (units nm)
Options: 0a3a Loss_of_Sig,Tx_Fault,Tx_Disable
BR Max: 231
BR Min: 0
Serial No: MAA12034C005585S
Date Code: 200818
DD Type: 0x68
Enh Options: 0xfa
Status/Ctrl: 0x30
Pwr On Time: 2.62 years (22931 hours)
E-Wrap Control: 0
O-Wrap Control: 0
Alarm flags[0,1] = 0x0, 0x0
Warn Flags[0,1] = 0x0, 0x0
Temperature: 45 Centigrade
Current: 7.092 mAmps
Voltage: 3337.5 mVolts
RX Power: -1.4 dBm (732.6uW)
TX Power: -0.8 dBm (827.2 uW)

State transitions: 1
Port Speed Capabilities 16Gbps 32Gbps 64Gbps

PEER Port SFP Info

Vendor Name: FINISAR CORP.
Serial num: P91B566
Vendor PN: FTLF8564D1BCW-QM
Vendor Rev: A
Date Code: 230615
Laser Type: Short Wave Laser
SFP Type: Optical Port Type
Connector Type: SFP+
Following SFP Parameters are Valid
Temperature: 45 Centigrade [Range -128 - +128 C]

3–Fibre Channel Configuration

Using Fibre Channel Diagnostics

```
Current:      8.498   mAmps   [Range 0 - 131 mAmps]
Voltage:     3321.9  mVolts  [Range 0 - 3600 mVolts]
Rx Power:    721.3   uW      [Range 0 - 6550 uW]
Tx Power:    852.3   uW      [Range 0 - 6550 uW]
Signal Loss (Upstream) : -9.8   dBm (105.9 uW) (-0.59 dB)
Signal Loss (Downstream): -9.2   dBm (119.7 uW) (-0.66 dB)
Port Speed Capabilities      16Gbps   32Gbps 64Gbps
```

	Alarm		Warn	
	low	high	low	high
Temperature alerts (Centigrade):	-5	75	0	70
Voltage alerts (mVolts)	: 2970	3630	3135	3465
Tx Bias alerts (uA)	: 1000	12000	2000	11500
Tx Power alerts (uW)	: 89	5011	177	2511
Rx Power alerts (uW)	: 44	3162	89	2511

Last poll time: 10-23-2023 GMT Mon 18:22:59

Running RDP Test from a Cisco 16/32G Switch

The following command shows the RDP test results for port 10 from a Cisco 16/32G switch (NX-OS 8.2.1 or later).

```
Cisco# sh rdp fcid 0x590223 vsan 10
```

```
-----
                        RDP frame details
-----

Link Service Request Info:
-----

Port Speed Descriptor Info:
-----

Port speed capabilities : 64G 32G 16G
Port Oper speed        : 64000 Mbps

Link Error Status:
-----

VN PHY port type       : FC
Link failure count     : 1
Loss of sync count     : 0
Loss of signal count   : 0
Primitive sequence proto error : 0
Invalid Transmission word : 0
Invalid CRC count      : 0

Port Name Descriptor:
-----

Node WWN               : 20:00:00:24:ff:8f:d8:88
Port WWN               : 21:00:00:24:ff:8f:d8:88
Attached Node WWN      : 20:0a:00:de:fb:83:80:81
Attached Port WWN      : 20:01:00:de:fb:83:80:80
```

```
SFP Diag params:
-----
SFP flags      : Optical
SFP Tx Type    : Long Wave LC

FEC Status:
-----
Corrected blocks : 0
Uncorrected blocks : 0

Buffer Credit Descriptor:
-----
Rx B2B credit   : 186
Tx B2B credit   : 32
Port RTT        : 0 ns

Optical Product Data:
-----
Vendor Name     : FINISAR CORP.
Model No.       : FTLF1432P3BCV-QL
Serial No.      : UVE2AN5
Revision        : A
Date            : 171106

-----
                Current      Alarms      Warnings
                Measurement   High      Low      High      Low
-----
Temperature    44.25 C      75.00 C   -5.00 C   70.00 C    0.00 C
Voltage        3.38 V       3.60 V    3.00 V    3.50 V     3.10 V
Current        37.92 mA      70.00 mA   1.00 mA   68.00 mA    2.00 mA
Tx Power       0.02 dBm      4.00 dBm -10.00 dBm 3.00 dBm   -9.00 dBm
Rx Power      -4.16 dBm      3.00 dBm -20.00 dBm 2.00 dBm  -18.01 dBm
-----

Note: ++ high-alarm; + high-warning; -- low-alarm; - low-warning
```

Management Component Transport Protocol Support

All QLE2870 and QLE2872 adapters support MCTP. However, refer to *Application Note—PLDM Support in Marvell QLogic Fibre Channel Adapters* (Part Number TD-001688) about SFP Polling and MCTP support in QLE2874 adapters.

MCTP Support for QLE2874 Adapters and SFP Polling

To verify that your 2800 Series Adapters support SFP Polling and Quad Port Adapter supports MCTP, compare the versions of the installed Flash components to the minimum versions below.

Do not use the Enablement Kit for Quad Port adapters or to enable SFP Polling. Enabling MCTP for 2800 Series Quad Port adapters or SFP Polling requires flashing the MBI using QCC.

■ Minimum versions for 2874 Adapters:

- ☐ Flash Image: 10.10.04
- ☐ MPI Firmware: 3.04.03
- ☐ Preload Fw Area: 3.03.09
- ☐ FC Brd Cfg: 3.00.03

4 NVM Express over Fibre Channel

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

- [Hardware and Software Requirements](#)
- [“Host Setup” on page 53](#)
- [“Limitations and Known Issues” on page 60](#)
- [“Unloading and Reloading the Linux Driver” on page 60](#)

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

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

Hardware and Software Requirements

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

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

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

Adapters

FC-NVMe requires one or more of the following adapters, including models QLE2870, QLE2872, QLE2874.

Operating Systems

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

- Operating systems with inbox drivers:
 - ☐ Windows Server 2022 or later
 - ☐ SLES15 SP7, SLES16
 - ☐ RHEL 9 and later
 - ☐ Oracle Linux 9.4 with UEK 7 U3 and later
 - ☐ ESXi 8.0 U3, ESXi 9.0 or later
 - ☐ Ubuntu 24.04.4 LTS, Ubuntu 26.04 LTS

- Operating systems with out-of-box drivers:
 - ☐ Windows Server 2019 and later
 - ☐ SLES15 SP7, SLES16
 - ☐ RHEL 9 and later
 - ☐ ESXi 8.0 U3, ESXi 9.0 or later

Host Setup

Procedures for setting up the initiator host include the following:

- [Installing the Drivers](#)
- [Finding Host NQN](#)
- [Programming the Host NQN onto the NVMe Storage Array](#)
- [Registering the Host NQN into the Fibre Channel Fabric](#)
- [Configuring System for NVMe Auto Discovery](#)
- [Verifying the Availability of FC-NVMe Devices](#)
- [Unloading and Reloading the Linux Driver](#)
- [Limitations and Known Issues](#)

Installing the Drivers

The following instructions are for configuring Linux, Windows, and ESX Host for FC-NVMe. Installation of the OS driver is a prerequisite for configuring FC-NVMe.

Finding Host NQN

The Host NQN string is needed to configure and connect to NVMe Storage. The default Host NQN is displayed at the HII UEFI driver level. For more information on configuring NVMe BFS see *User's Guide—UEFI Human Interface Infrastructure*, 2690 Series 16GFC, 2740/2760 32GFC, 2770 Series 32GFC Fibre Channel Adapters (part number BK3254602-00).

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

Linux

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

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

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

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

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

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

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

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

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

Windows

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

The following shows an example Host NQN output:

```
-----  
Host Name : WIN-4KHVTBV0147  
Host NQN : nqn.2018-08.com.marvell:nvme.host.sys.wwpn:21000024ff784720  
HBA Instance : 2  
HBA Model : QLE2872
```

NOTE

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

ESXi

Issue the following command to list the current Host NQN:

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

Programming the Host NQN onto the NVMe Storage Array

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

Registering the Host NQN into the Fibre Channel Fabric

The NVMe Host NQN of a specific adapter can be viewed, registered or deregistered using the QCC CLI.

NOTE

This feature is supported on Brocade switch FOS 9.2.2a and later for 32G and 64G fabrics only.

See *User's Guide—QConvergeConsole CLI* (part number SN0054667-00) for more details.

Configuring System for NVMe Auto Discovery

Auto discovery is natively supported for SLES15 SP7, SLES16; RHEL 9 and later; Windows Server 2022 or later; and ESXi 8.0 U3 or later. No additional configuration is needed.

The following procedure is for configuring auto discovery on RHEL 9.0 or later only.

To set up a Linux host:

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

```
# ls -l /etc/udev/rules.d/99-nvme-fc.rules
-rw-r--r--. 1 root root 310 Dec 31 15:14
/etc/udev/rules.d/99-nvme-fc.rules
# ls -l /usr/sbin/qla2xxx/qla2xxx_nvme_boot_connections.sh
-rwxr-xr-x. 1 root root 859 Jan 31 10:56
/usr/sbin/qla2xxx/qla2xxx_nvme_boot_connections.sh
# ls -l
/usr/lib/systemd/system/qla2xxx-nvmefc-connect@.service
-rwxr-xr-x. 1 root root 859 Jan 31 10:56
/usr/lib/systemd/system/qla2xxx-nvmefc-connect@.service
# ls -l
/usr/lib/systemd/system/qla2xxx-nvmefc-connect@.service
-rw-r--r--. 1 root root 279 Jan 31 10:54
/usr/lib/systemd/system/qla2xxx-nvmefc-connect@.service
```

2. If the Linux out-of-box driver does not discover NVMe targets after a system reboot, perform the following steps to ensure the service named `qla2xxx-nvmefc-boot-connection.service` is enabled.

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

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

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

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

Following is an example response that should be seen:

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

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

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

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

Verifying the Availability of FC-NVMe Devices

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

Verifying in Linux

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

```
# lsblk
or
# nvme list
```

Example: `lsblk` Command

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

NAME	MAJ:MIN	RM	SIZE	RO	TYPE	MOUNTPOINT
nvme0n1	259:2	0	1G	0	disk	
nvme0n1p1	259:25	0	1023M	0	part	
nvme0n2	259:8	0	1G	0	disk	
nvme0n3	259:12	0	1G	0	disk	
nvme0n4	259:16	0	1G	0	disk	
nvme0n5	259:18	0	1G	0	disk	
nvme0n6	259:20	0	1G	0	disk	
nvme0n7	259:22	0	1G	0	disk	
nvme0n8	259:24	0	1G	0	disk	

Example: `nvme list` Command

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

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

Verifying in Windows

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

```
gauccli -t 0 (HBA instance 0)
Using config file: C:\Users\Administrator\gauccli.cfg
Installation directory: C:\Users\Administrator
```

4–NVM Express over Fibre Channel

Host Setup

Working dir: C:\Users\Administrator

HBA Instance 0: QLE2872Port 1 WWPN 21:00:00:24:ff:8f:e3:9a PortID 01:21:00
Link: Online (FEC)

Path : 0
Target : 0
Device ID : 0x83
Product Vendor : NVMe
Product ID : SANBlaze VLUN P2
Product Revision : -dev
Serial Number : 6001_10d0_0247_2600_0302_0000_929d_14c6.
Node Name : 20:02:00:11:0d:47:26:00
Port Name : 20:02:00:11:0d:47:26:00
Port ID : 01:1b:00
Product Type : NVME Disk (Ensure that this shows as an FC-NVMe target.)
LUN Count(s) : 51
Status : Online

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




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

Verifying in ESXi

To verify the controllers using ESXi 8.0 U3/9.0, issue the following command:

```
$ esxcli nvme controller list
```

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

Figure 4-2. Listed NVMe Controllers

Display the NVMe namespace using the following command:

```
$ esxcli nvme namespace list
```

Name	Controller Number	Namespace ID	Block Size	Capacity in MB
eui.600110d00227b20003020000929d14c6	260	1	512	4096
eui.600110d00227b202030200022acfbac0c	267	1	512	1
eui.600110d00227b20403020004b6297c9a	268	1	512	1
eui.600110d00227b2060302000641833f28	269	1	512	1
eui.600110d00227b20003020000929d14c6	274	1	512	4096
eui.600110d00227b202030200022acfbac0c	281	1	512	1
eui.600110d00227b20403020004b6297c9a	282	1	512	1
eui.600110d00227b2060302000641833f28	283	1	512	1

Figure 4-3. NVMe Namespace List

NOTE

The ESX 8.0 U3/9.0 `qlnativefc` driver may create and register two logical devices, one for NVMe and one for FCP. The vmhbas do not represent individual physical FC ports, but a logical device with which the operating system can interact. The WWPN/WWNN shows that it is the same physical FC port.

Unloading and Reloading the Linux Driver

This section applies to a locally booted Linux operating system.

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

To unload and reload the Linux driver:

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

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

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

```
/dev/nvme0
```

```
/dev/nvme1
```

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

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

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

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

```
# modprobe -r qla2xxx
```

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

```
# modprobe -v qla2xxx
```

Limitations and Known Issues

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

Linux

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

Windows

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

Boot From SAN Setup

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

NOTE

NVMe BFS is only supported on RHEL 9.4 and later, Oracle Linux 9.4 with UEK 7 U3 and later, SLES15 SP4 and later, Windows Server 2022 or later, and ESXi 8.0 U3 or later. For more information on configuring NVMe BFS see *Marvell QLogic UEFI Human Interface Infrastructure 2690 Series Enhanced 16GFC, 2740/2760 Series 32GFC, 2770 Series Enhanced 32GFC, 2870 Series 64GFC Marvell QLogic Fibre Channel Adapters User's Guide* (part number BK3254602-00).

NOTE

Custom NQN configured in HII for NVMe BFS with SLES and Windows is not supported for SLES15 SP3, and Windows Server 2022 or later. The default Host NQN string is required for NVMe BFS installation to succeed using the above listed operating systems.

Once the UEFI driver is configured using HII, and the NVMe storage has been configured with the Initiator Host NQN, there are no additional NVMe BFS steps required to detect the NVMe storage in order to install SLES15 SP3, SLES15 SP4, Windows Server 2022 or later.

For ESX 8.0 U3 or later it is recommended to change the default Host NQN in HII to match the OS default Host NQN using the procedures in the following section.

Setup for NVME Boot-from-SAN Using ESXi 8.0 U3

For VMware Host NQN is derived from the fully qualified domain name (FQDN) of the host. Following is the format for host NQN under VMware ESXi 8.0 U3 and later:

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

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

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

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

NOTE

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

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

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

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

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

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

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

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

```
#localcli nvme info set --hostnqn=default
```

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

```
# localcli nvme info get
```

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

```
localcli storage san fc list
```

Following is a partial sample output:

```
FcDevice:
```

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

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

```
# localcli storage san fc reset -A vmhba65
# esxcfg-rescan vmhba65
```

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

8. List the namespaces using the following command:

```
# localcli nvme namespace list
```

9. List the NVME controllers using the following command:

```
# localcli nvme controller list
```

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

NOTE

For ESXi 8.0 U3, if the above steps do not list the NVMe namespaces then execute the following steps:

1. Fetch the Controller ID by executing the command:

```
localcli nvme controller list
```
 2. Clear all the existing connections/sessions with the older NQN

```
localcli nvme fabrics disconnect -a <NVMe device> -n <target NVMe controller>
```
 3. Reset the port by executing the command:

```
localcli storage san fc reset -A <vmhbaX>
```
 4. List the namespaces using the following command:

```
# localcli nvme namespace list
```
-

Setup for NVME Boot-from-SAN Using ESXi 9.0 and Later

From ESXi 9.0, VMware supports SMBIOS UUID based NVMe BFS for Marvell FC adapters.

Perform the following steps to configure NVMe BFS:

1. In system setup navigate to **Device Settings**, then to **Marvell Fibre Channel Adapter Settings** on the **HBA Configuration Page**.
2. Configure NVME target LUNs with SMBIOS UUID NQN.

Configure NVMeoFC Boot – SMBIOS UUID based host NQN

Enable NVMe Boot from SAN	☐ Disabled ☒ Enabled
Host NQN	-08.org.nvmexpress:uuid:4c4c4544-0039-4610-8053-b2c04f424d32
Host ID	44454C4C390010468053B2C04F424D32

Figure 4-4. UUID NVME BIOS SETTING

3. Load the ESXi 9.x ISO installer. NVME LUNs are visible without any additional settings.

A Adapter LEDs

The 2800 Series Adapter ports have three LEDs (see [Figure A-1](#)):

- The leftmost green LED is 64Gb.
- The middle green LED is 32Gb.

- The rightmost green LED is 16Gb.




Figure A-1. QLE2800 Adapter LEDs

Table A-1 defines the LED behavior for the QLE28xx 64Gb Adapter.

Table A-1. QLE2800, 64Gb Adapter LED Scheme

LED	Power Off	Power On (Before Firmware Initialization)	Power On (After Firmware Initialization)	Firmware Fault	16Gb Link Up and Active	32Gb Link Up and Active	64Gb Link Up and Active	Beaconing
Green LED (64Gbps)	Off	On	Flashing	Flashing in sequence	Off	Off	On or flashing	Flashing
Green LED (32Gbps)	Off	On	Flashing	Flashing in sequence	Off	On or Flashing	Off	Off
Green LED (16Gbps)	Off	On	Flashing	Flashing in sequence	On or Flashing	Off	Off	Flashing

B Troubleshooting

This appendix provides procedures for solving issues by running Fibre Channel diagnostics using QConvergeConsole CLI. In addition, a two-part flowchart walks you through troubleshooting problems that may occur with devices and LUNs on the 2800 Series Adapter.

Troubleshooting with Fibre Channel Diagnostics

The following sections describe how to use QConvergeConsole CLI to perform diagnostics on the 2800 Series Adapter.

Either of following tests can be conducted only while the HBA port is connected to a Fibre Channel switch:

- Read Write Buffer Test
- Port Beacon Test.

Loopback test may be conducted while the HBA port is connected to a Fibre Channel switch by:

- Selecting the External Loopback type causes the HBA firmware to transmit the Extended Link Services (ELS) ECHO requests
- Selecting Internal Loopback type causes no transmission of any frame types

Internal Loopback Test requires no SFP to be plugged into the HBA port.

Transceiver Diagnostics Monitoring Test requires an SFP to be plugged into the HBA port.

While a loopback connector is plugged into the HBA SFP, selecting Loopback Test with External Loopback type causes the HBA firmware to transmit data frames through the loopback connector.

NOTE

Do not perform an internal loopback test at the OS if the target device is the SAN target from which the system boots.

Fibre Channel Diagnostics Using Interactive QConvergeConsole CLI

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

1. On the **Main Menu**, select **4: Adapter Diagnostics**.
2. On the **FC Diagnostics** menu, select the test type to run:
Select **1: Loopback Test** to run the loopback test.
Select **2: Read Write Buffer Test** to run the read/write buffer test.
Select **3: Transceiver Diagnostics Monitoring Interface (DMI)** to view the transceiver diagnostics information.
Select **4: Port Beacon** to run the beacon test.

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

Fibre Channel Diagnostics Using Noninteractive QConvergeConsole CLI

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

Fibre Channel Troubleshooting Diagram

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




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




Figure B-2. Fibre Channel Troubleshooting Diagram (Part 2 of 2)

Troubleshooting in Diagnostics Port Mode

Use the diagnostics port feature on the 2800 Series Adapter to identify and isolate link failures resulting from faulty modules (link, cable, or SFP) by running diagnostics over a Fibre Channel link.

- The D_Port feature is supported on Windows, Linux, and ESXi.
- D_Port requires a Fibre Channel Adapter that is linked to a Brocade or Cisco Fibre Channel 16G, 32G, or 64G switch, and the supported Brocade or Cisco 16G, 32G, or 64G Fibre Channel switch firmware.

NOTE

All Marvell QLogic-branded 16Gbps, 32Gbps, 64Gbps SFP+ transceivers support diagnostics port tests. If you have a supported non-Marvell QLogic branded transceiver, ensure that it supports the optical wrap (OWrap) and electrical wrap (EWrap) capability that is required for diagnostics port.

Brocade has deprecated support for EWrap and OWrap from FOS 9.2.0 and later.

Before running diagnostics, verify that the FC switch supports Diagnostic Port mode, and then enable the HBA port Diagnostic Port mode. The following example enables and shows the Diagnostic Port mode on HBA port instance 0, where 0 specifies the HBA port instance number.

Execute this Non-Interactive QCC CLI command to enable D_Port on the HBA:

```
PS C:\Users\Administrator> qauccli -dport 1 --enable
Using config file: C:\Program Files\QLogic
Corporation\QConvergeConsoleCLI\qauccli.cfg
HBA Parameters update complete. Changes have been saved to HBA 0
and a reboot is required.
```

```
PS C:\Users\Administrator> qauccli -dport 1 --info
Using config file: C:\Program Files\QLogic
Corporation\QConvergeConsoleCLI\qauccli.cfg
-----
QLE2874 (SN AFD1916Y07458)
-----
=====
Port 1:
=====
DPort   : Enable
```

```
PS C:\Users\Administrator> qauccli -dport 3 --show
Using config file: C:\Program Files\QLogic
Corporation\QConvergeConsoleCLI\qauccli.cfg
Starting diagnostic port test of HBA 3 (QLE2874), please wait...
Start Time: Thu Sep 12 10:35:27 2024
End Time   : Thu Sep 12 10:35:47 2024
```

```
-----
HBA Instance 3: QLE2874 Port 4 WWPN 21:00:34:80:0d:65:c4:6f PortID
64:12:00
Link: Link Down
-----
```

HBA Port	Electrical Loopback	Optical Loopback	Link Traffic
-----	-----	-----	-----
Value	04	0xd2	0xd3
0xd5			
Status	0x01	0x01	0x01
0x02			
Result	Passed	Passed	Passed
Skipped			

```
Tx Power: 0x23c8
Rx Power: 0x17d4
-----
Details:
-----
Mbx1: 0x41
Mbx2: 0xe000
Roundtrip Link Latency: 574 ns
Cable Link Distance: 17 meters
Allowable Power Loss: 0.119000
Egress Power: Tx: -0.381045 dBm, Rx: -1.852198 dBm, Diff: 1.471153 dBm
Ingress Power: Tx: -1.612768 dBm, Rx: -2.146702 dBm, Diff: 0.533934 dBm
```

While D_Port is enabled on both the switch port and 2800 Series Adapter port, the HBA port responds to diagnostic test commands received from the switch port. The switch runs the following tests:

- Electric loopback test
- Optical loopback test

NOTE

When running Diagnostic Port in a Boot from SAN environment, you should run diagnostics port tests on multipath configured boot from a SAN or fabric-assigned boot LUN. Running diagnostics port tests on a single path boot from SAN or fabric-assigned boot LUN may cause issues with OS stability in true boot from SAN environments.

For more information about verifying diagnostics port mode using QConvergeConsole CLI, refer to the section about the Fibre Channel noninteractive commands in the *User's Guide—QConvergeConsole CLI* (part number SN0054667-00). For information about running diagnostics port from the host server adapter, see the *User's Guide—QConvergeConsole CLI*.

The following sections describe how to configure, run, and view diagnostics port mode from Fibre Channel switches:

- [Configuring, Running, and Viewing Results of a Diagnostics Port on a Brocade Fibre Channel Switch](#)
- [Configuring, Running, and Viewing Results of a Diagnostics Port on a Cisco Fibre Channel Switch](#)

Configuring, Running, and Viewing Results of a Diagnostics Port on a Brocade Fibre Channel Switch

Follow the steps given below to configure and run the test, then view results. For more detail about the switch, see the Brocade switch documentation.

To enable diagnostics port mode:

From the Fibre Channel switch, issue the following commands:

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

The diagnostics port test starts automatically.

The following example configures and starts a diagnostic test on port 17 of a Brocade Fibre Channel switch.

```
# portdisable 17
# portcfgdport --enable 17
# portenable 17
```

To view the diagnostics port test results from the Fibre Channel switch:

- Issue the `portdporttest` command as follows:

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

For example:

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

This command can be issued multiple times to monitor the status and to view the results of the test.

Following is an example output for this command after the test completes (for FOS versions 9.2.0 or earlier):

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

```

Status:                                     PASSED
=====
Test           Start time      Result      EST (HH:MM:SS)
Comments
=====
Electrical loopback    01:46:08      PASSED      -----
-----
Optical loopback      01:46:54      PASSED      -----
-----
Link traffic test     -----      SKIPPED      -----
-----

Roundtrip link latency:      157 nano-seconds
Estimated cable distance:    3 meters
Buffers required:           1 (for 2112 byte frames at 16Gbps
speed)

```

The following example shows the result of Diagnostic Port test on port 16 (FOS version 9.2.0 and later):

NOTE

Brocade FOS 9.2.0 and later omit both Electrical and Optical loopback phases. The result is the Brocade switch reports listed below.

```
GAD23826_EF5_G720:FID128:admin> portdporttest --show 16
```

```
D-Port Information:
```

```
=====
```

```
Port: 16
```

```
Remote WWPN: 21:00:34:80:0d:65:c4:4c
```

```
Mode: Automatic
```

```
No. of test frames: 1 Million
```

```
Test frame size: 1024 Bytes
```

```
Pattern: jCRPAT
```

```
FEC (enabled/option/active): Yes/No/Yes
```

```
CR (enabled/option/active): Yes/No/No
```

```
Start time: Fri Sep 29 01:25:01 2023
```

```
End time: Fri Sep 29 01:26:57 2023
```

```
Status: SKIPPED
```

```
=====
```

```
Test Start time Result EST (HH:MM:SS) Comments
```

```
=====
```

```
Link traffic test ----- SKIPPED ----- No remote port support
or
Skipped
```

To disable diagnostics port mode:

From the Fibre Channel switch, issue the following commands:

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

NOTE

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

Configuring, Running, and Viewing Results of a Diagnostics Port on a Cisco Fibre Channel Switch

Follow the steps given below to configure and run the test, then view results. For more detail about the switch, see the Cisco switch documentation.

To enable diagnostics port mode:

The following example configures and starts a diagnostic test on port 2 of a Cisco Fibre Channel switch:

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

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

```
CISCO32G1# show diagnostic result interface fc1/2 test link-dia
```

Following is an example output for this command after the test completes:

```
PWWN of peer port: 21:00:34:80:0d:73:05:83
Status: Supported (Reflector)
Reflector loopback capabilities: Xcvr-optical Electrical
Time of Test: Wed Aug 24 15:17:00 2022
Total time taken: 42 seconds
```

```
=====|=====|=====|=====|=====|=====|
|          |          | Discards |          | Latency (ns) |          |
```

B–Troubleshooting

Troubleshooting in Diagnostics Port Mode

Loopback Level	Tx Frames	Rx Frames	IN	OUT	BAD WORDS	In-Switch	External	Status
Remote-Switched(R)	0	0	0	0	0	0	0	-NA-
Mac(R)	0	0	0	0	0	0	0	-NA-
Xcvr-optical(R)	48122705	48122705	0	0	0	3644	2684	Success
Electrical(R)	20000	20000		-NA-		-NA-	-NA-	Success

Overall Status : Success

Cable Length (approx. +/- 5 metres) : 10.5 metres

To disable the diagnostic port, type the following commands:

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

NOTE

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

C Revision History

Document Revision History	
Revision 01, November 2, 2018	
Revision 02, October 7, 2022	
Revision 03, March 24, 2023	
Revision 04, April 4, 2023	
Revision 05 not released. All updates from revision 05 included in revision 06	
Revision 06, November 17, 2023	
Revision 07, June 10, 2024	
Revision 08, December 06, 2024	
Revision 09, June 10, 2025	
Revision 10, May 29, 2026	
Changes	Sections Affected
■ Updated ESXi versions from ESXi 9.0 to 9.x. ■ Removed support of RHEL 8.4 or later, RHEL 8.5, RHEL 8.6, RHEL 8.7 or later, RHEL 8.10 or later ■ Removed support of SLES15 SP3 and later Added a new feature: Host NVMe Qualified Name (NQN) registration (FOS version 9.2.2a or later, for 32G and 64G fabrics only) Updated: ■ QConvergeConsole CLI version to 3.2 and later ■ Marvell QLogic QConvergeConsole Extension for Windows Admin Center version to 2.5.0 and later ■ Marvell QLogic PowerKit version to 2.6.1 and later ■ Marvell ESXCLI Plug-in for VMware vSphere version to 1.9.0 and later	All sections “Features” on page 1 “On-The-Fly Firmware Updates” on page 7

Updated RHEL version to RHEL 9.0 and later Updated the note with: “This feature is also supported on Brocade switches running FOS 9.2.2a and later that support NPVC.” Removed sentences: ■ “The Fibre Channel Transport Services daemon (<code>fctxpd</code>) must be installed for marginal path detection to work. The transport daemon improves DM Multipath functionality to take corrective actions when link integrity events are received from the switch.” ■ “The steps are provided in FC Transport daemon Installation for RHEL 8.4, 8.5, and 8.6” Updated the FPIN link integrity marginal path detection support for the following OS versions: ■ SLES16 or later Removed section “FC Transport daemon Installation for RHEL 8.4, 8.5, and 8.6” Updated OS support for FC-NVMe. Operating systems with inbox drivers: ■ SLES15 SP7, SLES16 ■ ESXi 8.0 U3 ■ Ubuntu 24.04.4 LTS, Ubuntu 26.0 LTS Operating systems with out-of-box drivers: ■ SLES15 SP7, SLES16 ■ ESXi 8.0 U3 Added section “Registering the Host NQN into the Fibre Channel Fabric” Updated SLES and RHEL versions for configuring system for NVMe auto discovery to: ■ RHEL version 9 and later ■ SLES15 SP7, SLES16	“Boot from SAN” on page 13 “USCM Virtual Lanes” on page 42 “FPIN Link Integrity Marginal Path Detection Support” on page 43 “Operating Systems” on page 53 “Host Setup” on page 53 “Configuring System for NVMe Auto Discovery” on page 55
--	--

Glossary

adapter

The board that interfaces between the host system and the target devices. Adapter is synonymous with *Host Bus Adapter*, *Host Channel Adapter*, *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.

AL_PA

Arbitrated loop physical address. A unique one-byte value assigned during loop initialization to each node loop port (NL_Port) on a loop.

alias

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

arbitrated loop physical address

See [AL_PA](#).

bandwidth

A measure of the volume of data that can be transmitted at a specific transmission rate. A 1Gbps or 2Gbps Fibre Channel port can transmit or receive at nominal rates of 1 or 2Gbps, depending on the device to which it is connected. This corresponds to actual bandwidth values of 106MB and 212MB, respectively.

basic input output system

See [BIOS](#).

BIOS

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

boot code

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

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

boot device

The device, usually 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.

CRC

Cyclic redundancy check. A scheme to check data that have been transmitted or stored and to detect errors. A CRC cannot correct errors.

cyclic redundancy check

See [CRC](#).

device

A computer subsystem, such as an adapter card, that 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.

Ethernet

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

fabric device management interface

See [FDMI](#).

FCoE

Fibre Channel over Ethernet. A 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.

FDMI

Fabric device management interface. A standard enabling the management of devices such as adapters through the fabric.

FEC

Forward error correction. Automatically detects and recovers from bit errors. FEC enhances transmission reliability and performance.

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

See [FCoE](#).

firmware

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

Flash

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

forward error correction

See [FEC](#).

Host Bus Adapter

An adapter that connects a host system (the computer) to other network and storage devices.

LED

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

light-emitting diode

See [LED](#).

logical unit number

See [LUN](#).

loopback

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

LUN

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

message signaled interrupts

See [MSI](#), [MSI-X](#).

MSI, MSI-X

Message signaled interrupts. One of two PCI-defined extensions to support message signaled interrupts (MSIs), in PCI 2.2 and later and PCI Express. 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 2,048 and gives each interrupt separate data and address registers. Optional features in MSI (64-bit addressing and interrupt masking) are mandatory with MSI-X.

multiboot

The act of installing multiple operating systems on a computer, and being able to choose which one to boot when starting the computer. Multibooting may require a custom boot loader.

N_Port

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

node port

See [N_Port](#).

N_Port ID virtualization

See [NPIV](#).

non-volatile random access memory

See [NVRAM](#).

NPIV

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

NVRAM

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

PCI Express (PCIe)

A third-generation I/O standard that allows increased network throughput 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—a Fibre Channel port that supports point-to-point topology.
- NL_Port—a Fibre Channel port that supports loop topology.
- F_Port—a port in a fabric where an N_Port can attach.
- FL_Port—a port in a fabric where an NL_Port can attach.

port instance

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

QLflash

The executable file `QLflash.exe` is used by the `UPDATE.BAT` file to update the adapter [multiboot](#) code. The application `QLflash.exe` may be used to read, write, or verify either the multiboot image or the [NVRAM](#) on the adapter.

QoS

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

quality of service

See [QoS](#).

RAID

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

RDP

Read diagnostic parameters. A function that provides diagnostic information about a fiber optic cable and its connection.

Read diagnostic parameters

See [RDP](#).

redundant array of independent disks

See [RAID](#).

SAN

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

SCSI

Small computer systems interface. 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.

small computer systems interface

See [SCSI](#).

storage area network

See [SAN](#).

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.

target binding

The process in which the adapter driver binds a target ID using a target's world-wide port name ([WWPN](#)) or port ID. This binding enables the target ID to always connect to the WWPN or port ID across reboots regardless of [SAN](#) reconfiguration.

UEFI

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

unified extensible firmware interface

See [UEFI](#).

USCM

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

USCM Profiles

Settings to control the level of I/O throttling undertaken by the FC HBA in response to SAN congestion between a server and a target.

Virtual Lanes (VLs)

Marvell technology that enables a physical Fibre Channel link to be carved into multiple independent lanes, each with their own buffer-to-buffer credits. Virtual lanes work in conjunction with Cisco Extended Receiver Ready (ER_RDY) and Virtual Links. When there is peer congestion, the traffic for a designated slow device can be moved to a slow virtual lane without affecting the traffic to other devices.

virtual machine

See [VM](#).

virtual port

See [vPort](#).

VM

Virtual machine. A software implementation of a machine (computer) that executes programs like a real machine.

vPort

Virtual port. A port number or service name associated with one or more virtual servers. A virtual port number should be the same TCP or UDP port number to which client programs expect to connect.

worldwide node name

See [WWNN](#).

worldwide port name

See [WWPN](#).

WWNN

Worldwide node name. A unique 64-bit address assigned to a device.

WWPN

Worldwide port name. A unique 64-bit address assigned to each port on a device. One WWNN may contain multiple WWPN addresses.



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

Copyright © 2026 Marvell or its affiliates. All rights reserved. The MARVELL mark and M logo are registered and/or common law trademarks of Marvell or its affiliates in the US and/or other countries. Visit www.marvell.com for a complete list of Marvell trademarks. All third-party names, certification marks, trademarks (including logos and product names) and brands are the property of their respective owners. Use of such marks is for identification purposes only and does not imply endorsement or affiliation.

Doc. No. MA2854601-00 Rev. 10 Revised: May 29, 2026