

# Dell EMC Host Connectivity Guide for Windows

P/N 300-000-603

REV 61

MAY 2020

Copyright © 2015-2020 Dell Inc. or its subsidiaries. All rights reserved.

Dell believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS-IS.” DELL MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. USE, COPYING, AND DISTRIBUTION OF ANY DELL SOFTWARE DESCRIBED IN THIS PUBLICATION REQUIRES AN APPLICABLE SOFTWARE LICENSE.

Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners. Published in the USA.

Dell EMC  
Hopkinton, Massachusetts 01748-9103  
1-508-435-1000 In North America 1-866-464-7381  
[www.DellEMC.com](http://www.DellEMC.com)

# CONTENTS

	<b>PREFACE</b>	<b>5</b>
<b>Chapter 1</b>	<b>Introduction to Windows Infrastructure</b>	<b>7</b>
	Overview.....	8
	Terminology.....	8
	Windows functions and utilities.....	8
	Supported network connectivity in Windows environment.....	8
	Fibre Channel.....	9
	iSCSI.....	9
	Supported hardware.....	9
<b>Chapter 2</b>	<b>Connectivity</b>	<b>11</b>
	Fibre Channel connectivity.....	12
	Introduction.....	12
	Configuring HBAs for a Windows host.....	12
	iSCSI connectivity.....	17
	Introduction.....	17
	iSCSI discovery.....	18
	iSCSI Initiator cleanup.....	19
	NVMeoFC.....	20
	Booting from SAN.....	20
	Introduction.....	20
	Supported environments.....	20
	Limitations and guidelines.....	21
	Preparing host connectivity.....	21
	Configuring a SAN boot for an FC-attached host.....	22
	Direct-attached storage.....	25
	Network-attached storage.....	25
<b>Chapter 3</b>	<b>Managing Storage and Disk Paths in a Windows Environment</b>	<b>27</b>
	Introduction.....	28
	Dell EMC PowerPath for Windows.....	28
	PowerPath for Windows.....	28
	PowerPath verification and problem determination.....	30
	Microsoft Native MPIO.....	35
	MPIO with Windows Server 2012, 2012 R2, 2016, and 2019.....	35
	Veritas volume management software.....	38
<b>Chapter 4</b>	<b>Host Connectivity with Dell EMC Products</b>	<b>39</b>
	PowerMax/VMAX All Flash.....	40
	Dell EMC PowerMax.....	40
	Dell EMC VMAX All Flash.....	40
	PowerMax/VMAX All Flash/VMAX3 - Device types .....	41
	Local replication services .....	41
	Remote Replication Services .....	42
	Non-Disruptive Migration services .....	42
	PowerMax/VMAX All Flash-Storage provisioning .....	43


	Dell EMC PowerMax/VMAX All Flash/VMAX3 - Director bit settings ....	
	43	
Midrange Storage.....	44	
Unity storage.....	44	
Storage configuration.....	44	
Features.....	47	
Application considerations.....	50	
PowerStore.....	55	
VPLEX.....	60	
Overview.....	60	
Documentation.....	60	
Prerequisites.....	60	
Provisioning and exporting storage.....	61	
Storage volumes.....	63	
System volumes.....	64	
Required storage system setup.....	65	
Required VMAX series FA bit settings.....	65	
Supported storage arrays.....	66	
Initiator settings on back-end arrays.....	66	
Host connectivity.....	66	
Configuring Fibre Channel HBAs.....	66	
Windows Failover Clustering with VPLEX.....	72	
Setting up quorum on a Windows 2016/2019 Failover Cluster for VPLEX Metro clusters.....	72	
Setting up quorum on a Windows 2012 / 2012 R2 Failover Cluster for VPLEX Metro clusters.....	77	
Configuring quorum on Windows 2008/2008 R2 Failover Cluster for VPLEX Metro clusters.....	80	
XtremIO.....	84	
Best practices for zoning and subnetting.....	84	
Host Configuration for Microsoft Windows.....	91	
 <b>Chapter 5</b>	 <b>Operating System Features</b>	 <b>111</b>
General Windows information.....	112	
Booting Windows from external storage.....	112	
Boot-from-SAN.....	112	
Microsoft Windows failover clustering.....	114	
Operating system characteristics.....	115	
Hyper-V.....	116	



# PREFACE

As part of an effort to improve its product lines, Dell EMC periodically releases revisions of its software and hardware. Therefore, some functions described in this document might not be supported by all versions of the software or hardware currently in use. The product release notes provide the most up-to-date information on product features.

Contact your Dell EMC technical support professional if a product does not function properly or does not function as described in this document.

 **Note:** This document was accurate at publication time. Go to Dell EMC Online Support (<https://support.emc.com>) to ensure that you are using the latest version of this document.

## Purpose

This guide describes the features and setup procedures for Windows Server 2019, Windows Server 2016, Windows Server 2012 R2, and Windows Server 2012 host interfaces to Dell EMC storage arrays over Fibre Channel or iSCSI.

## Audience


This guide is intended for use by storage administrators, system programmers, or operators who are involved in acquiring, managing, or operating Dell EMC PowerMax, Dell EMC VMAX All Flash Family, Dell EMC VMAX3 Family, Dell EMC Unity Family, Dell EMC Unified VNX series, Dell EMC XtremIO, Dell EMC PowerStore, Dell EMC VPLEX, and host devices, and Windows Server 2019, Windows Server 2016, Windows Server 2012 R2, and Windows Server 2012.


## Related documentation


For the documentation referred in this guide, go to [Dell EMC Online Support](#).

## Special notice conventions used in this document

Dell EMC uses the following conventions for special notices:

 **WARNING** Indicates a hazardous situation which, if not avoided, could result in death or serious injury.

 **CAUTION** Indicates a hazardous situation which, if not avoided, could result in minor or moderate injury.

 **Note:** Presents information that is important, but not hazard-related.

## Typographical conventions

Dell EMC uses the following type style conventions in this document:

<b>Bold</b>	Used for names of interface elements, such as names of windows, dialog boxes, buttons, fields, tab names, key names, and menu paths (what the user specifically selects or clicks)
<i>Italic</i>	Used for full titles of publications referenced in text
Monospace	Used for: <ul style="list-style-type: none"><li>• System code</li><li>• System output, such as an error message or script</li><li>• Pathnames, filenames, prompts, and syntax</li></ul>

	<ul style="list-style-type: none"> <li>• Commands and options</li> </ul>
<i>Monospace italic</i>	Used for variables
<b>Monospace bold</b>	Used for user input
[ ]	Square brackets enclose optional values
	Vertical bar indicates alternate selections - the bar means “or”
{ }	Braces enclose content that the user must specify, such as x or y or z
...	Ellipses indicate nonessential information omitted from the example

## Where to get help

Dell EMC support, product, and licensing information can be obtained as follows:

### Product information

For documentation, release notes, software updates, or information about Dell EMC products, go to Dell EMC Online Support at <https://www.dell.com/support/>.

### Technical support

Go to Dell EMC Online Support and click Service Center. You will see several options for contacting Dell EMC Technical Support. Note that to open a service request, you must have a valid support agreement. Contact your Dell EMC sales representative for details about obtaining a valid support agreement or with questions about your account.

## Your comments

Your suggestions will help us continue to improve the accuracy, organization, and overall quality of the user publications. Send your opinions of this document to [techpubcomments@dell.com](mailto:techpubcomments@dell.com).

# CHAPTER 1

## Introduction to Windows Infrastructure

This chapter contains the following topics:

• <a href="#">Overview</a> .....	8
• <a href="#">Terminology</a> .....	8
• <a href="#">Windows functions and utilities</a> .....	8
• <a href="#">Supported network connectivity in Windows environment</a> .....	8
• <a href="#">Supported hardware</a> .....	9

## Overview

The Microsoft Windows Server family consists of the following versions including all Datacenter, Core, Essentials, and Enterprise:

- 2012/2012 R2
- 2016
- 2019

For information about support dates and service packs for your Windows Server version, see [Microsoft Lifecycle Policy](#).

The following chapters are designed to help a storage administrator or storage architect build a Windows Server infrastructure with Dell EMC storage.

## Terminology

The following table lists terms that are used in this guide.

Term	Definition
Array	Dell EMC storage, including Fibre Channel (FC), FC over Ethernet (FCoE), and Internet SCSI (iSCSI); external storage; and software-defined data center (SDDC) hyperconverged infrastructure (HCI) storage.
Free space	An unused and unformatted portion of a hard disk that can be partitioned or sub-partitioned.
Partition	A portion of a physical hard disk that functions as a physically separate unit.
Volume	A partition or collection of partitions that have been formatted for use by a file system. A volume is assigned with a drive letter.

## Windows functions and utilities

This section describes some Windows functions and utilities that you can use to define and manage Dell EMC VMAX/PowerMax series, Unity series, XtremIO series, and SDDC systems such as VPLEX and VxFLEX. The use of these functions and utilities is optional.

- **Disk Manager** - Graphical tool for managing disks; for example, partitioning, creating, and deleting volumes.
- **Registry Editor** - Graphical tool for displaying detailed hardware and software configuration information. This is normally not a part of the Administrative Tools group. The registry editor, Regedit.exe and Regedt32.exe, is in the Windows \system32 subdirectory.
- **Event Viewer** - Graphical tool for viewing system or application errors.

## Supported network connectivity in Windows environment

This section provides Fibre Channel and iSCSI support information specific to the Windows environment.

## Fibre Channel

Microsoft allows the host to connect to the external storage through a Fibre Channel (FC) host bus adapter (HBA) that is certified by Microsoft (See [Windows Server Catalog](#)) and running with Windows Server 2012, Windows Server 2012 R2, Windows Server 2016, and Windows Server 2019.

### Terminology

Term	Definition
Storage Area Network (SAN)	SAN is the most commonly used enterprise storage networking architecture for business-critical applications that need to deliver high throughput and low latency. A rapidly growing portion of SAN deployments uses all-flash storage to gain high performance, consistent low latency, and lower total cost when compared to spinning disks. With centralized shared storage, SANs enable organizations to apply consistent methodologies and tools for security, data protection, and disaster recovery.
World Wide Port Name (WWPN)	WWPN is a World Wide Name that is assigned to a port in a Fibre Channel fabric. It is a unique identifier that is used in SANs and performs a function that is equivalent to the MAC address in the Ethernet protocol.

## iSCSI

Microsoft iSCSI Initiator enables you connect a host computer that is running Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2, Windows Server 2016, and Windows Server 2019 to an external iSCSI-based storage array through an Ethernet network adapter. You can use Microsoft iSCSI Initiator in your existing network infrastructure to enable block-based SANs. SANs provide iSCSI target functionality without requiring an investment in additional hardware.

### Terminology

Term	Definition
Challenge Handshake Access Protocol (CHAP)	CHAP is an authentication method that is used during the iSCSI login in both the target discovery and the normal login.
iSCSI Network Portal	iSCSI Network Portal is the host NIC IP address that is used for the iSCSI driver to create a session with the storage.

## Supported hardware

See the [Dell EMC Simple Support Matrix](#) or contact your Dell EMC representative for the latest information on qualified hosts, host bus adapters (HBAs), and connectivity equipment.

Dell EMC recommends that you do not mix HBAs from different vendors in the same host.

The following Dell EMC storage products are Microsoft-certified and can be used in the Windows environment:

- **Unity series**

- **FC** - Unity 300 / Unity 400 / Unity 500 / Unity 600 / Unity 300F / Unity 400F / Unity 500F / Unity 600F / Unity 350F / Unity 450F / Unity 550F / Unity 650F/Unity XT 380 / Unity XT 480 /Unity XT 680 /Unity XT 880 /Unity XT 380F /Unity XT 480F /Unity XT 680F /Unity XT 880F
- **iSCSI**: Unity 300 iSCSI Target / Unity 400 iSCSI Target / Unity 500 iSCSI Target / Unity 600 iSCSI Target / Unity 300F iSCSI Target / Unity 400F iSCSI Target / Unity 500F iSCSI Target / Unity 600F iSCSI Target / Unity 350F iSCSI Target / Unity 450F iSCSI Target / Unity 550F iSCSI Target / Unity 650F iSCSI Target / Unity XT 380 iSCSI Target / Unity XT 480 iSCSI Target / Unity XT 680 iSCSI Target / Unity XT 880 iSCSI Target / Unity XT 380F iSCSI Target / Unity XT 480F iSCSI Target / Unity XT 680F iSCSI Target / Unity XT 880F iSCSI Target / UnityVSA iSCSI Target
- **XtremIO series**
  - **FC** - XtremIO X2-R / XtremIO X2-S / XtremIO X2-T
  - **iSCSI** - XtremIO X2-R iSCSI Target / XtremIO X2-S iSCSI Target / XtremIO X2-T iSCSI Target
- **VMAX/PowerMax series**
  - **FC** - VMAX 100K / VMAX 200K / VMAX 400K / VMAX250F / VMAX250FX / VMAX450F / VMAX450FX / VMAX850F / VMAX850FX / VMAX950F / VMAX950FX / PowerMax 2000 / PowerMax 8000
  - **iSCSI** - VMAX 100K iSCSI Target / VMAX 200K iSCSI Target / VMAX 400K iSCSI Target / VMAX250F iSCSI Target / VMAX250FX iSCSI Target / VMAX450F iSCSI Target / VMAX450FX iSCSI Target / VMAX850F iSCSI Target / VMAX850FX iSCSI Target / VMAX950F iSCSI Target / VMAX950FX iSCSI Target / PowerMax 2000 iSCSI Target / PowerMax 8000 iSCSI Target
- **PowerStore series**
  - **FC** - PowerStore-T(san) / PowerStore-X(HCI)
  - **iSCSI** - PowerStore-T(san) Target / PowerStore-X(HCI) Target
- **VPLEX series**
  - **FC** - VPLEX VS2 / VPLEX VS6

# CHAPTER 2

## Connectivity

This chapter contains the following topics:

• <a href="#">Fibre Channel connectivity</a> .....	12
• <a href="#">iSCSI connectivity</a> .....	17
• <a href="#">Booting from SAN</a> .....	20
• <a href="#">Direct-attached storage</a> .....	25
• <a href="#">Network-attached storage</a> .....	25

# Fibre Channel connectivity

This section provides details about Fibre Channel (FC) host bus adapter (HBA) configuration.

## Introduction

FC has some of the benefits of both channels and networks. A FC fabric is a switched network, providing generic, low-level services onto which host channel architectures and network architectures can be mapped. Networking and I/O protocols (such as SCSI commands) are mapped to FC constructs and then encapsulated and transported within FC frames. This process enables high-speed transfer of multiple protocols over the same physical interface.

As with direct-attach SCSI, FC provides block-level access to the devices that allows the host system to identify the device as a native device. The true power of native device identification is the ability to use all the current applications (for example, backup software, volume management, and raw disk management) without modification.

FC is a technology for transmitting data between computer devices at data rates of up to 32 GB/s at this time. FC is flexible, and devices can be as far as 10 kilometers (about 6 miles) apart if optical fiber is used as the physical medium.

FC supports connectivity over fiber optic cabling or copper wiring. FC devices using fiber optic cabling use two unidirectional fiber optic cables for each connection. One fiber optic cable is used for transmitting and the other for receiving.


## Configuring HBAs for a Windows host

This section describes the procedure for installing a Dell EMC approved Broadcom Emulex, and Marvell QLogic host bus adapter (HBA) into a Windows host environment and configuring the host for connection to a Dell EMC storage array over Fibre Channel.

This section addresses the following topics:

- Prerequisites for first-time installation
- Broadcom Emulex Fibre Channel (FC) HBA
- Marvell QLogic FC HBA

### Prerequisites for first-time installation

 **Note:** Dell EMC recommends that you do not mix different types of Fibre Channel (FC) host bus adapters (HBAs) (including different types from the same vendor) in a server.

- See the vendor's FC HBA product documentation to properly install HBA in your server.
- See the vendor's product documentation to verify and update HBA firmware and boot the BIOS to a Dell EMC-qualified versions.
- The latest Dell EMC-approved HBA drivers and software are available for download at the following websites:
  - <https://www.broadcom.com>
  - <http://www.marvell.com>

Dell EMC HBA installation and configurations guides are available at the Dell EMC download pages on these websites.



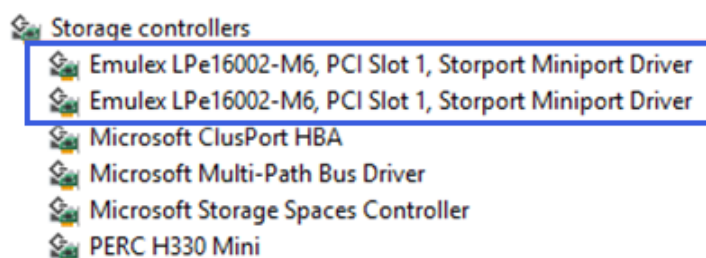
## Emulex FC HBA

Using a Broadcom Emulex Fibre Channel (FC) host bus adapter (HBA) with the Windows operating system requires adapter driver software. The driver functions at a layer below the Windows SCSI driver to present FC devices to the operating system as standard SCSI devices.

Download the driver from the Dell EMC section of the Broadcom website. Follow the links to your adapter for the appropriate operating system and version.

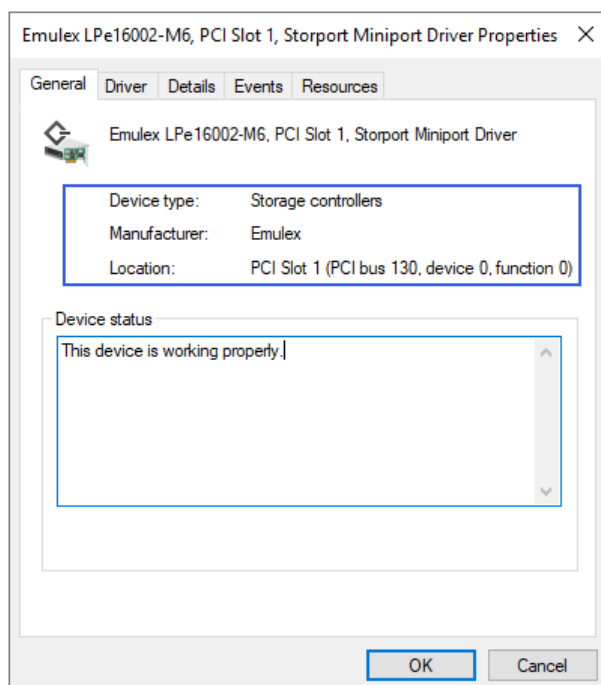
After the driver installation is complete, go to **Computer Management > Device Manager > Storage Controllers**. You can see the Emulex driver information as shown in the following illustration.

**Figure 1** Emulex driver information



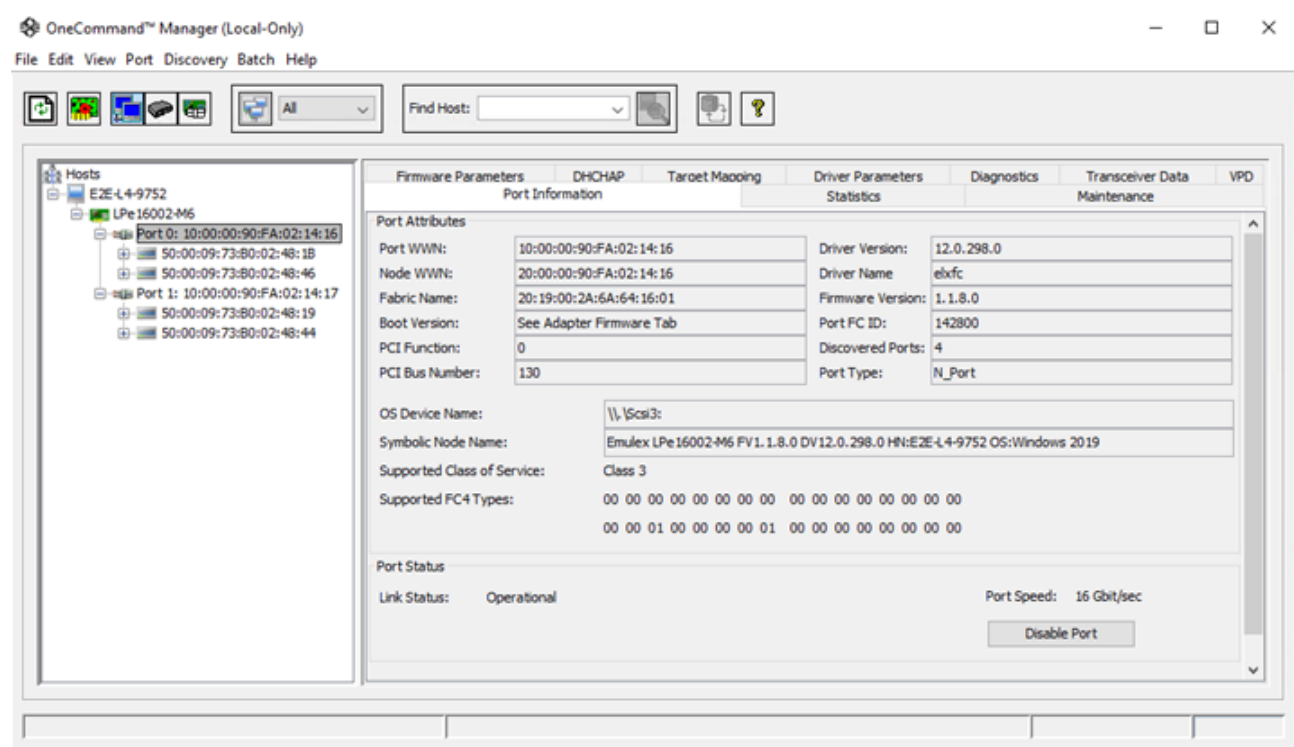
Right-click a port and review the driver properties, as shown in following illustration.

**Figure 2** Emulex driver properties window



Dell EMC recommends that you install the Emulex OneCommand Manager from the [Broadcom](#) website for better HBA management.

Figure 3 OneCommand Manager



You can also use the Dell EMCInquiry utility, INQ, as an alternate for HBA management.

**Figure 4** Viewing Emulex HBA details on INQ

[illegible]

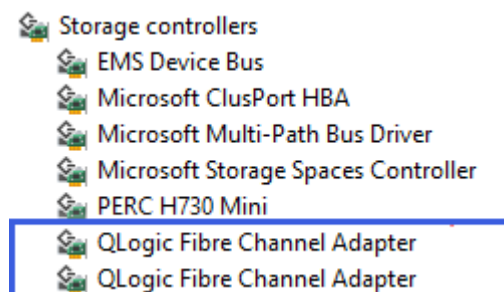
## Marvell QLogic FC HBA

Using the Qlogic Fibre Channel (FC) host bus adapter (HBA) with the Windows operating system requires adapter driver software. The driver functions at a layer below the Windows SCSI driver to present FC devices to the operating system as standard SCSI devices.

Download the driver from the Dell EMC section of the Qlogic website. Follow the links to your adapter for the appropriate operating system and version.

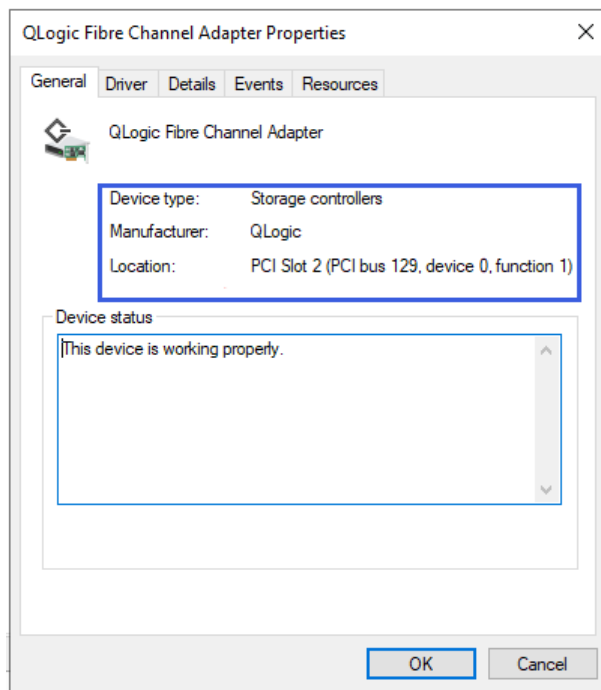
After the driver installation is complete, go to **Computer Management > Device Manager > Storage Controllers**. You can see the QLogic driver information as follows:

**Figure 5** Marvell QLogic driver information



Right-click a port and review the driver properties, as shown in the following figure.

**Figure 6** Marvell QLogic driver properties window



Dell EMC recommends that you install the QConvergeConsole from the [Marvell QLogic](#) website for better HBA management. The following figure shows the QConvergeConsole CLI.

**Figure 7** QConvergeConsole

```

QConvergeConsole CLI
QConvergeConsole
CLI - Version 2.3.0 (Build 11)

FC Adapter Information

1: FC Adapter Information
2: FC Port Information
3: FC VPD Information
4: FC Target/LUN Information
5: FC VFC Information
6: FC Flash Information

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

QConvergeConsole
CLI - Version 2.3.0 (Build 11)

Adapter Information

1: HBA Model: QLE2742 SN: RFD1821X77742
   Port 1 WWPN: 21-00-F4-E9-D4-S3-15-4C Online (FEC)
   Port 2 WWPN: 21-00-F4-E9-D4-S3-15-4D Online (FEC)

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

```

You can also use the Dell EMC Inquiry utility, INQ, as an alternative for HBA management.



architecture can be implemented between routers, host-to-switch, and storage array-to-storage array to provide asynchronous/synchronous data transfer.

iSCSI initiators come in two varieties:

- A software initiator is an operating system driver that is used with an Ethernet card. The iSCSI driver handles all requests by determining if the packet is an Ethernet packet that is then passed to the network stack, or if it is an iSCSI packet that will then have the SCSI packet stripped out and pasted to the SCSI stack. Using a software initiator requires more CPU and memory resources on the host.
- A hardware initiator is a host bus adapter (HBA) that offloads some or all the iSCSI packet processing, which saves CPU and memory resources.

These adapters reduce the load of the iSCSI stack in the operating system.

## iSCSI discovery

### About this task

**Note:** Microsoft iSCSI Initiator is native to Windows Server 2019, Windows Server 2016, Windows Server 2012 R2, and Windows Server 2012. On these operating systems, no installation steps are required. For more details, see the *Microsoft iSCSI Initiator Step-by-Step Guide* on [Microsoft TechNet](#). Before configuring the iSCSI initiator, ensure that you have identified the NIC and the target where it will connect.

**For example:** NIC1 and SPA-0 and SPB-0 are on one network subnet. NIC2 and SPA-1 and SPB-1 are on a different subnet. This example connects NIC1 to SPA-0 and SPB-0, and NIC2 to SPA-1 and SPB-1.

**Note:** NIC1 and NIC2 could also be on the same subnet, but Dell EMC does not recommend it except for PowerStore arrays, that support only one subnet for iSCSI connections.

To configure iSCSI Initiator, complete the following steps:

### Procedure

1. Log in to the server as Administrator and open Microsoft iSCSI Initiator through the **Control Panel** (showing all Control Panel items) or **Administrative Tools**.
2. Select **Server Manager > Tools > iSCSI Initiator**. In the **iSCSI Initiator Properties** dialog box, click the **Discovery** tab and click **Discover Portal**.

The **Discover Target Portal** dialog box appears.

3. Enter the IP address of the target storage and select **Advanced**.

The **Advanced Setting** dialog box appears.

4. In the **Local adapter** field, select **Microsoft iSCSI Initiator**.
5. Select the IP address of the NIC to be used.
6. Click **OK**, and then click **OK**.

The **iSCSI Initiator Properties** dialog box appears.

7. Click the **Targets** tab.
8. Select the first target iSCSI Qualified Name (IQN) and click **Connect**.

The **Connect To Target** dialog box appears.

9. If you are using Dell EMC PowerPath or Microsoft Multipath I/O (MPIO):
  - a. Select **Enable multi-path** and click **Advance**.

The **Advanced Settings** dialog box appears.

- b. In the **Local adapter** field, select **Microsoft iSCSI Initiator**.
- c. In the **Initiator IP** field, select the correct NIC IP address.
- d. In the **Target portal IP** field, select the IP address.
- e. Click **OK**.
10. Click **OK**.
11. Connect each of the other three targets in the list by following the preceding steps.
12. To see all the targets that are connected, go to the **iSCSI Initiator Properties** dialog box, and click the **Favorite Targets** tab.
13. If Unisphere/Navisphere Agent is not registered, manually register the NIC in Unisphere/Navisphere Manager.

If the host has Unisphere/Navisphere Agent installed, you can see it logged in and registered in Unisphere/Navisphere Manager.

14. Using Unisphere/Navisphere Manager, place the host in a Storage Group that has LUNs in it, and then go back to the host and do a device manager scan for hardware changes.

After a few minutes, the disk devices are displayed in the PowerPath UI or in Disk Management or both.

**Note:** The PowerPath UI shows only one adapter, even if you are using multiple NICs. The adapter that is shown in the PowerPath UI represents the Microsoft iSCSI Initiator, not the NICs that are installed in your system.

## iSCSI Initiator cleanup

### About this task


**Note:** If you are running Windows Failover Cluster, ensure that the host does not own any disk resources. Move resources to another node in the cluster or take the disk resources offline. Similarly, any LUNs being used on a standalone Windows host must be offline. Use Disk Management to take the disks offline.

To clean up iSCSI Initiator, complete the following steps.

### Procedure

1. Log in to the server as Administrator and open Microsoft iSCSI Initiator through **Control Panel** (showing all Control Panel items) or **Administrative Tools**.
2. Click the **Discovery** tab, select one of the addresses in the **Target portals** field, and click **Remove**.
3. In the **Remove Target Portal** dialog box, click **OK**.
4. Remove all the other target portals.
5. For each volume that you want to remove, in the **iSCSI Initiator Properties** dialog box, click the **Volumes and Devices** tab, select the volume in the **Volume List** field, and click **Remove**.
6. For each target that you want to remove, in the **iSCSI Initiator Properties** dialog box, click the **Favorite Targets** tab, select the target from the **Favorite targets** field, and click **Remove**.
7. In the **iSCSI Initiator Properties** dialog box, click the **Targets**, select one of the targets in the **Discovered targets** field, and click **Disconnect**.
8. In the warning message box, click **Yes**.

9. Repeat step 7 and step 8 for each of the targets to be disconnected.
10. If you are running PowerPath software, the PowerPath UI shows the status of all the devices as **dead**. To clean and remove these, complete the following steps:
  - a. Open PowerShell using the **Run as Administrator** option, as shown in the following figure:



- b. Type `powermt check` and, when prompted to remove dead device, select **a** for ALL.
11. Review the **Discovery** tab and ensure that no targets are connected.
12. Review each of the iSCSI initiator tabs and ensure that they are empty.

## NVMeoFC

NVMeoFC is currently not supported with Windows.

## Booting from SAN

### Introduction

All Dell EMC storage arrays support installing and booting Windows from a SAN environment. The Windows host operating system can reside on an external device that is managed by the Microsoft Multipath I/O (MPIO) utility or PowerPath software. This section describes the configuration process and how to prevent possible issues.

### Supported environments

Dell EMC storage such as PowerMax and VMAX series, VNX, Unity, PowerStore, VPlex, and XtremIO arrays supports boot-from-SAN.

See the [Dell EMC Simple Support Matrix](#) for a list of operating system kernels that support booting from SAN storage.

Dell EMC supported Broadcom Emulex and Marvell QLogic host bus adapters (HBAs) can be used to boot from SAN. To boot from storage that is attached to the SAN environment, the HBA's boot



BIOS must be installed and enabled on the adapter. See the driver manuals and configuration guides on the Dell EMC section of the Broadcom Emulex and Marvell QLogic, websites as well as operating system, HBA, and server vendor documentation.

 **Note:** iSCSI booting from SAN is supported in limited configurations. See the [Dell EMC Simple Support Matrix](#) for information about supported environments.

## Limitations and guidelines

Boot configurations must not deviate from the following limitations:

- The Dell EMC storage device must have enough disk space to hold the Windows operating system.
- Space reclamation, available with Enginuity 5874 and later, is prohibited for use on PowerMax and VMAX Virtual Provisioning (thin) devices that are used for system reservation volumes.
- 4K devices are not supported for SAN boot with XtremIO arrays.

## Preparing host connectivity

### Guidelines


- Maintain the simplest connectivity configuration between the host server and SAN environment before installing the operating system. You can alter the configuration after installation.
- If multiple host bus adapters (HBAs) are attached to the host, ensure that the HBA that is connected to the lowest-numbered PCI slot is zoned to the array.
- Unzone all arrays from the host server, except the array where the boot device resides. Place only the boot device that is attached to the host server on the array where the boot device resides.
- Assign the boot LUN as host LUN ID 0. If the boot LUN has a Host ID other than 0, the HBA BIOS installation might fail and have no visibility to the boot LUN.
  - The boot LUN's Host ID on a VNX series or Dell EMC CLARiiON system can be forced to 0 by removing all the other LUNs from the storage group and adding back only the boot LUN.
  - For Symmetrix, the Symmetrix LUN base/offset skip adjustment (symmask set lunoffset) capability can be used to assign LUN 0 to a boot LUN if necessary.
  - For Unity series or VNXe series, you can modify the Host LUN ID directly in the Unisphere UI.
  - For PowerStore series, you can modify the Host LUN ID directly in the PowerStore Manager UI.
  - For the PowerMax and VMAX series, when you add devices to a storage group, you can use the following command:
 

```
switch -lun 0. (symaccess -sid xxx -type stor -name host_sg_name
add devs lun_id -lun 0)
```
- In XtremIO version 4.0.0 or later, volumes are numbered by default, starting from LUN ID 1. Do not manually adjust the LUN ID to 0, because doing so might lead to issues with some operating systems. In XtremIO 3.x and earlier versions, LUN IDs start at 0 and remain accessible when the XtremIO cluster is updated from 3.0.x to 4.x.
- You can add LUNs after the operating system installation is completed.

## Single-path and dual-path configuration examples


The following figure shows an example of single-path configuration.

**Figure 9** Single-path configuration



The following figure shows an example of dual-path configuration.

**Figure 10** Dual-path configuration



## Configuring a SAN boot for an FC-attached host

### Installing and configuring Fibre Channel HBA

After you make a connection between the host and boot LUN, you must enable and install the host bus adapter (HBA) BIOS to enable a boot from a SAN. When the host operating system resides on an external device, it is not loaded during a boot by the host system's hard disk controller. Only the HBA BIOS can obtain the operating system image. To facilitate visibility of an external boot device by the HBA, therefore, the HBA BIOS must be installed to register the external device as the boot source.

The major HBA vendors, Broadcom and Marvell QLogic have embedded HBA configuration utilities in their product BIOS that can be accessed by a shortcut key during server boot. For details, see the HBA vendor's site.

### Updating HBA BIOS and firmware

#### About this task

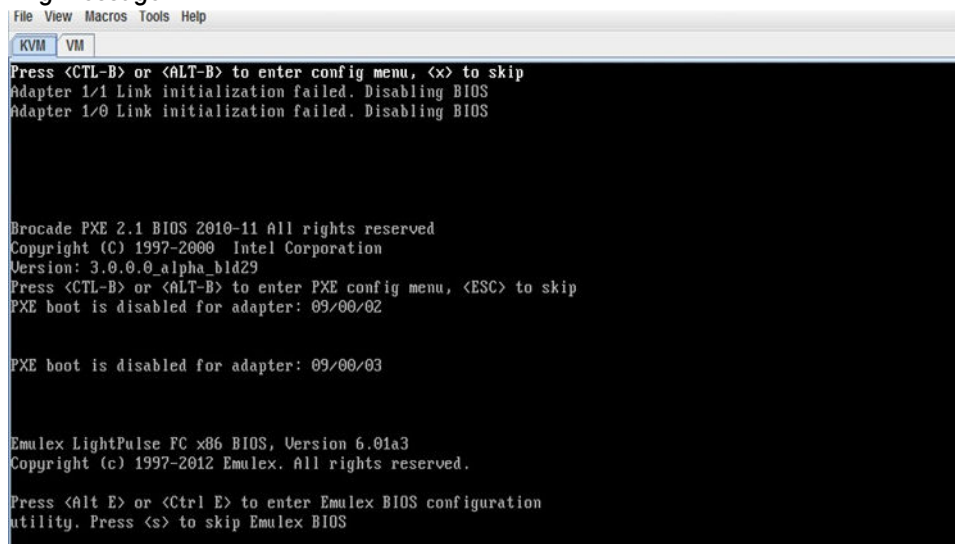
Before you configure a host bus adapter (HBA) for boot-from-SAN, verify that the HBA BIOS and firmware have been updated to the Dell EMC supported version. The Dell EMC supported version is available in the Dell EMC OEM section of the Broadcom and Marvell QLogic websites, and on Dell EMC Online Support.

The following procedure is an example of how to set up the Windows operating system with the Broadcom Emulex HBA and using single-path connection during installation. The procedure for

other vendor HBA and operating system versions might require slight modifications. For details, see the HBA and operating system vendor websites.

### Procedure

1. Boot the server, and press Alt+E to enter the Broadcom Emulex BIOS when you see the following message:



```
File View Macros Tools Help
KVM VM
Press <CTL-B> or <ALT-B> to enter config menu, <x> to skip
Adapter 1/1 Link initialization failed. Disabling BIOS
Adapter 1/0 Link initialization failed. Disabling BIOS

Broadcom PXE 2.1 BIOS 2010-11 All rights reserved
Copyright (C) 1997-2000 Intel Corporation
Version: 3.0.0.0_alpha_bld29
Press <CTL-B> or <ALT-B> to enter PXE config menu, <ESC> to skip
PXE boot is disabled for adapter: 09/00/02

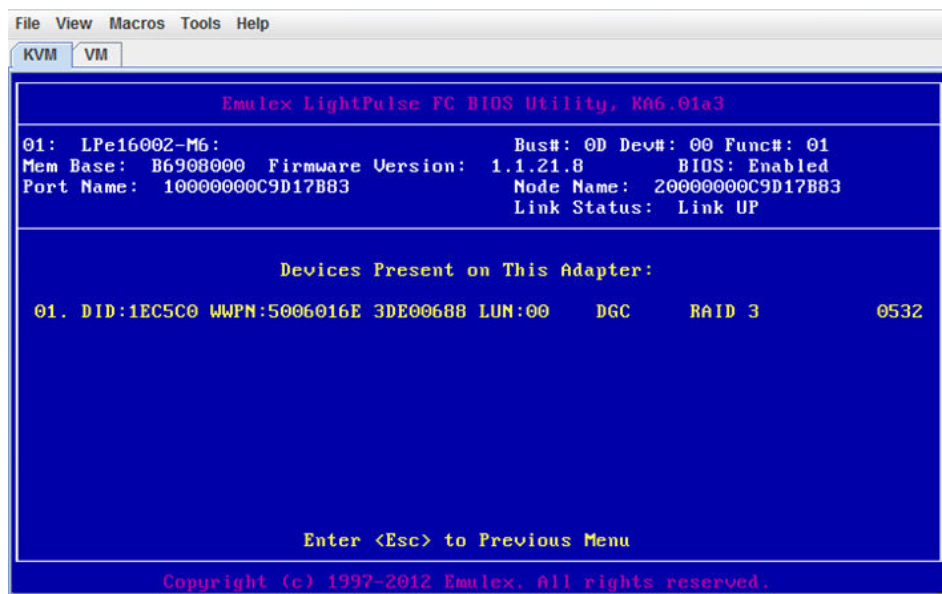
PXE boot is disabled for adapter: 09/00/03

Emulex LightPulse FC x86 BIOS, Version 6.01a3
Copyright (c) 1997-2012 Emulex. All rights reserved.

Press <Alt E> or <Ctrl E> to enter Emulex BIOS configuration
utility. Press <s> to skip Emulex BIOS
```

2. Select the adapter port that you want to configure, which, in a single-path configuration, is the HBA port zoned to storage.
3. If the link status is **Link UP**, select **Enable/Disable Boot from SAN** to enable the HBA BIOS.
4. To enable the boot BIOS on the adapter port, select **Enable**.
5. After the BIOS is enabled, select **Scan for Target Devices**.

The utility lists the attached LUNs. The example in following figure shows the attached boot LUN, which is a VNX RAID3 LUN.



```
File View Macros Tools Help
KVM VM
Emulex LightPulse FC BIOS Utility, KA6.01a3

01: LPe16002-M6: Bus#: 0D Dev#: 00 Func#: 01
Mem Base: B6908000 Firmware Version: 1.1.21.8 BIOS: Enabled
Port Name: 10000000C9D17B83 Node Name: 20000000C9D17B83
Link Status: Link UP

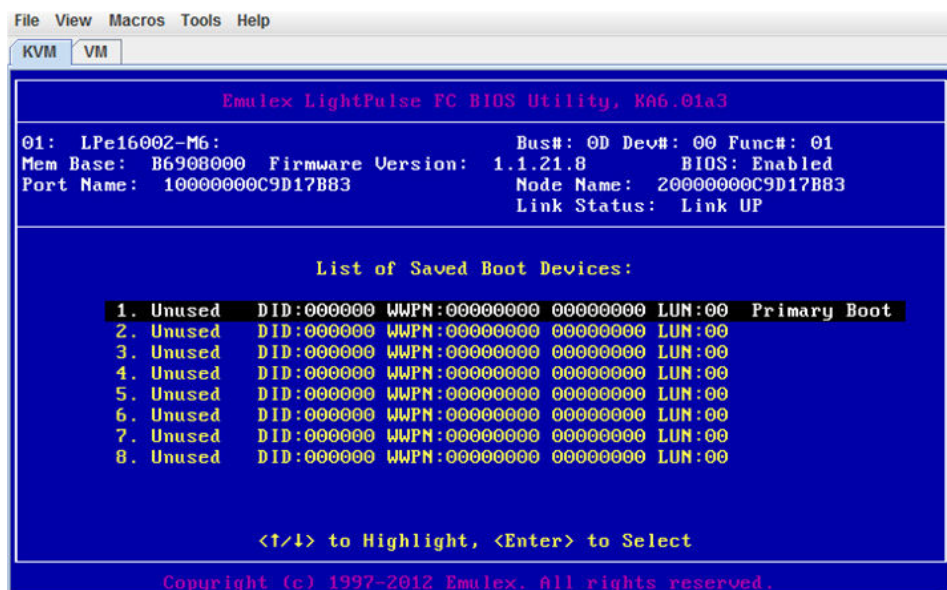
Devices Present on This Adapter:

01. DID:1EC5C0 WWP:5006016E 3DE00688 LUN:00 DGC RAID 3 0532

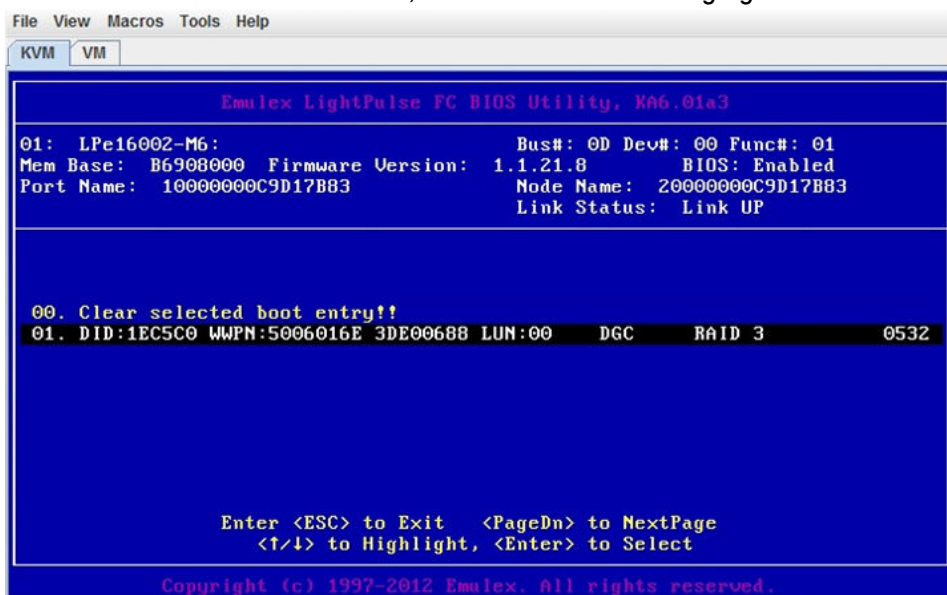
Enter <Esc> to Previous Menu

Copyright (c) 1997-2012 Emulex. All rights reserved.
```

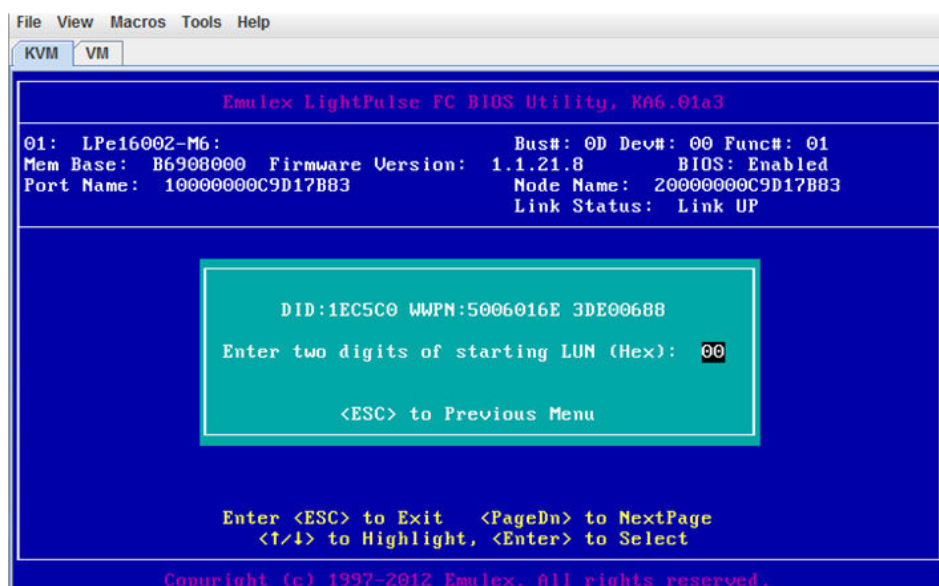
6. When the boot LUN is visible to the HBA, select **Configure Boot Devices**.
7. Review the list of boot devices and select the required LUN as the primary boot device.



You can view the device for boot details, as shown in the following figure.

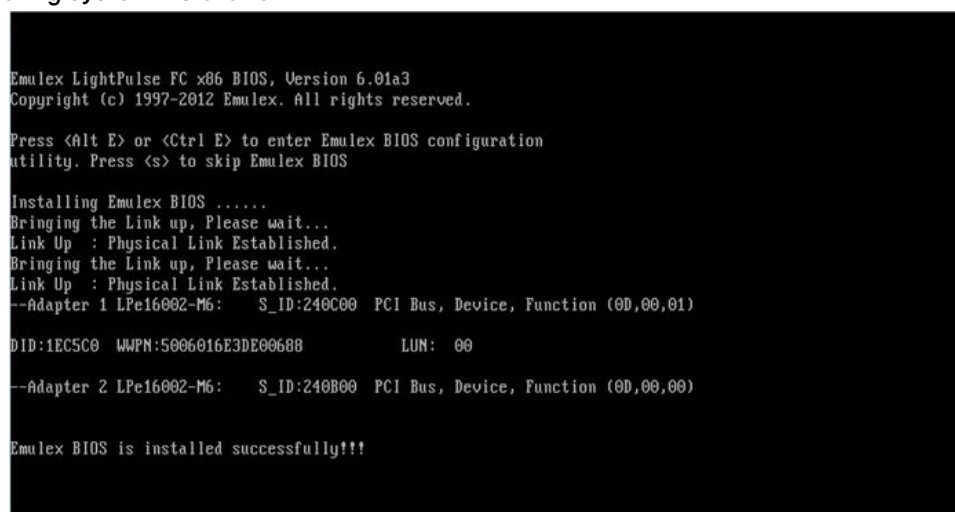


8. Set the boot LUN ID as 00.



9. Save the changes and reboot the system.

If configuration is completed successfully, the information for the selected boot device is displayed after the welcome banner when the server boots, with message that tells you the BIOS is successfully installed, as shown in the following figure. The server is now ready for operating system installation.



## Direct-attached storage

Direct-attached Storage (DAS) is a host system directly attached to a Dell EMC storage array ports, it does not have any network switches in-between the host and array.

## Network-attached storage

Network-attached storage (NAS) is a file-level data-storage server that is connected to a computer network that provides data access to a heterogeneous group of clients. NAS is specialized for serving files by its hardware, software, or configuration.

Some Dell EMC storage arrays such as VNX, Unity, PowerStore, and XtremIO arrays support Network-attached Storage. For more details, see *Dell EMC Support Matrix*, available on [Dell EMC E-Lab Navigator](#).

# CHAPTER 3

## Managing Storage and Disk Paths in a Windows Environment

This chapter contains the following topics:

• <a href="#">Introduction</a> .....	28
• <a href="#">Dell EMC PowerPath for Windows</a> .....	28
• <a href="#">Microsoft Native MPIO</a> .....	35
• <a href="#">Veritas volume management software</a> .....	38

## Introduction

Dell EMC supports various mechanisms to address multiple paths to a device. Having redundant access paths to a storage environment is an essential aspect in any storage topology. Online array code upgrades (nondisruptive upgrades, or NDU), online configuration changes, or any disturbances in the topology are best handled by a host when multiple paths are available, and when path management software is installed and configured on the host.

The advantages of path management software include:

- **Path failover and path monitoring**-Periodically assess the health of the host storage connectivity and routing over a preconfigured alternate path in case of a path failure, a component failure, or both.
- **Load balancing**-Share the I/O load across multiple channels to improve performance.
- **Device management**- Manage multiple native devices which are instances of a single device, and in active-passive array environments, route I/O to an active device.
- **Device persistence**-Achieve persistence of device names when the SCSI device is rescanned upon reboot.

## Dell EMC PowerPath for Windows

PowerPath for Windows software is available in two packages: PowerPath for Windows and PowerPath iSCSI for Windows.

### PowerPath for Windows

PowerPath for Windows supports both Fibre Channel (FC) and iSCSI environments. It is digitally certified by Microsoft for FC environments only. PowerPath for Windows supports failover path management and load-balancing up to 32 paths in heterogeneous storage environments. PowerPath for Windows is not currently supported by Microsoft for iSCSI implementations, although it is supported by Dell EMC for Dell EMC iSCSI storage systems.

### PowerPath iSCSI for Windows

PowerPath iSCSI for Windows supports Dell EMC PowerMax and VMAX series, VNX, Unity, PowerStore, and XtremIO iSCSI storage systems. It is Microsoft digitally certified and is built on the Microsoft Multipath I/O (MPIO) framework. PowerPath iSCSI for Windows supports failover path management up to 8 paths in iSCSI storage environments.

## PowerPath for Windows

### PowerPath and MSCS


If you are installing PowerPath software and Microsoft Cluster Server (MSCS) for the first time, install PowerPath first, and then install MSCS. Installing PowerPath first will avoid any disrupt to cluster services later.

### Integrating PowerPath with an existing MSCS cluster

You can integrate PowerPath software with an existing MSCS cluster without shutting down the cluster, if there is close coordination between the nodes and the storage system. Each node in a cluster can own a distinct set of resources. Node A is the primary node for its resources and the failover node for Node B resources. Conversely, Node B is the primary node for its resources and the failover node for Node A resources.



After you install PowerPath software on the cluster, if you test node failover by disconnecting all cables for a LUN or by disrupting the path between the active host and the array, Windows will log event messages indicating a hardware or network failure and possible data loss. If the cluster is working correctly, it will fail over to a node with an active path and you can ignore the messages from the original node as logged in the event log.

 **Note:** Check the application that is generating I/O to see if any failures have occurred. If no failures have occurred, everything is working normally.

Installing PowerPath software in a clustered environment requires:

- Moving all resources to Node A
- Installing PowerPath software on Node B
- Configuring additional paths between the storage array and Node B
- Moving all resources to Node B
- Installing PowerPath software on Node A
- Configuring additional paths between the storage array and Node A
- Returning Node A resources to Node A

The following sections describe these tasks in detail.

## Moving all resources to Node A

### Procedure


1. To start the Failover Clustering Administrator utility, click **Start > Server Manager > Tools > Failover Cluster Manager**.
2. In the left pane of the window, select **All roles and disks** owned by Node B.
3. Migrate the selected roles and disks to Node A. Right click the role and select **MOVE** action.
4. To pause Node B, click **Node B** and click **Action > Pause**.

This keeps the node from participating in the cluster during PowerPath software installation.

## Installing PowerPath software on Node B


### Procedure

1. Install PowerPath software on Node B.
2. Shut down Node B.

 **Note:** If the cluster has more than two nodes, install PowerPath software on the other nodes.  
For example, In a four-node cluster, replace Node B with Nodes B, C, and D in step 4 of the previous section (Moving resources to Node A), and in steps 1 and step 2 of this section.

## Configuring additional paths between the storage array and Node B

### About this task

 **Note:** If a cluster has more than two nodes, configure additional paths between the storage system and the other nodes.  
For example, in a four-node cluster, replace Node B with Nodes B, C, and D in this procedure.

### Procedure

1. If necessary, reconfigure the storage system so that its logical devices are displayed on multiple ports.
2. If necessary, install additional HBAs on Node B.
3. Connect the cables for new paths between Node B and the storage system.
4. Turn on Node B.
5. To resume Node B, click **Node B > Action > Resume**.

## Moving all resources to Node B

### Procedure

1. In the left pane of the window, select all roles and disks owned by Node A.
2. To migrate the selected roles and disks to Node B, right click on one role and select **MOVE**.
3. To pause Node A, click **Node A** and click **Action > Pause**.

## Installing PowerPath software on Node A

### Procedure

1. Install PowerPath software on Node A.
2. Shut down Node A.

## Configuring additional paths between the storage array and Node A

### Procedure

1. If necessary, configure the storage system so that its logical devices are displayed on multiple ports.
2. If necessary, install additional HBAs on Node A.
3. Connect the cables for new paths between Node A and the storage system.
4. Turn on Node A.
5. To resume Node A, click **Node A** and click **Action > Resume**.

## Returning Node A resources to Node A

### Procedure

1. Using the MSCS Cluster Administrator utility, select all groups that were previously owned by Node A.
2. Click **Action > Move**.

## PowerPath verification and problem determination

This section assumes that PowerPath has been installed properly. See the appropriate *PowerPath Installation and Administration Guide* on Dell EMC Online Support for information about how to install PowerPath. This section is designed to help you to verify that PowerPath was installed correctly and to recognize some common failures.

In the taskbar, click the PowerPath icon—the circled icon that is shown in the following figure, to access the PowerPath UI, PowerPath Administrator.

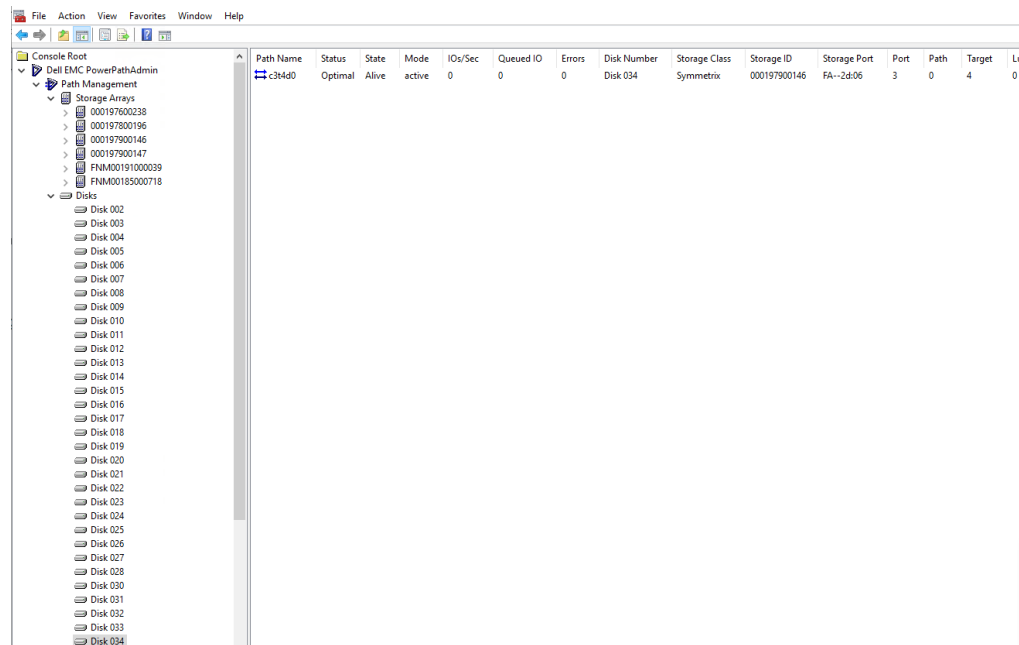
**Figure 11** PowerPath icon in taskbar

The following figure shows taskbar icons and the PowerPath status that each represents.

**Figure 12** PowerPath taskbar icons and status

Taskbar Icon	Status
	<b>Optimal</b> — All configured paths are operating properly.
	<b>Degraded</b> — Some but not all paths have failed.
	<b>All Failed</b> — All configured paths have failed.
	<b>Server Not Responding</b> — Cannot communicate with Monitor service. This occurs if the PowerPath Administrator Suite encountered errors during initialization.

The following figure shows how the PowerPath Administrator would look if PowerPath is installed correctly. In this case, one path is zoned between the host bus adapter (HBA) and one port on the storage device.

**Figure 13** One path

The following figure shows how the PowerPath Administrator would look when multiple paths are zoned to your storage device.

**Figure 14** Multiple paths







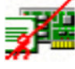

The screenshot shows the Dell EMC PowerPath Administrator console. The left pane displays a tree view with 'Console Root' expanded, showing 'Dell EMC PowerPathAdmin', 'Path Management', 'Storage Arrays', and 'Disks'. The right pane shows a table of paths to a disk device.

Path Name	Status	State	Mode	IOs/Sec	Queued IO	Errors	Disk Number	Storage Class	Storage ID	Storage Group	Storage Port	Port	Path	Target	Lun
c3054	Optimal	Alive	active	0	0	0	Disk 005	PowerStore	FNM0018500...	-	A0-NA-f3	3	0	0	4
c3654	Optimal	Alive	active	0	0	0	Disk 005	PowerStore	FNM0018500...	-	A0-NB-f3	3	0	6	4
c40d1	Optimal	Alive	active	0	0	0	Disk 005	PowerStore	FNM0018500...	-	A0-NB-f2	4	0	0	1
c4k3d1	Optimal	Alive	active	0	0	0	Disk 005	PowerStore	FNM0018500...	-	A0-NA-f2	4	0	3	1















## Problem determination

Use of PowerPath Administrator can help you determine what has caused a loss of connectivity to the storage device. The PowerPath Administrator UI shows array ports that are offline, defective host bus adapters (HBAs), and broken paths in several ways. The following table shows the known possible failure states. Referencing this table can greatly reduce problem determination time.

**Table 1** Possible failure states


Object	Degraded icon	Failed icon	Unlicensed and degraded icon	Unlicensed and failed icon
PowerPath device	 <p>One or more (but not all) paths to the disk device have failed.</p>	 <p>All paths to the disk device have failed. The disk is not available.</p>	 <p>PowerPath Administrator is unlicensed. One or more (but not all) paths to the disk device have failed.</p>	 <p>PowerPath Administrator is unlicensed. All paths to the disk device have failed. This disk is not available.</p>
Adapter	 <p>One or more (but not all) adapters have either failed or degraded. The status is displayed in the Adapters folder in the scope pane.</p> <p>One or more (but not all) paths have failed on the adapter. The status is displayed in the individual adapter folder under Adapters in the scope pane or in the result pane when Adapters is selected from the scope pane.</p>	 <p>All paths on this adapter to this disk have failed.</p>	 <p>PowerPath Administrator is unlicensed.</p> <p>One or more (but not all) adapters have either failed or degraded. The status is displayed in the Adapters folder in the scope pane.</p> <p>One or more (but not all) paths have failed on the specific adapter. The status is displayed in the individual adapter folder under Adapters in the scope pane or in the result pane when Adapters is selected from the scope pane.</p>	 <p>PowerPath Administrator is unlicensed.</p> <p>All paths on this adapter to the disk have failed.</p>

**Table 1** Possible failure states (continued)

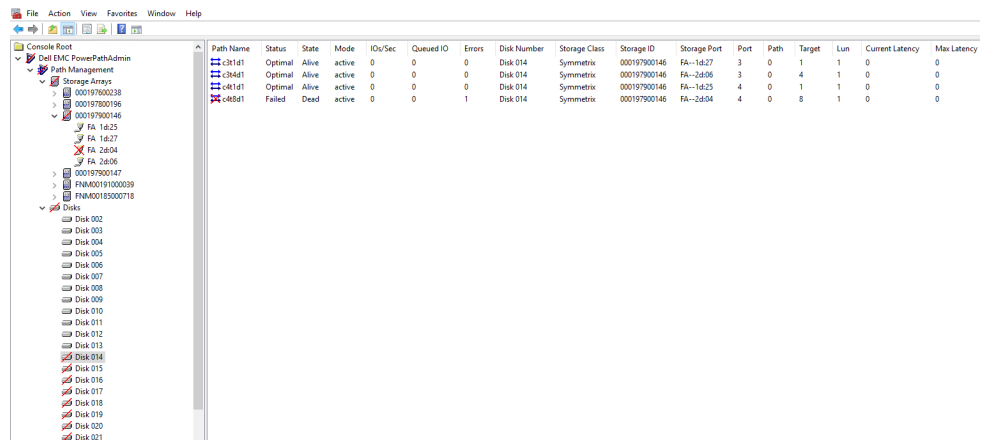
Object	Degraded icon	Failed icon	Unlicensed and degraded icon	Unlicensed and failed icon
Path	Not applicable.	 One or more hardware components that make up the path have failed; therefore, the entire path failed.	Not applicable.	 PowerPath Administrator is unlicensed and one or more of the path hardware components has failed; therefore the entire path failed.
Root node	 PowerPath Administrator is identifying a degradation within the managed environment.	 PowerPath Administrator is in a failed state.	 PowerPath Administrator is unlicensed and identifying a degradation within the managed environment.	 PowerPath Administrator is unlicensed and in a failed state.
Storage array	 One or more (but not all) paths to the storage array have failed or are in a degraded state.	 All paths to the array have failed. This array is not available.	 PowerPath Administrator is unlicensed. One or more (but not all) paths to the array have failed or are in a degraded state.	 PowerPath Administrator is unlicensed. All paths to the array have failed. This disk is not available.
Storage array port	 One or more (but not all) paths to the storage array port have failed or are in a degraded state.	 All paths to the array port have failed. This array port is not available.	 PowerPath Administrator is unlicensed. One or more (but not all) paths to the array port have failed or are in a degraded state.	 PowerPath Administrator is unlicensed. All paths to the array port have failed. This disk is not available.

The following list provides some failure examples:

- The following figure shows error with an array port or the path leading to the array port. The error is symbolized by the red X through one of the array ports. Access to the disk device is still available when the array port is down.

 **Note:** Degraded access is noted by a red slash (/).

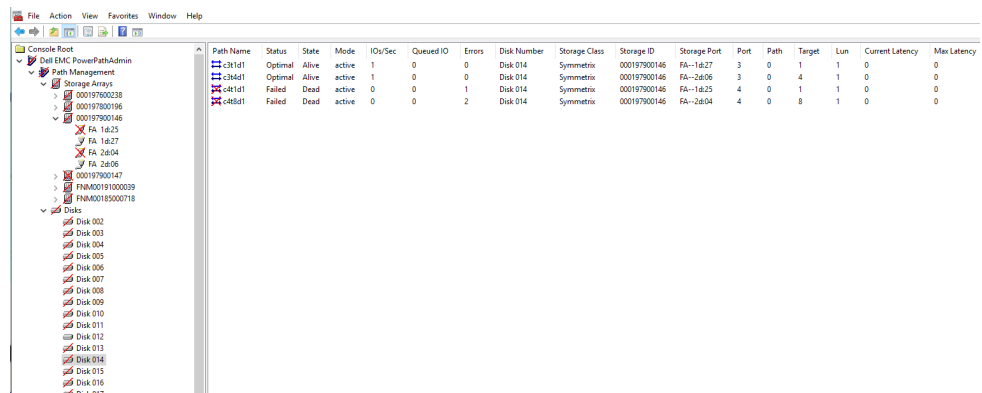
**Figure 15** Error with an array port



Path Name	Status	State	Mode	IOu/Sec	Queued IO	Errors	Disk Number	Storage Class	Storage ID	Storage Port	Port	Path	Target	Lun	Current Latency	Max Latency
c3d1d1	Optimal	Alive	active	0	0	0	Disk 014	Symmetrix	000197900146	FA--16:27	3	0	1	1	0	0
c3d4d1	Optimal	Alive	active	0	0	0	Disk 014	Symmetrix	000197900146	FA--26:06	3	0	4	1	0	0
c4d1d1	Optimal	Alive	active	0	0	0	Disk 014	Symmetrix	000197900146	FA--16:25	4	0	1	1	0	0
c4d4d1	Failed	Dead	active	0	0	1	Disk 014	Symmetrix	000197900146	FA--26:04	4	0	8	1	0	0

- The following figure shows the result of a problem with one of the HBAs or the path leading to the HBA. The failed HBA or path is marked with a red X. Access to the disk devices still exists even when it is degraded.

**Figure 16** Failed HBA path



Path Name	Status	State	Mode	IOu/Sec	Queued IO	Errors	Disk Number	Storage Class	Storage ID	Storage Port	Port	Path	Target	Lun	Current Latency	Max Latency
c3d1d1	Optimal	Alive	active	1	0	0	Disk 014	Symmetrix	000197900146	FA--16:27	3	0	1	1	0	0
c3d4d1	Optimal	Alive	active	1	0	0	Disk 014	Symmetrix	000197900146	FA--26:06	3	0	4	1	0	0
c4d1d1	Failed	Dead	active	0	0	1	Disk 014	Symmetrix	000197900146	FA--16:25	4	0	1	1	0	0
c4d4d1	Failed	Dead	active	0	0	2	Disk 014	Symmetrix	000197900146	FA--26:04	4	0	8	1	0	0

## Making changes to your environment

You must reconfigure PowerPath software after making configuration changes that affect the host-to-storage-system connectivity or logical device identification. Such configuration changes include:

- Fibre Channel (FC) switch zone changes
- Adding or removing FC switches
- Adding or removing host bus adapters (HBAs) or storage-system ports
- Adding or removing logical devices

In most cases, when you make changes to your environment, PowerPath software automatically detects it. Depending on the type of HBA you are using, you might have to scan for new devices in Device Manager. On some occasions and depending on the operating system version you are running, you might have to reboot your system.

## PowerPath messages

For a complete list of PowerPath messages and their meanings, see the *PowerPath Product Guide*.

# Microsoft Native MPIO

## MPIO with Windows Server 2012, 2012 R2, 2016, and 2019

### Support for MPIO in Windows Server 2012

Microsoft MultiPath I/O (MPIO) is a native feature of the Windows Server 2012 operating system. Dell EMC PowerMax and VMAX series, VNX, Unity, PowerStore, VPLEX, XtremIO, and CLARiiON arrays support MPIO.

The following list provides necessary guidelines:

- To use the Microsoft default Device Specific Module (DSM), storage must be compliant with SCSI Primary Commands-3 (SPC-3).
- Default DSM timer counters are supported.
- Hosts running Windows Server 2012 must be manually configured so that the initiators are registered using failover mode 4 (ALUA).
- CLARiiON systems must be running Flare 30 or later to support MPIO.
- VNX arrays require VNX OE for Block 31 or later.

### Configuring MPIO for Windows Server 2012

This section explains how to configure MultiPath I/O (MPIO) for Dell EMC storage arrays.

MPIO is installed as an optional feature of Windows Server 2012.

#### Configuring MPIO and installing DSM


When MPIO is installed, the MPIO control panel is also installed. You can use the control panel to do the following:

- Configure MPIO functionality
- Install additional storage Device Specific Modules (DSMs)
- Create MPIO configuration reports

You can open the MPIO control panel either by using the Windows Server 2012 control panel or Administrative Tools.

To open the MPIO control panel using the Windows Server 2012 control panel, click **Start > MPIO**.

To open the MPIO control panel using Administrative Tools, click **Start > Administrative Tools > MPIO**.

 **Note:** To access the MPIO control panel on Server Core installations, open a command prompt and type `MPIOCPL.EXE`.

The MPIO control panel opens to the **MPIO Properties** dialog box.

Once installed, native MPIO must be configured to manage PowerMax and VMAX series, VNX, Unity, PowerStore, VPLEX, XtremIO, and CLARiiON systems. Open **Control Panel**, then open the **MPIO** applet.

## Device discovery and claiming devices for MPIO

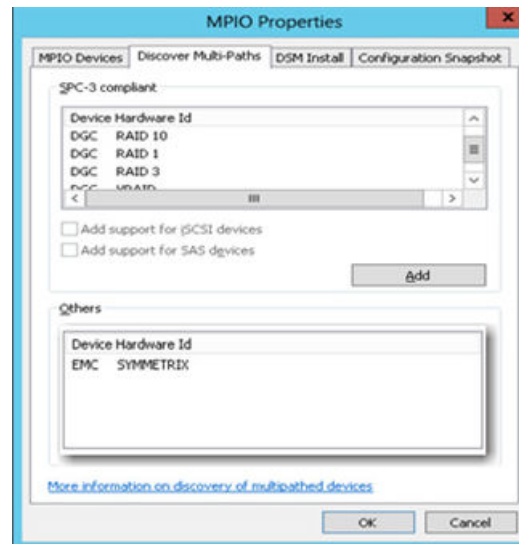
**CAUTION** When configuring MultiPath I/O (MPIO) on your system, do not exceed the MPIO maximum of 32 paths per LUN. If you exceed this number, the host will fail and produce a blue-screen error message.

1. In the **MPIO Properties** dialog box of the MPIO control panel applet, click the **Discover Multi-Paths** tab.  
Only arrays that are connected with at least two logical paths are available to be added in this tab.

- The **SPC 3-compliant** section of the applet lists only devices from:
  - VNX systems running VNX OE for Block 31
  - CLARiiON systems that are running Flare 30 or later and that are configured for failover mode 4 (ALUA)
- The **Others** section of the applet lists devices from the PowerMax and VMAX series, VNX, Unity, PowerStore, VPLEX, XtremIO, and CLARiiON system arrays.

2. Select the array and device types to be claimed and managed by MPIO by selecting the Device Hardware ID and clicking **Add**.

**Figure 17** MPIO Properties dialog box



**Note:** Although the operating system prompts you to reboot for each device type that you add, a single reboot after you add multiple device types is sufficient.

## Path management in MPIO

This section describes path management in MPIO for the following arrays: PowerMax and VMAX series, VNX, Unity, PowerStore, VPLEX, XtremIO, and CLARiiON systems.

After the reboot, once all device types have been claimed by MPIO, Device Manager shows each disk on the following arrays as a MultiPath Disk Device: PowerMax and VMAX series, VNX, Unity, PowerStore, VPLEX, XtremIO, and CLARiiON systems.

When a device is managed by MPIO, an **MPIO** tab is available under **Disk Device Properties** of the selected disk device. Under the **MPIO** tab, the number of logical paths that are configured between the host and array are reported.

The **MPIO** tab also enables you to change the MPIO load balance policy for a disk device.



**Note:** Some MPIO load balance policies might not be available for specific array disk types. For example, the Round Robin policy is not available for VNX disk devices, but Round Robin with Subset is available.

Available policy options are as follows:

**Fail Over Only**-This policy does not perform load balancing. It uses a single active path, and the rest of the paths are standby paths. The active path is used for sending all I/O. If the active path fails, then one of the standby paths is used. When the path that failed is reactivated or reconnected, the standby path that was activated returns to standby.

**Round Robin**-This load balance policy allows the Device Specific Module (DSM) to use all available paths for MPIO in a balanced way. This is the default policy that is chosen when the storage controller follows the active-active model and the management application does not specifically choose a load balance policy.

**Round Robin with Subset**-This load balance policy allows the application to specify a set of paths to be used in a round robin fashion and with a set of standby paths. The DSM uses paths from a primary pool of paths for processing requests as long as at least one of the paths is available. The DSM uses a standby path only when all the primary paths fail. For example: With four paths: A, B, C, and D, paths A, B, and C are listed as primary paths and D is the standby path. The DSM chooses a path from A, B, and C in round robin fashion as long as at least one of them is available. If all three paths fail, the DSM uses D, which is the standby path. If path A, B, or C becomes available, the DSM stops using path D and switches to the available paths among A, B, and C.

**Least Queue Depth**-This load balance policy sends I/O down the path with the fewest currently outstanding I/O requests. For example, consider that, of two I/Os, one I/O is sent to LUN 1 on Path 1, and the other I/O is sent to LUN 2 on Path 1. The cumulative outstanding I/O on Path 1 is 2, and on Path 2, it is 0. Therefore, the next I/O for either LUN will process on Path 2.

**Weighted Paths**-This load balance policy assigns a weight to each path. The weight indicates the relative priority of a path. The larger the number, the lower ranked the priority. The DSM chooses the least-weighted path from the available paths.

**Least Blocks**-This load balance policy that sends I/O down the path with the least number of data blocks currently being processed. For example, consider that, of two I/Os, one is 10 bytes and the other is 20 bytes. Both are in process on Path 1, and both have completed Path 2. The cumulative outstanding amount of I/O on Path 1 is 30 bytes. On Path 2, it is 0. Therefore, the next I/O will process on Path 2.

The default load balance policy for each disk device (as reported in the **MPIO** tab) depends upon the type of disk device that is presented.

In Windows Server 2012:

- VMAX and PowerMax series devices report a default load balance policy of **Round Robin**, where all paths are listed as **Active/Optimized**
- VNX series, PowerStore, and CLARiiON devices report a load balance policy of **Round Robin With Subset**, where all paths to the storage processor owning the device are **Active/Optimized**, and all paths to the storage processor not owning the LUN are **Active/Unoptimized**.
- VNX series, PowerStore, and CLARiiON devices that are attached to the host in ALUA mode (as is required when using MPIO) report the path state. The reported path state is used directly by the host running MPIO and cannot be overridden by changing the load balance policy.
- VPLEX devices report a default load balance policy of **Round Robin**, with all active paths as **Active/Optimized**.

Change load balance policies based on your environment. Usually, the default policy is suitable for your I/O load needs. However, some environments might require a change to the load balance policy to improve performance or better spread I/O load across storage front-end ports. Dell EMC

does not requires a specific load balance policy for any environment, and customers can change their load balance policies to meet their environment's needs.

For more information about Microsoft MPIO, see [Microsoft website](#) and [Microsoft TechNet](#).

## Veritas volume management software

Veritas Storage Foundation/Volume Manager (InfoScale) from Veritas replaces the native volume manager of the Windows operating system to enable management of physical disks as logical devices. The added features of this software are designed to enhance data storage management by controlling space allocation, performance, data availability, device installation, and system monitoring of private and shared systems. Veritas Storage Foundation (InfoScale) bundles path management software (DMP) as part of the Veritas product offering on Windows.

For more information about Veritas Storage Foundation, including documentation and software availability, see the [Veritas webiste](#)

To determine which Veritas Storage Foundation/Volume Manager configurations are supported and which service packs are required for your configuration, see the [Dell EMC Simple Support Matrix](#).

# CHAPTER 4

## Host Connectivity with Dell EMC Products

This chapters contains the following topics:

• <a href="#">PowerMax/VMAX All Flash</a> .....	40
• <a href="#">Midrange Storage</a> .....	44
• <a href="#">VPLEX</a> .....	60
• <a href="#">XtremIO</a> .....	84

## PowerMax/VMAX All Flash

This section provides support information about the PowerMax, VMAX All Flash, and VMAX3 series.

### Dell EMC PowerMax

Dell EMC PowerMax provides a platform that is ready for current and next generation data center data storage requirements. The PowerMax Family uses NVMe technology for customer application data.

PowerMax is built using a 100% NVMe storage back-end, allowing it to reach optimal IOPS densities by eliminating the flash media issues that are found in traditional SAS and SATA interfaces. NVMe over Fibre Channel 32 Gb front-end connections are supported starting with the PowerMaxOS 5978.444 release running on PowerMax storage systems allows for end-to-end NVMe accessibility, the PowerMaxOS 5978.444 and later releases running on PowerMax storage systems. It also introduces capability for traditional Fibre Channel FCP connections with an operation link speed of 32 Gb. NVMe over Fibre Channel and 32 Gb FCP both require the installation of the 32 Gb Fibre Channel I/O module. The PowerMaxOS 5978.479 and later releases leverages the 32 Gb Fibre Channel I/O module for 32 Gb SRDF link connection support. All PowerMaxOS releases maintain downward compatibility with the legacy 16 Gb Fibre Channel I/O module.

There are two PowerMax models as follows:

*PowerMax 2000*-PowerMax 2000 is designed to provide high efficiency and flexibility to the data center, providing 1.7 million IOPS (8 K RRH) and up to 1 PB of effective capacity in 20U total space. It is the entry NVMe scale out array sold with the Essentials and Pro software packages.

*PowerMax 8000*-PowerMax 8000 designed to provide optimal scalability, performance, and IOPS density. It can consolidate disparate workloads on a mass scale as 8 Bricks can support 10 million IOPS (8 K RRH) and provide up to 4 PB of effective capacity in just two floor tiles of space. It is the flagship NVMe scale out array sold with the Essentials and Pro software packages.

### Dell EMC PowerMax Operating System environment

PowerMaxOS 5978 is the only supported software operating environment for PowerMax storage arrays. It manages the storage and controls communications with the host systems and introduces support of dedupe and other new feature enhancements.

For detailed product specifications including the 32G Frontend I/O and 32G SRDF protocol support information, see Dell EMC PowerMax Family documentation available on <https://shop.dell EMC.com> and *Dell EMC PowerMax Family Product Guide* on [Dell EMC Online Support](#).

### Dell EMC VMAX All Flash

Dell EMC VMAX All Flash range of storage arrays use only high-density flash drives. The range contains four models that combine high scale, low latency, and rich data services:

- VMAX 250F with a maximum capacity of 1.16 PBe (Petabytes effective)
- VMAX 450F with a maximum capacity of 2.3 PBe
- VMAX 850F with a maximum capacity of 4.4 PBe
- VMAX 950F with a maximum capacity of 4.42 PBe

### Dell EMC VMAX All Flash Operating Environment

Dell EMC VMAX All Flash storage systems are supported with both HYPERMAX OS 5977 and PowerMaxOS 5978, however not all new feature enhancements that are introduced with

PowerMax storage systems running PowerMaxOS 5978 would be available with Dell EMC VMAX All Flash storage systems also running PowerMaxOS 5978.

HYPERMAX OS 5977 provides emulations that perform specific data service and control functions in the HYPERMAX environment, and it introduces an open application platform for running data services and provide file system storage with eNAS and embedded management services for Unisphere. The entire feature set available with HYPERMAX OS 5977 running on Dell EMC VMAX All Flash storage systems would be available with PowerMaxOS 5978 running on Dell EMC VMAX All Flash storage systems except FCoE front-end connections.

For detailed product specifications including the Frontend I/O protocol support information, see *VMAX All Flash Product Guide* and *VMAX All Flash: Family Overview* available on <https://www.dell.com/>.

## PowerMax/VMAX All Flash/VMAX3 - Device types

There are two classes of devices available on PowerMax/VMAX All Flash/VMAX3 storage arrays:

- **Compatibility device** - The Compatibility device is the default VMAX device type, advertises T10 SPC-3 compliancy and behaves almost like the legacy VMAX devices.
- **Mobility device** - The Mobility device is an optionally configurable device type that is supported with HYPERMAX OS 5977.811.784 and later, advertises T10 SPC-4 compliancy, and has numerous INQUIRY and VPD page changes implemented intended to allow for support of future feature enhancements. Implicit Asymmetric Logical Unit Access (ALUA) is advertised by the Mobility device, however the Mobility device is configured into only a single active target port group per PowerMax or VMAX storage system and a second target port group is configured on a remote PowerMax or VMAX system only when the Mobility devices are configured in an SRDF Metro configuration which is restricted only to Windows and Solaris operating system environments for now. For all supported SRDF/Metro Mobility ID/ALUA, see the Support Matrix available on [Dell EMC E-Lab Navigator](#). Mobility devices are also not currently supported with NDM migration.

## Local replication services

### TimeFinder SnapVX

Dell EMC TimeFinder SnapVX creates and manages point-in-time snapshots of critical data that can be used for backups, decision support, and to refresh data warehouse, test, and development environments. SnapVX snapshots do not require target volumes. SnapVX snapshots share back-end allocations with the source volume and other snapshots on the source volume.

TimeFinder SnapVX is supported on VMAX All Flash arrays running HYPERMAX OS 5977 and later, and snapshots are always consistent. Consistency across multiple arrays is achieved when source devices are in a composite group.

SnapVX provides very low impact snapshots and clones for VMAX LUNs. SnapVX supports up to 256 snapshots per source volume, which are tracked as versions with less overhead and simple relationship tracking. Users can assign names to identify their snapshots, and can set automatic expiration dates on each snapshot. SnapVX provides the ability to manage consistent point-in-time copies for storage groups with a single operation. Up to 1024 target volumes can be linked per source volume, providing read/write access as pointers or full-copy clones. TimeFinder in HYPERMAX OS also provides compatibility modes for users who rely on their TimeFinder Mirror, Clone, or VP Snap command scripts. This allows users to leverage their existing scripts while learning how to take advantage of the new features of SnapVX.

For details, see the *TimeFinder SNAPVX Local Replication* document available on PowerMax and VMAX All Flash Technical Documentation page on <https://www.dell.com/> or <https://www.dell.com/support>.

## Remote Replication Services

### SRDF

The Symmetrix Remote Data Facility (SRDF) maintains real-time or near real-time copies of data on a production storage array at one or more remote storage arrays and only one SRDF device mirror can be read/write enabled at any point in time.

### SRDF/Metro

HYPERMAX OS 5977.691.684 and Solutions Enabler/Unisphere for VMAX 8.1 first introduced the support for SRDF/Metro for VMAX3 and VMAX All Flash families of storage arrays and all versions of PowerMaxOS on Dell EMC PowerMax arrays.

With SRDF/Metro, the SRDF secondary device is read/write accessible to the host and takes on the external identity of the primary device (geometry, device WWN, and so on). By providing this external identity on the secondary device, both the primary and secondary devices may then appear as a single virtual device across the two SRDF paired arrays for presentation to a single host or host cluster.

With both the local and remote devices being read/write accessible concurrently, the host or hosts (in a cluster) can read and write to both primary and secondary devices with SRDF/Metro ensuring that each copy remains current, consistent, and addressing any write conflicts which may occur between the paired SRDF devices. A single PowerMax/VMAX3/VMAX All Flash array may simultaneously support multiple SRDF groups that are configured for SRDF/Metro operations and multiple SRDF groups that are configured for non-SRDF/Metro operations. PowerMaxOS Q3 2019 release with SRDF/Metro supports online device expansion through which devices could be expanded within an active SRDF/Metro group. These operations can be accomplished using the latest version of Unisphere for PowerMax or Solutions enabler. Both SRDF and SRDF/Metro support FC or GiGE links for connectivity between arrays.

For more detailed information about SRDF/metro configuration, see *Dell EMC SRDF Introduction* and *SRDF/Metro Overview and Best Practices Technical Note* available on <https://www.dell.com/support>.

## Non-Disruptive Migration services

Dell EMC's Non-Disruptive Migration (NDM) allows user to perform online data migrations that are simple and completely non-disruptive to the host and application. NDM is designed to help automate the process of migrating hosts and applications to a new PowerMax array with no downtime. NDM leverages SRDF replication technologies to move the application data to the new array. It also uses auto-provisioning, with PowerPath or a supported host multipathing solution, to manage host access to the data during the migration process.

NDM provides PowerMax user the following benefits:

- Allows non-disruptive migration with hosts and applications online. Potentially supported source storage systems are legacy VMAX v2, VMAX3 series, VMAX All Flash systems, and PowerMax. Potentially supported target systems are VMAX3 series, VMAX All Flash systems, and PowerMax.
- Ease of use with control operations that automate the setup and configuration of the migration environment.
- Managed by familiar and simple user interfaces using Solutions Enabler and Unisphere.
- Migrations can be easily canceled and failed back to the source array for any reason prior to commit.
- Completely built in and does not require any additional software or licensing costs.


For more detailed features, see the *Dell EMC PowerMax and VMAX: Non-Disruptive Migration Best Practices and Operational Guide* white paper available on <https://www.dell.com/whitepapers>.

See Dell EMC Simple Support Matrix available on [Dell EMC E-Lab Navigator](#) for host interoperability with various operating system platforms and multipathing software supported for Non-Disruptive migrations.

Use advanced query option on [Dell EMC E-lab Navigator](#) for specific configuration search.

## PowerMax/VMAX All Flash-Storage provisioning

PowerMax and VMAX All flash series initial configuration is performed by a Dell EMC Customer Engineer (CE) through the PowerMax or VMAX Management Module Control Station (MMCS). The CE will configure the storage arrays settings for each Fibre Channel port. The procedures in this document assume that any switches and storage systems to be used in this configuration have been installed, and that the front-end adapter ports have been connected to the appropriate switch ports and switch zoning is completed.

 **Note:** It is recommended to use Access Logic (ACLX) to mask volumes.

Storage provisioning operations can be accomplished by Solutions enabler or Unisphere for PowerMax software.

### Dell EMC Solutions Enabler

Dell EMC Solutions Enabler installation provides your host with SYMAPI, CLARAPI, and STORAPI shared libraries for use by Solutions Enabler applications, and the Symmetrix Command Line Interface (SYMCLI) for use by storage administrators and systems engineers. SYMCLI is a specialized library of UNIX-formatted commands that can be invoked one at a time. It supports single command-line entries and scripts to map and perform control operations on devices and data objects toward the management of your storage complex. It also monitors device configuration and status of devices that make up the storage environment. The target storage environments are typically PowerMax and VMAX arrays.

For detailed procedure, see *Dell EMC Solutions Enabler Array Controls and Management Guide* available on <https://www.dell.com/support>.

### Dell EMC Unisphere for PowerMax

Dell EMC Unisphere for PowerMax enables management and monitoring of PowerMax arrays along with legacy VMAX All Flash and VMAX3 arrays. Unisphere is installed on a dedicated Windows or Linux server, or deployed as a Virtual Appliance (vAPP).

For detailed procedure, see *Dell EMC Unisphere for PowerMax Installation Guide* available on [Dell EMC online Support](#) under PowerMax and VMAX technical documents and videos section.

## Dell EMC PowerMax/VMAX All Flash/VMAX3 - Director bit settings

Dell EMC PowerMax/VMAX All Flash/VMAX3 supports various director bit settings against different front-end protocol such as Fibre Channel FCP, NVMe over Fibre Channel (NVMeOF), iSCSI, and FCoE. They vary between switch-based connectivity and direct attach connectivity with various operating system types.

For detailed director bit settings, see the Simple Support Matrix available on [Dell EMC E-Lab Navigator](#).

## Dell EMC PowerMax Fibre Channel FCP connectivity

Dell EMC PowerMax arrays running with PowerMaxOS Q3 2019 introduces 32 Gb/s Fibre Channel modules support. Earlier versions of PowerMaxOS on PowerMax arrays and HYPERMAX OS on VMAX All Flash arrays supports 16 Gb/s front-end module for host connectivity.

For details, see the *Understanding storage provisioning* and *Understanding Host Management* sections in *Dell EMC Unisphere for PowerMax Product Guide* available on <https://www.dell.com/> or [Dell EMC Online Support](#).

## Dell EMC PowerMax NVMe-OF connectivity

Dell EMC PowerMax NVMe-OF connectivity is not supported with Windows environment.

## Dell EMC PowerMax iSCSI connectivity

Dell EMC PowerMax array supports high-density quad port 10 Gb/s interface module. For detailed iSCSI host connectivity information, see *Dell EMC PowerMax: iSCSI Implementation for Dell EMC Storage Arrays Running PowerMaxOS* available on <https://www.dell.com/> or [Dell EMC Online Support](#) under PowerMax and VMAX All Flash Technical Documentation page.

# Midrange Storage

This section describes host connectivity of the Dell EMC Midrange storage arrays, including Unity and PowerStore.

## Unity storage

## Storage configuration

### General recommendations for storage pools

Dell EMC Unity supports two types of storage pools: traditional pools and dynamic pools. The following recommendations are applicable to both types of pool.

Dell EMC recommends using fewer storage pools within Dell EMC Unity to reduce complexity and increase flexibility. However, it may be appropriate to configure multiple storage pools to:

- Separate workloads with different I/O profiles
- Separate pools where FAST Cache is active and not active
- Dedicate resources to meet specific performance goals
- Separate resources for multi-tenancy
- Create smaller failure domains

### Storage pool capacity

Storage pool capacity is used for multiple purposes:

- To store all data written into storage objects (LUNs, file systems, data stores, and VVols) in that pool.
- To store data that is needed for Snapshots of storage objects in that pool.
- To track changes to replicated storage objects in that pool



- To perform efficient data relocations for FAST VP

Storage pools must maintain free capacity to operate properly. By default, Dell EMC Unity will raise an alert if a storage pool has less than 30% free capacity, and will begin to automatically invalidate Snapshots and Replication sessions if the storage pool has less than 5% free capacity. Dell EMC recommends that a storage pool always have at least 10% free capacity.

## Raid protection

Dell EMC Unity applies RAID protection to the storage pool to protect user data against drive failures. Choose the RAID type that best suits your needs for performance, protection, and cost.

- RAID-1/0 provides the highest level of performance from a given set of drive resources, with the lowest CPU requirements; however, only 50% of the total drive capacity is usable.
- RAID-5 provides the best usable capacity from a set of drive resources, but at lower overall performance and availability than RAID-1/0.
- RAID-6 provides better availability than RAID-5 and better usable capacity than RAID-1/0, but has the lowest performance potential of the three RAID types.

## Traditional pools

Traditional Storage Pools apply RAID protection to individual groups of drives within the storage pool. Traditional pools are the only type of pool available on Dell EMC Unity hybrid systems, and are also available on all-Flash systems.

## Raid protection

For traditional pools, Dell EMC recommends RAID-5 for drives in Extreme Performance and Performance tiers, and RAID-6 for drives in the Capacity tier.

Assuming that roughly the same number of drives will be configured in a traditional pool, Dell EMC recommends smaller RAID widths as providing the best performance and availability, at the cost of slightly less usable capacity.

Example: When configuring a traditional pool tier with RAID-6, use 4+2 or 6+2 as opposed to 10+2 or 14+2.

When choosing RAID-1/0, 1+1 can provide better performance with the same availability and usable capacity as larger RAID widths (assuming that the same total number of drives are used), and also provides more flexibility.

## All-flash pool

All-flash pools provide the highest level of performance in Dell EMC Unity. Use an all-flash pool when the application requires the highest storage performance at the lowest response time.

Snapshots and Replication operate most efficiently in all-flash pools. Data Reduction is only supported in an all-flash pool.

FAST Cache and FAST VP are not applicable to all-flash pools.

Dell EMC recommends using only a single drive size and a single RAID width within an all-flash pool.

## Dynamic pools

Dynamic Storage Pools apply RAID protection to groups of drive extents from drives within the pool, and allow for greater flexibility in managing and expanding the pool. Dynamic pools are only available on Dell EMC Unity all-Flash systems, and therefore must be all-Flash pools; dynamic pools cannot be built with HDDs.

## RAID protection

At the time of creation, dynamic pools use the largest RAID width possible with the number of drives that are specified, up to the following maximum widths:

- RAID-1/0: 4+4
- RAID-5: 12+1
- RAID-6: 14+2

With dynamic pools, there is no performance or availability advantage to smaller RAID widths. To maximize usable capacity with parity RAID, Dell EMC recommends to initially create the pool with enough drives to guarantee the largest possible RAID width.

- For RAID-5, initially create the pool with at least 14 drives.
- For RAID-6, initially create the pool with at least 17 drives.

## Spare capacity

Hot spares are not needed with dynamic pools. A dynamic pool automatically reserves the capacity of one drive, as spare space in the pool, for every 32 drives. If a drive fails, the data that was on the failed drive is rebuilt into the spare capacity on the other drives in the pool. Also, unbound drives of the appropriate type can be used to replenish the spare capacity of a pool, after the pool rebuild has occurred.

Example: For an All-Flash pool, use only 1.6 TB SAS Flash 3 drives, and configure them all with RAID-5 8+1.

## Hybrid pool

Hybrid pools can contain HDDs (SAS and NL-SAS drives) and flash drive, and can contain more than one type of drive technology in different tiers. Hybrid pools typically provide greater capacity at a lower cost than all-flash pools, but also typically have lower overall performance and higher response times. Use hybrid pools for applications that do not require consistently low response times, or that have large amounts of mostly inactive data.

Performance of a hybrid pool can be improved by increasing the amount of capacity in the flash drive tier, so that more of the active dataset resides on and is serviced by the flash drives. See the [FAST VP](#) section.

Hybrid pools can have up to three tiers (Extreme Performance, Performance, and Capacity). Dell EMC recommends using only a single drive speed, size, and RAID width within each tier of a hybrid pool.

Example:

- For the Extreme Performance tier, use only 800 GB SAS flash 2 drives, and configure them all with RAID-5 8+1.
- For the Performance tier, use only 1.2 TB SAS 10K RPM drives, and configure them with RAID-5 4+1.
- For the Capacity tier, use only 6 TB NL-SAS drives, and configure them all with RAID-6 6+2.

## Storage object types

By default, Dell EMC Unity creates thin storage objects. Thin storage objects are virtually provisioned and space efficient. In general, Dell EMC recommends using thin storage objects, as they provide the best capacity utilization, and are required for most features. Thin storage objects are recommended when any of the following features will be used:

- Data Reduction
- Snapshots
- Thin Clones
- Asynchronous Replication

Thick storage objects will reserve capacity from the storage pool, and dedicate it to that particular storage object. Thick storage objects guarantee that all advertised capacity is available for that object. Thick storage objects are not space efficient, and therefore do not support the use of space-efficient features. If it is required to enable a space-efficient feature on a thick storage object, it is recommended to first migrate the thick storage object to a thin storage object, and enable the feature during the migration (for Data Reduction) or after migration has completed (for Snapshots, Thin Clones, and Asynchronous Replication).

In addition to capacity for storing data, storage objects also require pool capacity for metadata overhead. The overhead percentage is greater on smaller storage objects. For better capacity utilization, Dell EMC recommends configuring storage objects that are at least 100GB in size, and preferably at least 1TB in size.

## Features

### FAST VP

Fully Automated Storage Tiering (FAST) for Virtual Pools (VP) accelerates performance of a specific storage pool by automatically moving data within that pool to the appropriate drive technology, based on data access patterns. FAST VP is applicable to hybrid pools only within a Dell EMC Unity hybrid system.

The default and recommended FAST VP policy for all storage objects is **Start High then Auto-tier**. This policy places initial allocations for the storage object in the highest tier available, and monitors activity to this storage object to determine the correct placement of data as it ages.

FAST VP is most effective if data relocations occur during or immediately after normal daily processing. Dell EMC recommends scheduling FAST VP relocations to occur before backups or nightly batch processing. For applications which are continuously active, consider configuring FAST VP relocations to run constantly.

Dell EMC recommends maintaining at least 10% free capacity in storage pools, so that FAST VP relocations can occur efficiently. FAST VP relocations cannot occur if the storage pool has no free space.

### FAST Cache

FAST Cache is a single global resource that can improve performance of one or more hybrid pools within a Dell EMC Unity hybrid system. FAST Cache can only be created with SAS Flash 2 drives, and is only applicable to hybrid pools. Dell EMC recommends to place a Flash tier in the hybrid pool before configuring FAST Cache on the pool. FAST Cache can improve access to data that is resident in the HDD tiers of the pool.

Enable FAST Cache on the hybrid pool if the workload in that pool is highly transactional, and has a high degree of locality that changes rapidly.

For applications that use larger I/O sizes, have lower skew, or do not change locality as quickly, it may be more beneficial to increase the size of the Flash tier rather than enable FAST Cache.

FAST Cache can increase the IOPS achievable from the Dell EMC Unity system, and this will most likely result in higher CPU utilization (to service the additional I/O). Before enabling FAST Cache on additional pools or expanding the size of an existing FAST Cache, monitor the average system

CPU utilization to determine if the system can accommodate the additional load. See [Table 3](#) for recommendations.

## Data Reduction

Dell EMC Unity Data Reduction by compression is available for Block LUNs and VMFS datastores in an all-flash pool starting with Dell EMC Unity OE 4.1. Data reduction via compression is available for file systems and NFS datastores in an all-flash pool starting with Dell EMC Unity OE 4.2. Beginning with Dell EMC Unity OE 4.3, data reduction includes both compression and deduplication.

Be aware that data reduction increases the overall CPU load on the system when storage objects service reads or writes of reduceable data, and may increase latency. Before enabling data reduction on a storage object, Dell EMC recommends to monitor the system and ensure that the system has available resources to support data reduction (See [Table 3](#) to the Hardware Capability Guidelines). Enable data reduction on a few storage objects at a time, and then monitor the system to be sure it is still within recommended operating ranges, before enabling data reduction on more storage objects.

For new storage objects, or storage objects that are populated by migrating data from another source, Dell EMC recommends to create the storage object with data reduction enabled, before writing any data. This provides maximum space savings with minimal system impact.

## Advanced Deduplication

Dell EMC Unity Advanced Deduplication is an optional extension to Data Reduction, that you can enable to increase the capacity efficiency of data reduction enabled storage objects. Beginning with Dell EMC Unity OE 4.5, advanced deduplication is available for storage objects in dynamic pools on Dell EMC Unity 450F, 550F, and 650F All-Flash systems.

As with data reduction, advanced deduplication is only applied to data when it is written to the storage object. LUN Move can be utilized to deduplicate existing data on Block storage objects.

For new storage objects, or storage objects that will be populated by migrating data from another source, it is recommended to create the storage object with advanced deduplication enabled, before writing any data. This provides maximum space savings with minimal system impact.

## Snapshots

Dell EMC recommends including a Flash tier in a hybrid pool where snapshots will be active.

Snapshots increase the overall CPU load on the system, and increase the overall drive IOPS in the storage pool. Snapshots also use pool capacity to store the older data being tracked by the snapshot, which increases the amount of capacity used in the pool, until the snapshot is deleted. Consider the overhead of snapshots when planning both performance and capacity requirements for the storage pool.

Before enabling snapshots on a storage object, it is recommended to monitor the system and ensure that existing resources can meet the additional workload requirements (See [Table 2](#) for Hardware Capability Guidelines ). Enable snapshots on a few storage objects at a time, and then monitor the system to be sure it is still within recommended operating ranges, before enabling more snapshots.

Dell EMC recommends to stagger snapshot operations (creation, deletion, and so on). This can be accomplished by using different snapshot schedules for different sets of storage objects. It is also recommended to schedule snapshot operations after any FAST VP relocations have completed.

Snapshots are deleted by the system asynchronously; when a snapshot is in the process of being deleted, it will be marked as *Destroying*. If the system is accumulating Destroying snapshots over time, it may be an indication that existing snapshot schedules are too aggressive; taking snapshots less frequently may provide more predictable levels of performance. Dell EMC Unity will throttle

snapshot delete operations to reduce the impact to host I/O. Snapshot deletes will occur more quickly during periods of low system utilization.

## Thin Clones

Dell EMC recommends including a flash tier in a hybrid pool where thin clones will be active.

Thin clones use snapshot technology to provide space-efficient clones of block objects. Consider the overhead of snapshots when planning performance and capacity requirements for a storage pool which will have thin clones.

## Asynchronous replication

Dell EMC recommends including a Flash tier in a hybrid pool where asynchronous replication is active. This is applicable to both the source and the destination pools.

Dell EMC recommends configuring multiple replication interfaces per SP, and distributing replication sessions across them. Link Aggregation Control Protocol (LACP) can also be used to aggregate bandwidth for a replication interface. Configure Jumbo frames (MTU 9000) when possible.

Asynchronous replication takes snapshots on the replicated storage objects to create the point-in-time copy, determine the changed data to transfer, and maintain consistency during the transfer. Consider the overhead of snapshots when planning performance and capacity requirements for a storage pool that has replicated objects.

When possible, fill the source storage object with data before creating the replication session. The data will then be transmitted to the destination storage object during initial synchronization. This is typically the fastest way to populate the destination storage object with asynchronous replication.

Setting smaller RPO values on replication sessions will not make them transfer data more quickly; but smaller RPOs result in more frequent snapshot operations. Choosing larger RPOs, or manually synchronizing during nonproduction hours, may provide more predictable levels of performance.

## Synchronous replication/Metrosync for file

Dell EMC recommends including a Flash tier in a hybrid pool where synchronous replication will be active. This is applicable to both the source and the destination pools.

Synchronous replication transfers data to the remote system over the first Fibre Channel port on each SP. When planning to use synchronous replication, it may be appropriate to reduce the number of host connections on this port. When the CNA ports are configured as FC, CNA port 4 is defined as the replication port. If the CNA ports are configured as Ethernet, then port 0 of the lowest numbered FC I/O Module is the replication port.

When possible, create the synchronous replication session before filling the source storage object with data, as this alleviates the need to perform initial synchronization of the replication session. This is typically the fastest way to populate the destination storage object with synchronous replication.

When sizing a disaster recovery solution using synchronous replication, consider provisioning a destination system that has similar performance capabilities as the source system. This can help maintain the same level of application performance after a failover event.

## SAN Copy

SAN Copy provides one-time migration of Block resources from a third-party array, using either iSCSI or FC connections. When using FC, note that SAN Copy must use different ports than the FC ports which are designated for Synchronous Replication. This is true even if Synchronous Replication is not actively being used.

To lessen the impact of SAN Copy migrations on other host activity, consider reducing the number of host connections on the FC ports used for SAN Copy.


## NDMP

Dell EMC Unity supports 2-way NDMP for file data, which enables the system to send file data directly to a backup device using FC connections. Make sure that NDMP uses different ports than the FC ports which are designated for Synchronous Replication. This is true even if Synchronous Replication is not actively being used.

To lessen the impact of 2-way NDMP backups on other host activity, consider reducing the number of host connections on the FC ports that are used for NDMP.

## Data at Rest Encryption

Data at Rest Encryption (D@RE) is Controller Based Encryption that does not impact performance; therefore Dell EMC recommends ordering Dell EMC Unity systems as encryption-enabled, if appropriate for your environment.

 **Note:** Encryption can only be enabled at the time of system installation with the appropriate license.

If encryption is enabled, Dell EMC recommends making external backups of the encryption keys after system installation, and immediately following any change in the system's drives (such as, creating or expanding a storage pool, adding new drives, replacing a faulted drive, and so on).

## Host I/O limits

Dell EMC recommends setting Host I/O Limits on workloads which might monopolize pool resources and starve other applications of their required performance. Consider some of the following opportunities to utilize Host I/O Limits:

- Limit the bandwidth available to large-block applications, such as backup jobs, which may be increasing the latency on other small-block workloads.
- Limit the IOPS capability of Thin Clones which are used in Test/Dev environments, so that they do not impact the capability of the associated source objects.
- Limit the IOPS / bandwidth that is available to non-essential applications that are hosted on the same Dell EMC Unity system as your critical applications.

## Application considerations

### Host alignment for block LUNs

Alignment only needs to be done for host operating systems which still use a 63-block disk header. If alignment is required, perform the alignment using a host-based method, and align with a 1MB offset.

See the *Host Connectivity Guide* on [Dell EMC Online Support](#) to determine if alignment is required for your operating system, and how to perform the alignment.

## Hosts

Under the **ACCESS** category in the main navigation menu, users can configure hosts (Windows or Linux/UNIX) for storage access. Before a network host can access block storage or NFS file systems, the user must define a configuration for the host and associate it with a storage resource. SMB file systems can automatically be accessed by authorized users once provisioned. Users can use the Hosts page, as shown in the following figure to configure host configurations.

This can be done on an individual host-by-host basis or through subnet and netgroup configurations that allow access to multiple hosts or network segments. For block resources, before the user starts to configure a host, the user should ensure that initiator interfaces are configured and initiator registration completed. Once a host configuration is completed, users can go to the properties of a storage resource and specify the hosts, subnets, or netgroups from which they want the resource to be accessed.

**Figure 18 Hosts**

Name	Network Addresses	Operating System	Type	Tenant	LUNs	Initiators	Initiator Pat...
RHEL17229			Manual	--	0	0	0
Windows-KAP			Manual	--	6	1	2

## Initiators

To ensure that hosts can access block storage resources, the user must register initiators between the storage system and configured hosts. On the Initiators page, as shown in the following figure, users can manually register one or more Fibre Channel or iSCSI initiators. Initiators are endpoints from which Fibre Channel and iSCSI sessions originate, where each initiator is uniquely identified by its World Wide Name (WWN) or iSCSI Qualified Name (IQN). The link between a host initiator and a target port on the storage system is called the initiator path. Each initiator can be associated with multiple initiator paths. The Initiator Paths tab shows all data paths that are currently available to the initiators connected to the system either by FC or iSCSI. For iSCSI paths to show up, iSCSI interfaces must be configured on the Block Page. These initiators can then be discovered and registered by hosts using the iSCSI initiator tool (that is, the Microsoft iSCSI Initiator). For Fibre Channel paths, FC zoning on the appropriate switch is needed for the initiator paths to be seen as available by the system. Once the paths are available, users can configure their connected hosts on the Hosts Page.

**Figure 19 Initiators**

Initiator IQN/WWN	Host	Host Type	Protocol	Ignore	iSCSI Type	Bound	CHAP User...
20:00:00:00:C9:81:3F:B8:10:00:00:00:C9:81:3F:B8		Auto	FC	No	--	--	--
20:00:00:00:C9:81:3F:B9:10:00:00:00:C9:81:3F:B9		Auto	FC	No	--	--	--
20:00:00:1B:32:6A:61:7C:21:00:00:1B:32:6A:61:...		Auto	FC	No	--	--	--
20:00:00:90:FA:0C:69:E0:10:00:00:90:FA:0C:69:E0		Auto	FC	No	--	--	--
20:00:00:90:FA:0C:69:E1:10:00:00:90:FA:0C:69:E1		Auto	FC	No	--	--	--
iqn.1991-05.com.microsoft:win-8d1n17ne9d.prod...	Windows-KAP	Manual	iSCSI	No	--	--	--
iqn.1998-01.com.vmware:566eadee-80e6-72c4-...		Auto	iSCSI	No	Software	Yes	--

With the release of Dell EMC Unity OE version 4.3, Initiators can now have advanced settings customized through Unisphere. In order to access these settings, select an Initiator and then click

the pencil icon to bring up the Edit Initiator window. Clicking Advanced at the bottom to reveal the Initiator Source Type, Fail-over Mode, Unit Serial Number, and LunZ Enabled settings, as shown in the following figure. For more information about configuring Host Initiator Parameters, please reference the Online Help through Unisphere.

**Figure 20 Edit Initiator - Advanced**

The screenshot shows a dialog box titled "Advanced" for editing initiator settings. It includes the following fields:

- Initiator Source Type:** A dropdown menu currently showing "Open\_Native".
- Fail-over Mode:** A dropdown menu currently showing "Asymmetric Active Active".
- Unit Serial Number:** Two radio buttons, with "Array" selected and "Lun" unselected.
- Is LunZ Enabled:** An unchecked checkbox.

At the bottom right of the dialog are two buttons: "Close" and "OK".

## Snapshot schedule

Dell EMC Unity enables you to take point-in-time snapshots for all storage resources (block or file) to meet protection and recovery requirements in the event of corruption or accidental deletion.

The Snapshot Schedule page, as shown in the [figure](#), enables users to set the system to periodically take snapshots of storage resources automatically. Automating these administrative tasks takes away some of the management aspects of data protection. After enabling a snapshot schedule for a resource, each snapshot that is taken is time-stamped with the date and time of when it was created, and contains a point-in-time image of the data in the storage resource. The default snapshot schedules available on the system are:

- **Default protection** - A snapshot is taken at 08:00 (UTC) every day, and the snapshot is retained for 2 days.
- **Protection with shorter retention** - A snapshot is taken at 08:00 (UTC) every day, and the snapshot is retained for 1 day.
- **Protection with longer retention** - A snapshot is taken at 08:00 (UTC) every day, and the snapshot is retained for 7 days.

**Note:** Times are displayed in a user's local time in a 12-hour format and that default snapshot schedules cannot be modified; but custom snapshot schedules can be configured by selecting the intervals, times, and days for the system to take snapshots regularly.

With the Dell EMC Unity OE version 4.4 or later, user-defined Snapshot Schedules can be replicated using the Synchronous Replication connection that is established between two physical systems. Reference the new **Sync Replicated** column in the Snapshot Schedule page, as shown in the following figure. Applying a replicated Snapshot Schedule is only enabled in synchronously replicated file resources.



**Figure 21** Snapshot schedule

Name	Type	Sync Replicated	In Use
Default Protection	System Defined	No	No
Protection with shorter retention	System Defined	No	No
Protection with longer retention	System Defined	No	No
Replicated_Schedule	User Defined	Yes	Yes

Description: Note: Times are displayed in Local Time (UTC -04:00 in 12-hour format)

For more information about the snapshot technology available on Dell EMC Unity systems, see the *Dell EMC Unity: Snapshots and Thin Clones* and *Dell EMC Unity: MetroSync for File* white papers on [Dell EMC Online Support](#).

### Front-end connectivity

Dell EMC Unity provides multiple options for front-end connectivity, using on-board ports directly on the DPE, and using optional I/O Modules. This section discusses recommendations for the different types of connectivity.

In general, front-end ports need to be connected and configured symmetrically across the two storage processors (SPs), to facilitate high availability and continued connectivity if there is SP failure.

Example - A NAS Server is configured so that NAS clients connect using port 0 of the first I/O Module on SPA; therefore port 0 of the first I/O Module on SPB must be cabled so that it is accessible to the same networks.

For best performance, Dell EMC recommends using all front-end ports that are installed in the system, so that workload is spread across as many resources as possible.

Example - configuring the 4-port Fibre Channel I/O Module, zone different hosts to different ports so that all eight ports across the 2 SPs are used; do not zone all hosts to the first port of each I/O Module.

### Fibre Channel

When configured for Fibre Channel (FC), Dell EMC Unity CNA ports and I/O Module ports can be configured with 8 GB or 16 GB SFPs. All FC ports can negotiate to lower speeds. 16 GB FC is recommended for the best performance.

Dell EMC recommends single-initiator zoning when creating zone sets. For high availability purposes, a single host initiator should be zoned to at least one port from SPA and one port from SPB. For load balancing on a single SP, the host initiator can be zoned to two ports from SPA and two ports from SPB. When zoning additional host initiators, zone them to different SP ports when possible, to spread the load across all available SP ports.

Utilize multipathing software on hosts that are connected using FC, such as Dell EMC PowerPath, which coordinates with the Dell EMC Unity system to provide path redundancy and load balancing.

### iSCSI

Dell EMC Unity supports iSCSI connections on multiple 1 Gb/s and 10 GB/s port options. 10GBase-T ports can autonegotiate to 1 GB/s speeds. 10 GB/s is recommended for the best performance. If possible, configure Jumbo frames (MTU 9000) on all ports in the end-to-end network, to provide the best performance.

To achieve optimal iSCSI performance, use separate networks and VLANs to segregate iSCSI traffic from normal network traffic. Configure standard 802.3x Flow Control (Pause or Link Pause) on all iSCSI Initiator and Target ports that are connected to the dedicated iSCSI VLAN.

Dell EMC Unity supports 10 GbE and 1GBase-T ports that provide iSCSI offload. Specifically, the CNA ports (when configured as 10GbE or 1GBase-T) and the 2-port 10GbE I/O Module ports provide iSCSI offload. Using these modules with iSCSI can reduce the protocol load on SP CPUs by 10-20%, so that those cycles can be used for other services.

Utilize multipathing software on hosts that are connected using iSCSI, such as Dell EMC PowerPath, which coordinates with the Dell EMC Unity system to provide path redundancy and load balancing.

### Network-attached storage (NAS)

Dell EMC Unity supports NAS (NFS, FTP, and/or SMB) connections on multiple 1 GB/s and 10 GB/s port options. 10GBase-T ports can auto-negotiate to 1 GB/s speed. 10 GB/s is recommended for the best performance. If possible, configure Jumbo frames (MTU 9000) on all ports in the end-to-end network, to provide the best performance.

Dell EMC recommends configuring standard 802.3x Flow Control (Pause or Link Pause) on all storage ports, switch ports, and client ports that are used for NAS connectivity.

Dell EMC Unity provides network redundancy for NAS using Link Aggregation Control Protocol (LACP) and Fail-Safe Networking (FSN). Combine FSN and LACP with redundant switches to provide the highest network availability. In addition to redundancy, LACP can also improve performance with multiple 1GBase-T connections, by aggregating bandwidth. LACP can be configured across any Ethernet ports that have the same speed, duplex, and MTU.

**Note:** LACP cannot be enabled on ports that are also used for iSCSI connections.

While LACP creates a link aggregation with multiple active links, FSN provides redundancy by configuring a primary link and a standby link. The standby link is inactive unless the entire primary link fails. If FSN is configured with links of different performance capability (such as a link aggregation of 10 GB/s ports, and a stand-alone 1 GB/s port), Dell EMC recommends that you configure the highest performing link as the primary.

NAS Servers are assigned to a single SP. All file systems that are serviced by that NAS Server have I/O processed by the SP on which the NAS Server is resident. For load-balancing, Dell EMC recommends that you create at least two NAS Servers per Dell EMC Unity system: one on SPA, and one on SPB. Assign file systems to each NAS Server such that front-end workload is approximately the same for each SP.

### Connectivity Options

The following tables provide maximum expected IOPS and bandwidth from the different ports that are available in the Dell EMC Unity system. (The capability of a port does not guarantee that the system can reach that level, nor does it guarantee that performance will scale with additional ports. System capabilities are highly dependent on other configuration parameters.)

SAS ports are used by the SPs to move data to and from the back-end drives; all other ports can be used to provide access to hosts.

**Table 3** provides maximum expected IOPS and bandwidth from a 12Gb SAS port. The base Dell EMC Unity configuration contains four ports.

The following tables provide maximum expected IOPS from the different ports that are available in the Dell EMC Unity system. (The capability of a port does not guarantee that the system can reach that level, nor does it guarantee that performance will scale with additional ports. System capabilities are highly dependent on other configuration parameters.)

**Table 2** Ports and expected IOPS

Port	Maximum IOPS per Port
16 GB FC	45,000

**Table 2** Ports and expected IOPS (continued)

Port	Maximum IOPS per Port
CNA or 4-port I/O Module	
8 GB FC CNA	45,000
10 GbE iSCSI CNA or 2-port I/O Module	25,000
10 GbE iSCSI 4-port I/O Module	30,000
10GBase-T iSCSI On-board or 4-port I/O Module	30,000
1GBase-T iSCSI CNA, On-board or 4-port I/O Module	3,000

The following table provides maximum expected IOPS from the front-end ports which provide File protocols (NFS and SMB).

**Table 3** Front-end ports and expected IOPS

Port	Maximum IOPS per Port
10 GbE NAS CNA or 2-port I/O Module	60,000
10 GbE NAS 4-port I/O Module	60,000
10GBase-T NAS On-board or 4-port I/O Module	60,000
1GBase-T NAS CNA, On-board or 4-port I/O Module	6,000

## PowerStore

### Storage configuration

#### Host connectivity

For the latest information about qualified hosts and boot bus adapters (HBA), see the [Dell EMC Support Matrix](#) or contact your Dell EMC representative.

#### Logical devices

The logical devices that are presented by PowerStore arrays are the same on each enclosure node. A logical unit (LU) reports itself Active Optimized on one node and Active Not Optimized on the other node.

## Configure native multipathing on PowerStore volumes presented to the Windows host

### Before you begin

Before configuring the native multipathing, you must enable MPIO on the server by adding the MPIO feature to Windows.

### Procedure

1. Open PowerShell on the host.
2. If MPIO is not installed, perform the following tasks to install and enable MPIO :

- a. Run the following command to install MPIO :

```
Get-WindowsOptionalFeature -Online -FeatureName MultiPathIO
```

- b. Run the following command to enable MPIO:

```
Enable-WindowsOptionalFeature -Online -FeatureName MultiPathIO
```

3. Run the following commands to set *vid/pid*:

```
New-MSDSMSupportedHW -VendorId DelleMC -ProductId PowerStore
```

4. Run the following commands to set *RoundRobin* failover policy:

```
Set-MSDSMGlobalDefaultLoadBalancePolicy -Policy RR
```

5. Run the following commands to set MPIO timeout values:

```
Set-MPIOSetting -NewPathVerificationState Enabled
```

```
Set-MPIOSetting -NewPathVerificationPeriod 30
```

```
Set-MPIOSetting -NewPDORemovePeriod 20
```

```
Set-MPIOSetting -NewRetryCount 3
```

```
Set-MPIOSetting -NewRetryInterval 3
```

```
Set-MPIOSetting -custompathrecovery enabled
```

```
Set-MPIOSetting -newpathrecoveryinterval 10
```

```
Set-MPIOSetting -NewDiskTimeout 30
```

## Configure PowerPath multipathing with PowerStore

PowerStore supports multipathing using Dell EMC PowerPath on Windows. PowerPath versions 6.5 and later provide supports on PowerStore array.

For details about installing and configuring PowerPath with PowerStore support on your host, see the *Dell MC PowerPath on Windows Installation and Administration Guide* for the PowerPath version you are installing. This guide provides the required information for placing PowerStore volumes under PowerPath control.

**Note:** The PowerPath with PowerStore support installation procedure is fully storage-aware. All required PowerPath settings with PowerStore storage are automatically done when PowerPath is installed on your host. This includes settings such as the PowerPath multipathing policy that does not require manual setting.


## Provisioning PowerStore LUNs to Windows host

Use PowerStore Web Manager to provision LUNs to Windows hosts.

### Procedure

- For each PowerStore system connected to Windows hosts, use PowerStore Web Manager map LUNs to add the host or host group (for shared access like clustering):

- Click **Compute > Hosts and Host Groups > ADD HOST**.



- Enter the hostname and select **Windows** as operating system.

Add Host ×

Host Details      Host Protocol      Host Initiators      Summary

Give the host a friendly name, provide the operating system type, and optionally provide a description.

Host Friendly Name:

Description (Optional):

Operating System:

- Select **iSCSI** or **Fibre Channel** as Protocol Type on the **Host Protocol** page.

Add Host ×

Host Details      Host Protocol      Host Initiators      Summary

Select the protocol you want to use for the host connection.

Protocol Type: ☒ iSCSI ☐ Fibre Channel

Auto Discovery of iSCSI Initiators

Initiators are automatically discovered, if they are configured on the host side to use the cluster's iSCSI Auto Discovery IP address (172.16.11.40).  
Click Next to see the discovered initiators.

- Select initiator's WWPNs or add iSCSI iqn on the **Host Initiators** page.

Add Host ×

Host Details Host Protocol **Host Initiators** Summary

Select the initiators you want to use for configuring host access to the storage in this cluster. You can select from the automatically discovered initiators. To view a list of all available initiators, ensure that the FC embedded modules are configured properly and connected to the cluster.

132 Items Filter Sort Refresh

<input type="checkbox"/> Identifier	Origin
<input type="checkbox"/> 10:00:00:00:c9:71:7c:2e	
<input type="checkbox"/> 10:00:00:00:c9:71:7c:2f	
<input type="checkbox"/> 10:00:00:00:c9:71:be:d0	
<input type="checkbox"/> 10:00:00:00:c9:71:be:d1	
<input type="checkbox"/> 10:00:00:00:c9:b0:74:84	

CANCEL BACK NEXT

e. Review the summary and click **ADD HOST**.

Add Host ×

Host Details Host Protocol Host Initiators **Summary**

Review host details about the new host and when ready, click ADD HOST button.

Properties		Protocol	Initiators
Name	Windows	Selected	10:00:00:00:c9:71:7c:2e
OS	Test	Fibre	10:00:00:00:c9:71:7c:2f
Description	Windows	Protocol	Channel

CANCEL BACK ADD HOST

f. Repeat step a to step e to add multiple hosts and then add Host Group with added hosts if necessary.

2. Create volumes (LUNs) and map to created host:

a. Select **Storage > Volumes**, and click **CREATE**.

Dashboard Monitoring Compute **Storage** Protection Migration Hardware

## Volumes

+ CREATE MODIFY DELETE MORE ACTIONS

932 Volumes

b. On the **Properties** page, enter volume name (Or Prefix), quantity, and size.

## Create Volumes



## Properties

## Host Mappings

## Summary

## General

Name (Or Prefix)

Windows\_Volume\_Test

Description (Optional)

Quantity

10

Size

100

GB

## Additional Properties

Associated Volume Group (Optional)

None Selected

SELECT

Volume Protection Policy (Optional)

None

Volume Performance Policy

Medium

CANCEL

FINISH

NEXT

If necessary, you can also choose Volume Group, Protection Policy or Performance Policy for your volumes.

- c. On the **Host Mappings** page, select the Host or Host Group to map the volumes.

You can also choose to manually provide a LUN for the host or generate a LUN automatically.

The screenshot shows the Dell EMC PowerStore interface. On the left, the 'Hosts & Host Groups' table lists various hosts. On the right, the 'Map Volume' dialog is open, showing a list of available volumes to be mapped to the selected host.

Name	Logical Used	Provisioned
[Redacted]	0 GB	50.0 GB
[Redacted]	0 GB	15.0 GB
[Redacted]	12.2 GB	400.0 GB
[Redacted]	7.4 GB	400.0 GB
[Redacted]	0 GB	5.0 GB
[Redacted]	0 GB	5.0 GB
[Redacted]	0 GB	5.0 GB

- d. Review the **Summary** page and click **CREATE**.

3. Check if the volumes are successfully created and mapped:

- a. Go to **Compute > Hosts and Host Groups**, and click the created Host/Host Group.

- b. Select the **MAPPED VOLUMES** tab and ensure that all volumes are correctly created.

The screenshot shows the 'Hosts > Windows\_Test' page. The 'MAPPED VOLUMES' tab is selected, displaying a table of mapped volumes.

Name	Description	WWN	Logical Used	Provisioned	Protection P.	Performanc.	Appliance
automation-6c53e9f5-28d0-47f1-b2f1-19fda42955a8	--	naa.68cfc098009e540905fbc05a401c0e8	66.7 MB	5.0 GB	--	Medium	FINM006C
automation-9caa06de-f122-4b31-9692-37805e7847a1	--	naa.68cfc098009e5454ce534260a767c713	66.7 MB	5.0 GB	--	Medium	FINM006C

- c. You can also use Export the table option to export the details to csv or xlsx.
4. Scan host paths to discover the newly mapped volumes:  
Scan the host ports using the **Computer Management > Storage > Disk Management > Rescan Disks**.

## LUN ownership

PowerStore LUN ownership is set at LUN creation automatically. It does not support Explicit Failover, so you cannot change LUN ownership manually.

## Reclaim\Unmap

PowerStore currently does not support Reclaim\Unmap feature on Windows Server.

# VPLEX

## Overview

For detailed information about VPLEX, See documentation available at [Dell EMC Online Support](#).

## Documentation

See the following documents for configuration and administration operations:

- *Dell EMC VPLEX with GeoSynchrony 6.x Product Guide*
- *Dell EMC VPLEX with GeoSynchrony 6.x CLI Guide*
- *Dell EMC VPLEX with GeoSynchrony 6.x Configuration Guide*
- *Dell EMC VPLEX Hardware Installation Guide*
- *Dell EMC VPLEX Release Notes*
- *Implementation and Planning Best Practices for Dell EMC VPLEX Technical Notes*
- *VPLEX online help*, available on the Management Console GUI
- *VPLEX Procedure Generator*, available at [Dell EMC Online Support](#)
- *Dell EMC Simple Support Matrix, Dell EMC VPLEX and GeoSynchrony*, available on [Dell EMC E-Lab Navigator](#).

For the up-to-date support information, See [Dell EMC E-Lab Navigator](#).

## Prerequisites

Before configuring VPLEX in the Windows environment, complete the following on each host:

- Confirm that all necessary remediation has been completed. This ensures that OS-specific patches and software on all hosts in the VPLEX environment are at supported levels according to the [Dell EMC E-Lab Navigator](#).
- Confirm that each host is running VPLEX-supported failover software and has at least one available path to each VPLEX fabric.



**Note:** See [Dell EMC E-Lab Navigator](#) for the most up-to-date support information and prerequisites.



- If a host is running PowerPath, confirm that the load-balancing and failover policy is set to **Adaptive**.

## Provisioning and exporting storage

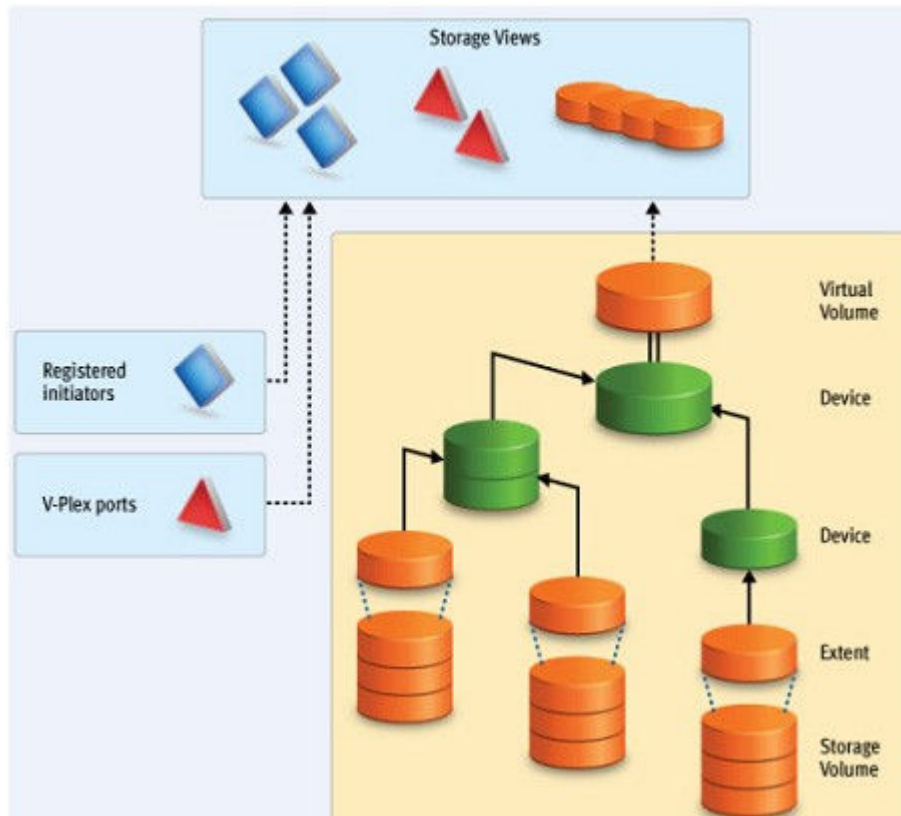
### VPLEX with GeoSynchrony v4.x

To begin using VPLEX, you must provision and export storage so that hosts and applications can use the storage. Storage provisioning and exporting refers to the following tasks required to take a storage volume from a storage array and make it visible to a host:

#### Procedure

1. Discover available storage.
2. Claim and name storage volumes.
3. Create extents from the storage volumes.
4. Create devices from the extents.
5. Create virtual volumes on the devices.
6. Create storage views to allow hosts to view specific virtual volumes.
7. Register initiators with VPLEX.
8. Add initiators (hosts), virtual volumes, and VPLEX ports to the storage view. You can provision storage using the GUI or the CLI. For more information, see *Dell EMC VPLEX Management Console Help* or the *Dell EMC VPLEX CLI Guide*, located on [Dell EMC Online Support](#).

The following figure shows the provisioning and exporting process:

**Figure 22** Provisioning and exporting process

### VPLEX with GeoSynchrony v5.x

VPLEX allows easy storage provisioning among heterogeneous storage arrays. After a storage array LUN volume is encapsulated within VPLEX, all of its block-level storage is available in a global directory and coherent cache. Any front-end device that is zoned properly can access the storage blocks.

Two methods are available for provisioning: EZ provisioning and Advanced provisioning. For more information, see the *Dell EMC VPLEX with GeoSynchrony 5.5 Product Guide* on [Dell EMC Online Support](#).

### VPLEX with GeoSynchrony v6.x

VPLEX provides easy storage provisioning among heterogeneous storage arrays. Use the web-based GUI to simplify everyday provisioning or to create complex devices.

Use the following ways to provision storage in VPLEX:

- Integrated storage provisioning (VIAS-VPLEX Integrated Array Services based provisioning)
- EZ provisioning
- Advanced provisioning

All provisioning features are available in the Unisphere for VPLEX GUI.

For more information, see the *Dell EMC VPLEX with GeoSynchrony 6.1 Product Guide* on [Dell EMC Online Support](#).

## Storage volumes

A storage volume is a LUN exported from an array. When an array is discovered, the storage volumes view shows all exported LUNs on that array. You must claim, and optionally name, these storage volumes before you can use them in a VPLEX cluster. Once claimed, you can divide a storage volume into multiple extents (up to 128), or you can create a single full-size extent using the entire capacity of the storage volume.

**Note:** To claim storage volumes, the GUI supports only the Claim Storage wizard, which assigns a meaningful name to the storage volume. Meaningful names help you associate a storage volume with a specific storage array and LUN on that array and are useful during troubleshooting and performance analysis.

This section contains the following topics:

- Claiming and naming storage volumes
- Extents
- Devices
- Distributed devices
- Rule sets
- Virtual volumes

### Claiming and naming storage volumes

You must claim storage volumes before you can use them in the cluster (except the metadata volume, which is created from an unclaimed storage volume). Only after claiming a storage volume you can use it to create extents, devices, and then virtual volumes.

### Extents

An extent is a slice (range of blocks) of a storage volume. You can create a full-size extent using the entire capacity of the storage volume, or you can carve the storage volume up into several contiguous slices. Extents are used to create devices and then virtual volumes.

### Devices

Devices combine extents or other devices into one large device with specific RAID techniques, such as mirroring or striping. Devices can only be created from extents or other devices. A device's storage capacity is not available until you create a virtual volume on the device and export that virtual volume to a host.

You can create only one virtual volume per device. There are two types of devices:

- Simple device - A simple device is configured using one component, which is an extent.
- Complex device - A complex device has more than one component, combined using a specific RAID type. The components can be extents or other devices (both simple and complex)

### Distributed devices

Distributed devices are configured using storage from both clusters and are only used in multi-cluster plexes. A distributed device's components must be other devices and those devices must be created from storage in different clusters in the plex.

### Rule sets

Rule sets are predefined rules that determine how a cluster behaves when it loses communication with the other cluster, for example, during an inter-cluster link failure or cluster failure. In these situations, until communication is restored, most I/O workloads require specific sets of virtual volumes to resume on one cluster and remain suspended on the other cluster.

VPLEX provides a Management Console on the management server in each cluster. You can create distributed devices using the GUI or CLI on either management server. The default rule set used by

the GUI causes the cluster used to create the distributed device to detach during communication problems, allowing I/O to resume at the cluster. For more information about creating and applying rule sets, see the *Dell EMC VPLEX CLI Guide* on [Dell EMC Online Support](#).

There are cases in which all I/O must be suspended resulting in a data unavailability. VPLEX with functionality of VPLEX Witness.

VPLEX with functionality of VPLEX Witness: When a VPLEX Metro configuration is augmented by VPLEX Witness, the resulting configuration provides the following features:

- High availability for applications in a VPLEX Metro configuration (no single points of storage failure)
- Fully automatic failure handling in a VPLEX Metro configuration
- Improved failure handling in a VPLEX configuration
- Better resource utilization

For information about VPLEX Witness, see the *Dell EMC VPLEX with GeoSynchrony 5.5, 6.1 Product Guide* on [Dell EMC Online Support](#).

### Virtual volumes

Virtual volumes are created on devices or distributed devices and presented to a host through a storage view. You can create virtual volumes only on top-level devices and always use full capacity of the device.

## System volumes

VPLEX stores configuration and metadata on system volumes that are created from storage devices. There are two types of system volumes:

- Metadata volumes
- Logging volumes

Each of these volumes is briefly discussed in the following sections:

### Metadata volumes

VPLEX maintains its configuration state, referred as metadata, on storage volumes provided by storage arrays. Each VPLEX cluster maintains its own metadata, which describes the local configuration information for this cluster and any distributed configuration information that is shared between clusters.

For more information about metadata volumes for VPLEX with GeoSynchrony v4.x, see *Dell EMC VPLEX CLI Guide*, on [Dell EMC Online Support](#).

For more information about metadata volumes for VPLEX with GeoSynchrony v5.x, see *Dell EMC VPLEX with GeoSynchrony 5.0 Product Guide*, on [Dell EMC Online Support](#).

For more information about metadata volumes for VPLEX with GeoSynchrony v6.x, see the *Dell EMC VPLEX with GeoSynchrony 5.0 Product Guide*, on [Dell EMC Online Support](#).

### Logging volumes

Logging volumes are created during initial system setup. It is required in each cluster to track any blocks written during a loss of connectivity between clusters. After an inter-cluster link is restored, the logging volume is used to synchronize distributed devices by sending only changed blocks over the inter-cluster link.

For more information about logging volumes for VPLEX with GeoSynchrony v4.x, see *Dell EMC VPLEX CLI Guide*, on [Dell EMC Online Support](#).

For more information about logging volumes for VPLEX with GeoSynchrony v5.x, see *Dell EMC VPLEX with GeoSynchrony 5.0 Product Guide*, on [Dell EMC Online Support](#).

For more information about logging volumes for VPLEX with GeoSynchrony v6.x, see *Dell EMC VPLEX with GeoSynchrony 5.0 Product Guide*, on [Dell EMC Online Support](#).

## Required storage system setup

Product documentation and installation procedures for connecting a VMAX, Symmetrix, Unity, VNX series, and CLARiiON storage system to a VPLEX instance are available on [Dell EMC Online Support](#).

## Required VMAX series FA bit settings

For VMAX series to VPLEX-connections, configure the VMAX series FC directors (FAs) as shown in following table:

**Note:** Dell EMC recommends that you download the latest information before installing any server.

**Table 4** Required Symmetrix FA bit settings for connection to VPLEX

Set	Do not set	Optional
SPC-2 Compliance (SPC2)	Disable Queue Reset on Unit Attention (D)	Linkspeed
SCSI-3 Compliance (SC3)	AS/400 Ports Only (AS4)	Enable Auto-Negotiation (EAN)
Enable Point-to-Point (PP)	Avoid Reset Broadcast (ARB)	VCM/ACLX <sup>a</sup>
Unique Worldwide Name (UWN)	Environment Reports to Host (E)	
Common Serial Number (C)	Soft Reset (S)	
	Open VMS (OVMS)	
	Return Busy (B)	
	Enable Sunapee (SCL)	
	Sequent Bit (SEQ)	
	Non-Participant (N)	
	OS-2007 (OS compliance)	

- a. You must set VCM/ACLX bit, if VPLEX is sharing VMAX series directors with hosts that require conflicting bit settings. For any other configuration, the VCM/ACLX bit can be either set or not set.

**Note:** When setting up a VPLEX-attach version 4.x or earlier with a VNX series or CLARiiON system, you must set the initiator type to CLARiiON Open and Failover Mode to 1. ALUA is not supported.

When setting up a VPLEX-attach version 5.0 or later with a VNX series or CLARiiON system, the initiator type can be set to CLARiiON Open and the Failover Mode set to 1 or Failover Mode 4 since ALUA is supported.

If you are using the LUN masking, set the VCM/ACLX flag. You must use VCM/ACLX, if sharing array directors with hosts which require conflicting flag settings.

**Note:** The FA bit settings that are listed in [Table 4](#) are for connectivity of VPLEX to Dell EMC VMAX series only. For host to Dell EMC VMAX series FA bit settings, see the [Dell EMC E-Lab Navigator](#).

## Supported storage arrays

For lists the storage arrays that are qualified to use with VPLEX, see the [Dell EMC Simple Support Matrix](#).

See *VPLEX Procedure Generator*, on [Dell EMC Online Support](#), to verify supported storage arrays.

VPLEX automatically discovers storage arrays that are connected to the back-end ports. All arrays connected to each director in the cluster are listed in the storage array view.

## Initiator settings on back-end arrays

See the *VPLEX Procedure Generator*, on [Dell EMC Online Support](#), to verify the initiator settings for storage arrays when configuring the arrays for use with VPLEX.

## Host connectivity

For Windows host connectivity recommendations and best practices with VPLEX configurations, see *Implementation and planning best practices for EMC VPLEX technical notes*, available on [Dell EMC Online Support](#).

For the most up-to-date information about qualified switches, hosts, HBAs, and software, see the [Dell EMC Simple Support Matrix](#), or contact your Dell EMC Customer Support.

The latest Dell EMC-approved HBA drivers and software are available for download at the following websites:


- [www.broadcom.com](http://www.broadcom.com)
- [www.QLogic.com](http://www.QLogic.com)
- [www.brocade.com](http://www.brocade.com)

The Dell EMC HBA installation and configurations guides are available at the Dell EMC download pages of these websites.

 **Note:** Direct connect from an HBA to a VPLEX engine is not supported.


## Configuring Fibre Channel HBAs

This section provides detail about Fibre Channel (FC) Host Bus Adapter (HBA) configuration.

 **Note:** The values provided are required and optimal for most scenarios. However, in host IO profiles with large-block reads, the values may need to be tuned if the performance of the VPLEX shows high front-end latency in the absence of high back-end latency, which has visible impact on host application(s). This may be an indication that there are too many outstanding IOs at a given time per port. If the recommended settings do not perform well in the environment, contact Dell EMC Support for additional recommendations.

For more information on how to monitor VPLEX performance, see the *VPLEX Administration Guide*. If host application(s) is seeing a performance issue with the required settings, contact Dell EMC Support for further recommendations.

## Setting queue depth and execution throttle for QLogic

 **CAUTION** Changing the HBA queue depth is designed for advanced users. Increasing the queue depth may cause hosts to over-stress other arrays connected to the Windows host, resulting in performance degradation while performing IO.

The execution throttle setting controls the amount of outstanding I/O requests per HBA port. The HBA execution throttle should be set to the QLogic default value, which is 65535. This can be done at the HBA firmware level using the HBA BIOS or the QConvergeConsole CLI or GUI.

The queue depth setting controls the amount of outstanding I/O requests per a single path. On Windows, the HBA queue depth can be adjusted using the Windows Registry.

**Note:** When the execution throttle in the HBA level is set to a value lower than the queue depth, it might limit the queue depth to a lower value than the set value.

The following procedures detail how to adjust the queue depth setting for QLogic HBAs:

- [Setting the queue depth for the QLogic FC HBA](#)
- [Setting the execution throttle on the QLogic FC HBA](#)
- [Setting the queue depth and queue target on the Emulex FC HBA](#)

Follow the procedure according to the HBA type. For any additional information, see the HBA vendor's documentation.

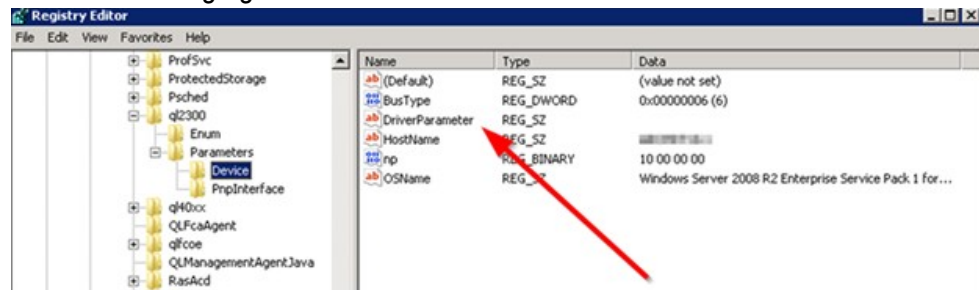
## Setting the queue depth for the QLogic FC HBA

### Procedure


1. On the desktop, click Start, select Run, and open the REGEDIT (Registry Editor).

**Note:** Some driver versions do not create the registry by default. In those cases, the user must manually create the registry.

2. Select **HKEY\_LOCAL\_MACHINE** and follow the tree structure down to the QLogic driver as shown in the following figure and double-click **DriverParameter**.



The Edit String dialog box displays:




3. Change the value of qd to 20. The value is set in hexadecimal; 20 is 32 in decimal.

If additional driver parameters are already set, and the string qd= does not exist, append the following text to the end of the string using a semicolon (";") to separate the new queue depth value from previous entries: **;qd=20**

4. Click **OK**.

The Registry Editor appears as follows:





5. Exit the Registry Editor and reboot the Windows host.

## Setting the execution throttle on the Qlogic FC HBA

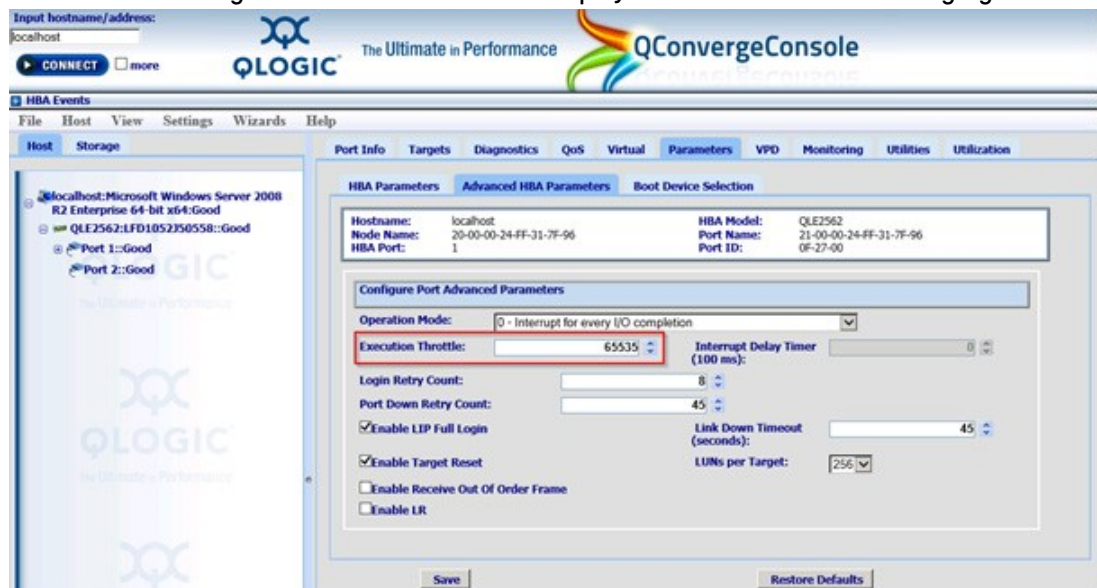
### Procedure

1. Install the QConvergeConsole GUI or CLI.
2. Follow the GUI or CLI directions, outlined in this section to set the execution throttle.

### Using QConvergeConsole GUI

#### Procedure

1. Start the QConvergeConsole GUI. The GUI is displayed as shown in the following figure:



2. Select one of the adapter ports in the navigation tree on the left.
3. Select **Host > Parameters > Advanced HBA Parameters**.
4. Set the Execution Throttle to 65535.
5. Click **Save**.
6. Repeat the steps for each port on each adapter connecting to VPLEX.

### Using QConvergeConsole CLI

#### Procedure

1. Select **2: Adapter Configuration** from the main menu.



```

QConvergeConsole CLI
Working dir: C:\Program Files\QLogic Corporation\QConvergeConsoleCLI
Using config file: C:\Program Files\QLogic Corporation\QConvergeConsoleCLI\netsec
li.cfg
Loading CNA Data ...
Scanning for QLogic adapters, please wait...
Scanning for QLogic adapters, please wait...
Scanning for QLogic adapters, please wait...

QConvergeConsole
CLI - Version 1.1.0 <Build 51>

Main Menu
1: Adapter Information
2: Adapter Configuration
3: Adapter Updates
4: Adapter Diagnostics
5: Adapter Statistics
6: Refresh
7: Help
8: Exit

Please Enter Selection: 2

```

2. Select 3. HBA Parameters:

```

QConvergeConsole CLI

Please Enter Selection: 2

QConvergeConsole
CLI - Version 1.1.0 <Build 51>

Fibre Channel Adapter Configuration
1: Adapter Alias
2: Adapter Port Alias
3: HBA Parameters
4: Target Persistent Binding
5: Boot Devices Configuration
6: Virtual Ports <NPIV>
7: Target Link Speed <iidMA>
8: Driver Parameters
9: Selective LUNs
10: QoS
11: Export <Save> Configuration
12: Generate Reports

```

3. Select the HBA (1: Port 1 in the following example).

```

QConvergeConsole CLI

6: Virtual Ports <NPIV>
7: Target Link Speed <iidMA>
8: Driver Parameters
9: Selective LUNs
10: QoS
11: Export <Save> Configuration
12: Generate Reports

<p or 0: Previous Menu; m or 98: Main Menu; ex or 99: Quit>
Please Enter Selection: 3

QConvergeConsole
CLI - Version 1.1.0 <Build 51>

Fibre Channel Adapter Configuration

HBA Model QLE2562 SN: LFD1052J50558
1: Port 1: WWPN: 21-00-00-24-FF-31-7F-96 Online
2: Port 2: WWPN: 21-00-00-24-FF-31-7F-97 Online

<p or 0: Previous Menu; m or 98: Main Menu; ex or 99: Quit>
Please Enter Selection: 1

```

4. Select 2: Configure HBA Parameters.

```

QConvergeConsole CLI

QConvergeConsole
Version 1.1.0 <Build 51>

HBA Parameters Menu
=====
HBA      : 0 Port: 1
SN       : LFD1052J50558
HBA Model : QLE2562
HBA Desc. : QLE2562 PCI Express to 8Gb FC Dual Channel
FW Version : 4.03.01
WWPN     : 21-00-00-24-FF-31-7F-96
WWNN     : 20-00-00-24-FF-31-7F-96
Link      : Online
=====

1: Display HBA Parameters
2: Configure HBA Parameters
3: Restore Defaults

<p or 0: Previous Menu; m or 98: Main Menu; ex or 99: Quit>
Please Enter Selection: 2

```

5. Select 11: Execution Throttle.

```

QConvergeConsole CLI
1: Connection Options
2: Data Rate
3: Frame Size
4: Enable HBA Hard Loop ID
5: Hard Loop ID
6: Loop Reset Delay <seconds>
7: Enable BIOS
8: Enable Fibre Channel Tape Support
9: Operation Mode
10: Interrupt Delay Timer <100ms>
11: Execution Throttle
12: Login Retry Count
13: Port Down Retry Count
14: Enable LIP Full Login
15: Link Down Timeout <seconds>
16: Enable Target Reset
17: LUNs per Target
18: Enable Receive Out Of Order Frame
19: Enable LR
20: Commit Changes
21: Abort Changes

<p or 0: Previous Menu; m or 98: Main Menu; ex or 99: Quit>
Please Enter Selection: 11

```

6. Set the value to 65535.

**Note:** The current value is in the second set of square brackets. The first is the allowable range.

```

QConvergeConsole CLI
2: Data Rate
3: Frame Size
4: Enable HBA Hard Loop ID
5: Hard Loop ID
6: Loop Reset Delay <seconds>
7: Enable BIOS
8: Enable Fibre Channel Tape Support
9: Operation Mode
10: Interrupt Delay Timer <100ms>
11: Execution Throttle
12: Login Retry Count
13: Port Down Retry Count
14: Enable LIP Full Login
15: Link Down Timeout <seconds>
16: Enable Target Reset
17: LUNs per Target
18: Enable Receive Out Of Order Frame
19: Enable LR
20: Commit Changes
21: Abort Changes

<p or 0: Previous Menu; m or 98: Main Menu; ex or 99: Quit>
Please Enter Selection: 11
Enter Execution Throttle [1-65535] [65535]:

```

7. Verify the options.

```

QConvergeConsole CLI
QConvergeConsole
Version 1.1.0 (Build 51)

HBA Parameters Menu
=====
HBA       : 0 Port: 1
SN        : LFD1052J50558
HBA Model : QLE2562
HBA Desc. : QLE2562 PCI Express to 8Gb FC Dual Channel
FW Version: 4.03.01
WWPN      : 21-00-00-24-FF-31-7F-96
WWNN      : 20-00-00-24-FF-31-7F-96
Link      : Online
=====

1: Display HBA Parameters
2: Configure HBA Parameters
3: Restore Defaults

<p or 0: Previous Menu; m or 98: Main Menu; ex or 99: Quit>
Please Enter Selection: 1

```

8. Validate that the **Execution Throttle** is set to the expected value of **65535**.

```

QConvergeConsole CLI
HBA Instance 0: QLE2562 Port 1 WWPN 21-00-00-24-FF-31-7F-96 PortID 0F-27-00
Link: Online

-----
Connection Options      : 2 - Loop Preferred, Otherwise Point-to-Point
Data Rate               : Auto
Frame Size              : 2048
Hard Loop ID            : 0
Loop Reset Delay <seconds> : 5
Enable Host HBA BIOS    : Disabled
Enable Hard Loop ID     : Disabled
Enable FC Tape Support  : Enabled
Operation Mode          : 0 - Interrupt for every I/O completion
Interrupt Delay Timer <100ms> : 0
Execution Throttle      : 65535
Login Retry Count       : 8
Port Down Retry Count   : 45
Enable LIP Full Login   : Enabled
Link Down Timeout <seconds> : 45
Enable Target Reset     : Enabled
LUNs Per Target         : 256
Out Of Order Frame Assembly : Disabled
Enable LR               : Disabled

Press <Enter> to continue:

```

- Repeat the steps for each port on each adapter connecting to VPLEX.

## Setting queue depth and queue target for Emulex

**CAUTION** Changing the HBA queue depth is designed for advanced users. Increasing the queue depth might cause hosts to overstress other arrays that are connected to the Windows host, resulting in performance degradation while performing IO.

Queue depth setting controls the amount of outstanding I/O requests per a single LUN/target. On Windows, the Emulex HBA queue depth can be adjusted using the Emulex UI (OneCommand).

Queue target controls I/O depth limiting on a per target or per LUN basis. If set to 0 = depth, limitation is applied to individual LUNs. If set to 1 = depth, limitation is applied across the entire target. On Windows, the Emulex HBA queue depth can be adjusted using the Emulex UI (OneCommand).

Adjust the queue depth and queue target settings for Emulex HBAs as follows:

- Set the Emulex HBA adapter queue depth in Windows to 32.
- Set the Emulex HBA adapter queue target in Windows to 0.

**Note:** This means 32 outstanding IOs per ITL. So, if a host has four paths, then there are 32 outstanding IOs per path, resulting in a total of 128.

Follow the procedure as per the HBA type. For any additional information, see the HBA vendor's documentation.

### Setting the queue depth and queue target on the Emulex FC HBA

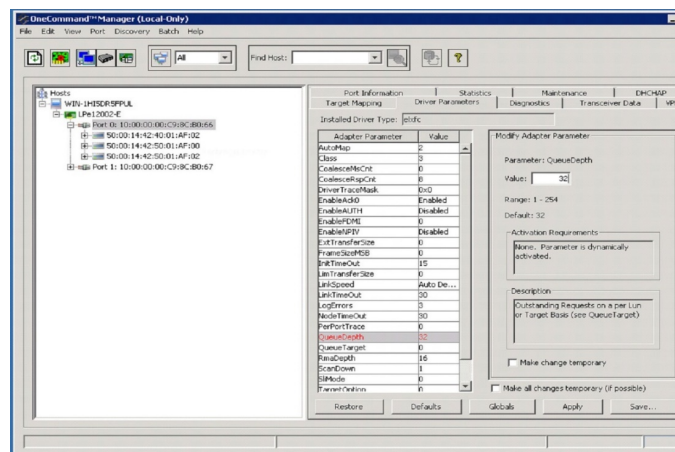
#### About this task

Setting the queue depth on the Emulex FC HBA is done using Emulex UI (OneCommand). OneCommand detects the active Emulex driver and enables changing the relevant driver's parameters, specifically queue depth.

**Note:** Setting the queue-depth as per this procedure is not disruptive.

#### Procedure

- Install OneCommand.
- Launch the OneCommand UI.



- Select the relevant host name from the **Hosts** list.
- Expand the HBA in the navigational tree and select the HBA port.
- Select the **Driver Parameters** tab.
- In the **Adapter Parameter** list, locate the **QueueDepth** parameter and set its value to 32.

7. In the same list, locate the **QueueTarget** parameter and set its value to 0.
8. Click **Apply**.
9. Repeat the steps for each port on the host that has VPLEX Storage exposed.

## Windows Failover Clustering with VPLEX

Microsoft strongly recommends that you validate your configuration before you create a failover cluster, by running all tests in the **Validate a Configuration Wizard**. This will confirm that your hardware and settings are compatible with failover clustering.

- Note:** With Windows Server 2012 or Windows Server 2012 R2, cluster validation storage tests may not discover VPLEX distributed devices when servers are geographically dispersed and configured on different VPLEX sites. This is because a storage validation test selects only shared LUNs. A LUN is determined to be shared if its disk signatures, device identification number (page 0x83), and storage array serial number are the same on all cluster nodes.
- Note:** When you have site-to-site mirroring configured (VPLEX distributed device), a LUN in one site (site A) has a mirrored LUN in another site (site B). These LUNs have the same disk signatures and device identification number (page 0x83), but the VPLEX storage array serial numbers are different. Therefore, they are not recognized as shared LUNs.

The following is an example of what is reported in the cluster validation logs:

Cluster validation message: List Potential Cluster Disks

Description: List disks that will be validated for cluster compatibility. Clustered disks which are online at any node will be excluded.

Start: 11/17/2013 5:59:01 PM.

- Physical disk 84d2b21a is visible from only one node and will not be tested. Validation requires that the disk be visible from at least two nodes. The disk is reported as visible at the following node:  
WNH6-H5.elabqual.emc.com
- Physical disk 6f473a9f is visible from only one node and will not be tested. Validation requires that the disk be visible from at least two nodes. The disk is reported as visible at the following node:  
WNH6-H13.elabqual.emc.com

To resolve the issue, run all the cluster validation tests before you configure distributed devices to geographically dispersed servers.

- Note:** If the validation test is needed later for support situations, LUNs that are not selected for storage validation tests are supported by Microsoft and Dell EMC Shared LUNs (distributed devices).

For more information, see the Microsoft KB article [Storage tests on a failover cluster may not discover all shared LUNs](#).

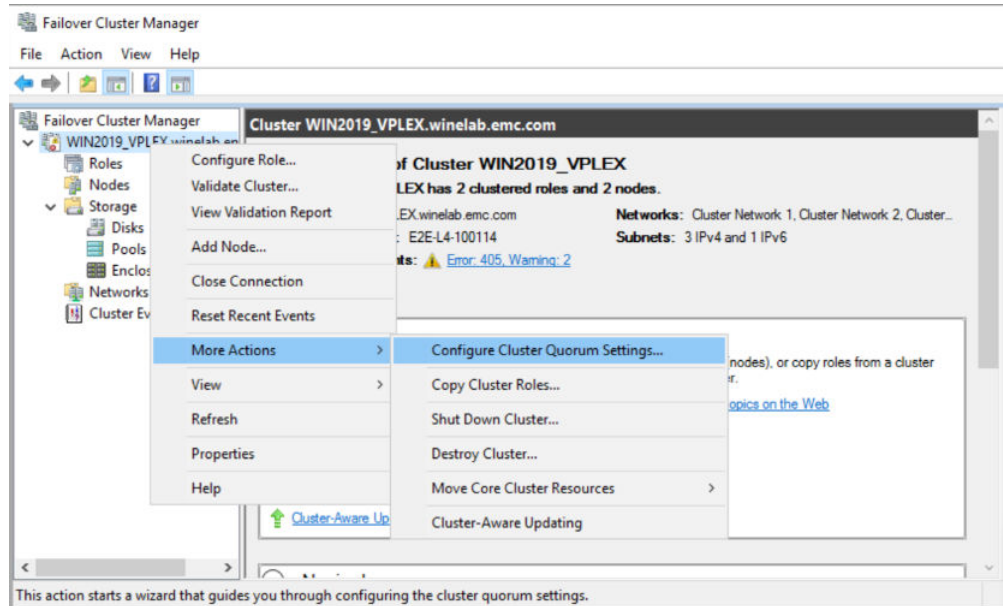
## Setting up quorum on a Windows 2016/2019 Failover Cluster for VPLEX Metro clusters

### About this task

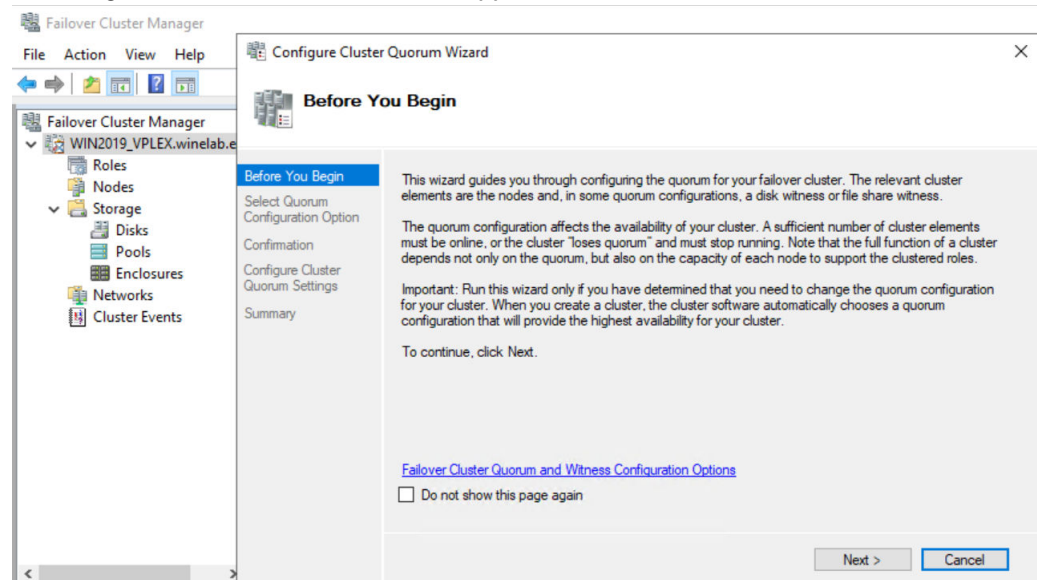
The recommended Windows Failover Clustering Quorum for stretched or cross connected clusters is the File Share Witness. To set up the File Share Witness quorum, complete the following steps.

## Procedure

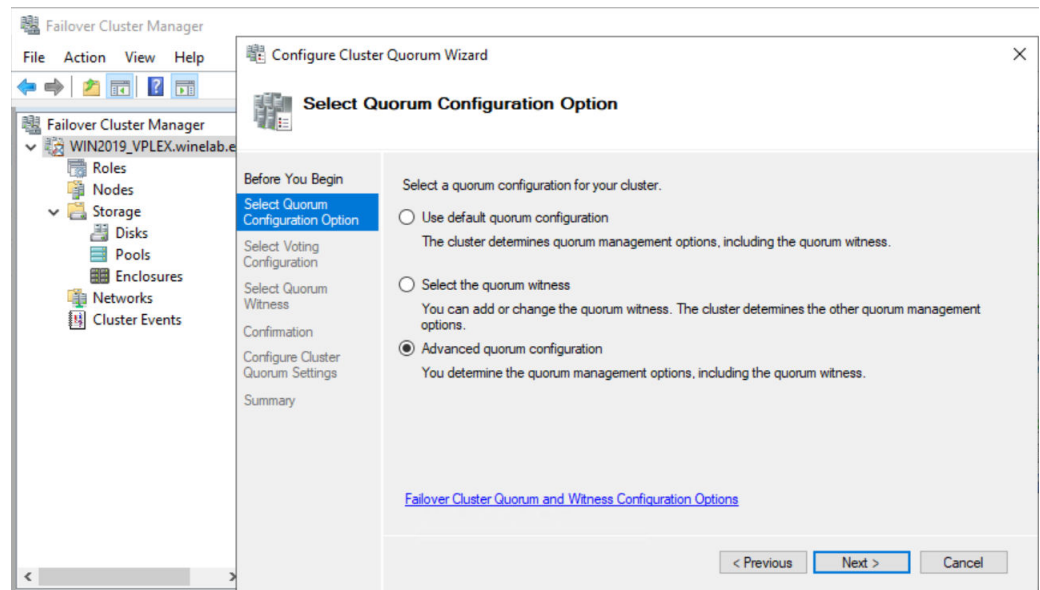
1. In Failover Cluster Manager, select the cluster and from the drop-down menu and select **More Actions > Configure Cluster Quorum Settings**.



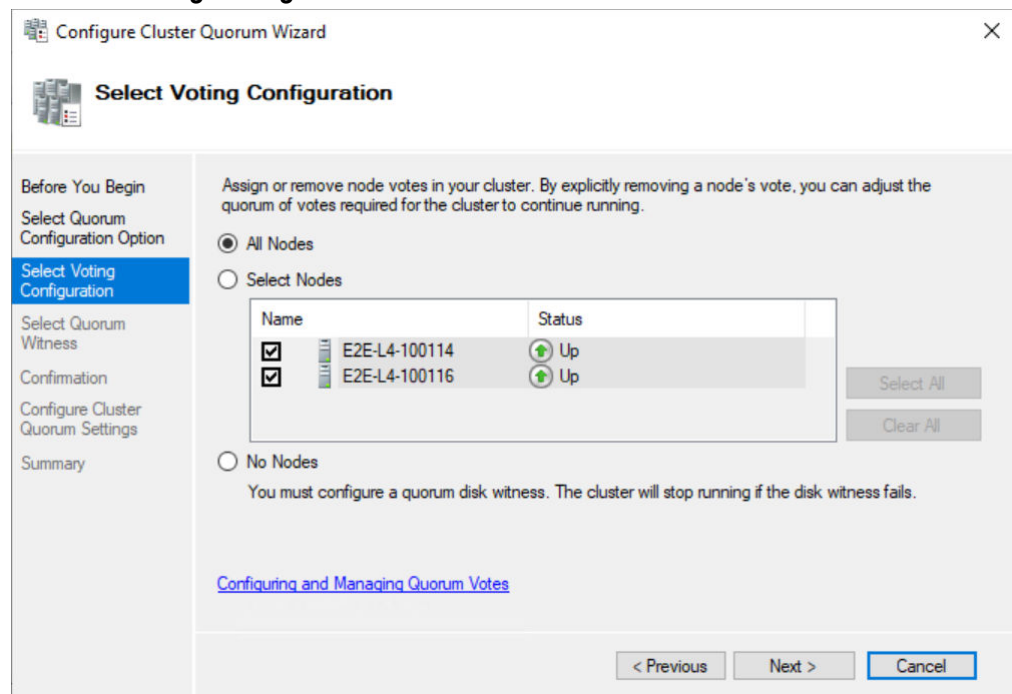
The **Configure Cluster Quorum Wizard** appears as follows:



2. Click **Select Quorum Configuration Option** and choose **Advanced quorum configuration**.

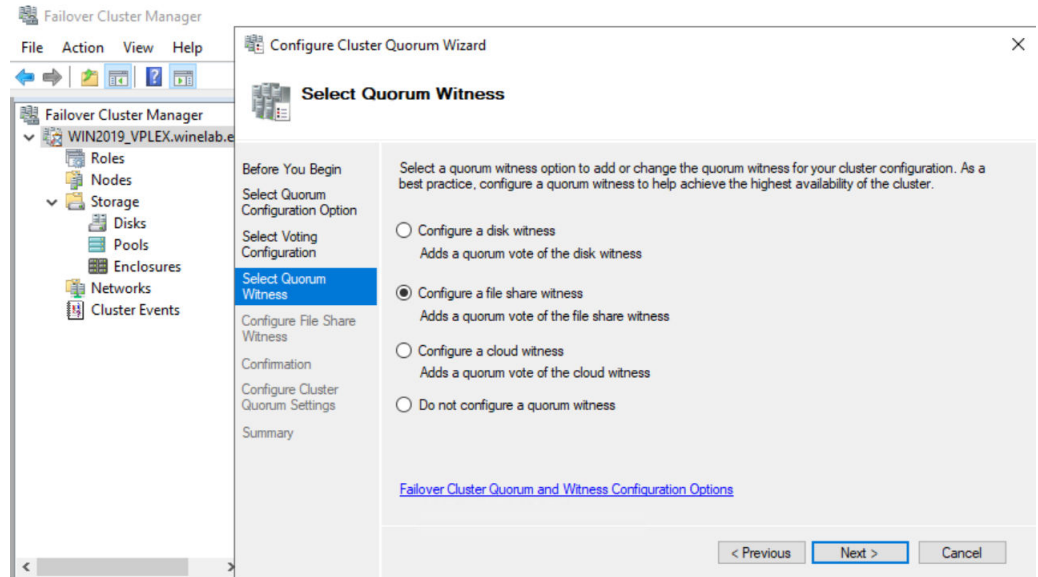


3. Click **Select Voting Configuration** and choose **ALL Nodes**.

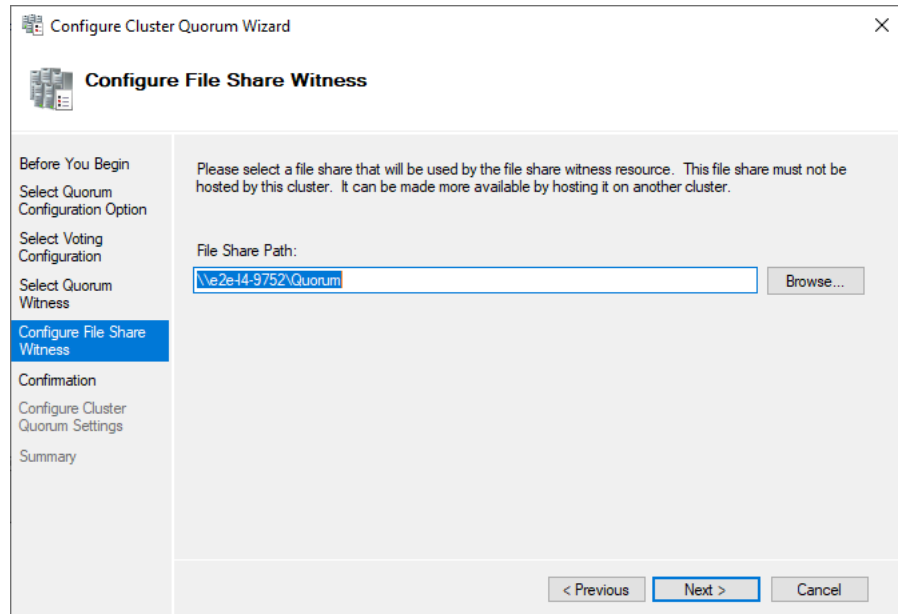


4. Click **Select Quorum Witness** and choose **Configure a file share witness**.

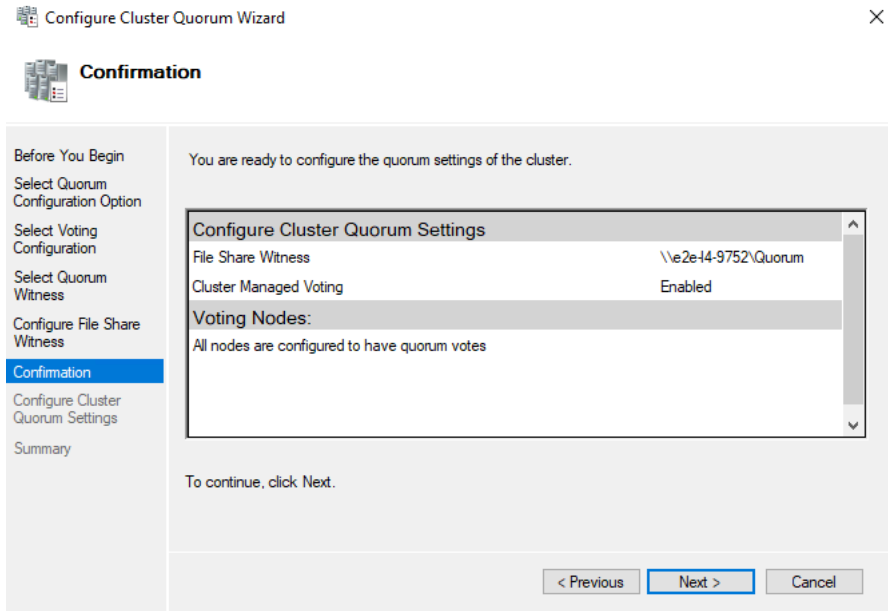




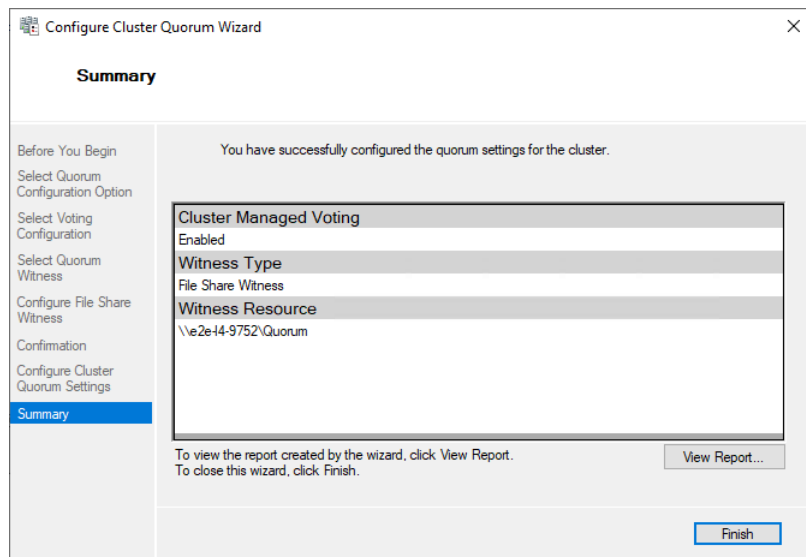
5. Fill the File Share Path and click **Next**.



6. In the **Confirmation** tab, Click **Next**.



## 7. Click **Finish**.



For the server hosting the file share, follow these requirements and recommendations:

- Must have a minimum of 5 MB of free space.
- Must be dedicated to the single cluster and not used to store user or application data.
- Must have write permissions enabled for the computer object for the cluster name.

The following are additional considerations for a file server that hosts the file share witness:

- A single file server can be configured with file share witnesses for multiple clusters.
- The file server must be on a site that is separate from the cluster workload. This enables equal opportunity for any cluster site to survive if site-to-site network communication is lost. If the file server is on the same site, that site becomes the primary site, and it is the only site that can reach the file share.
- The file server can run on a virtual machine if the virtual machine is not hosted on the same cluster that uses the file share witness.



- For high availability, the file server can be configured on a separate failover cluster.

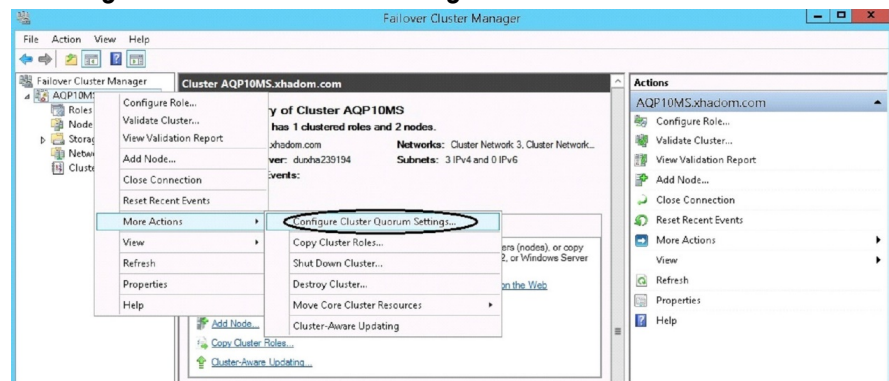
## Setting up quorum on a Windows 2012 / 2012 R2 Failover Cluster for VPLEX Metro clusters

### About this task

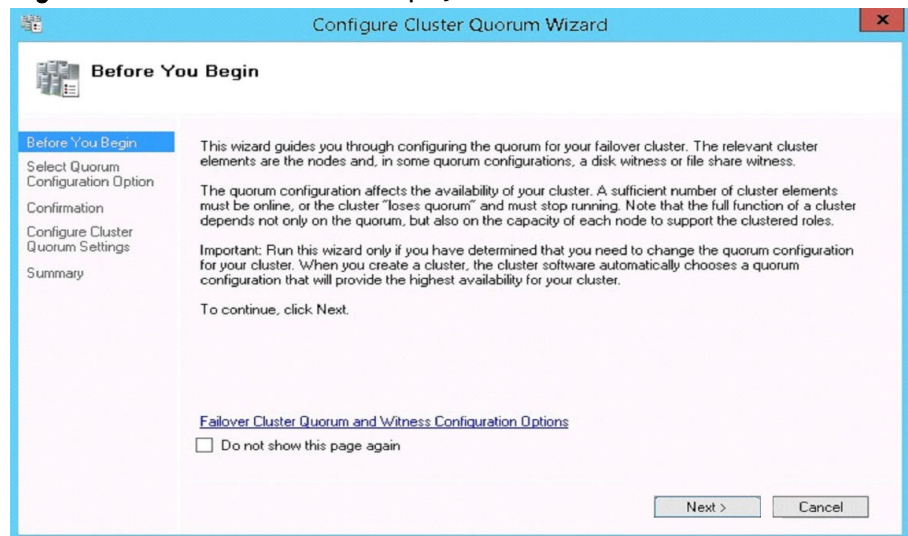
The recommended Windows Failover Clustering Quorum for stretched or cross connected clusters is the File Share Witness. To set up the File Share Witness quorum, complete the following steps.

### Procedure

1. In Failover Cluster Manager, select the cluster and from the drop-down menu select **More Actions > Configure Cluster Quorum Settings**.

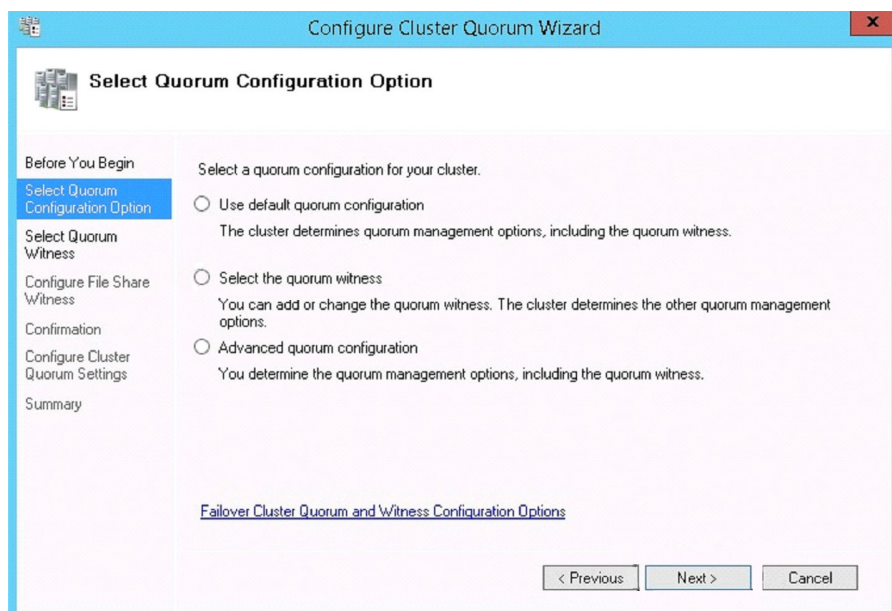


The **Configure Cluster Quorum Wizard** displays:



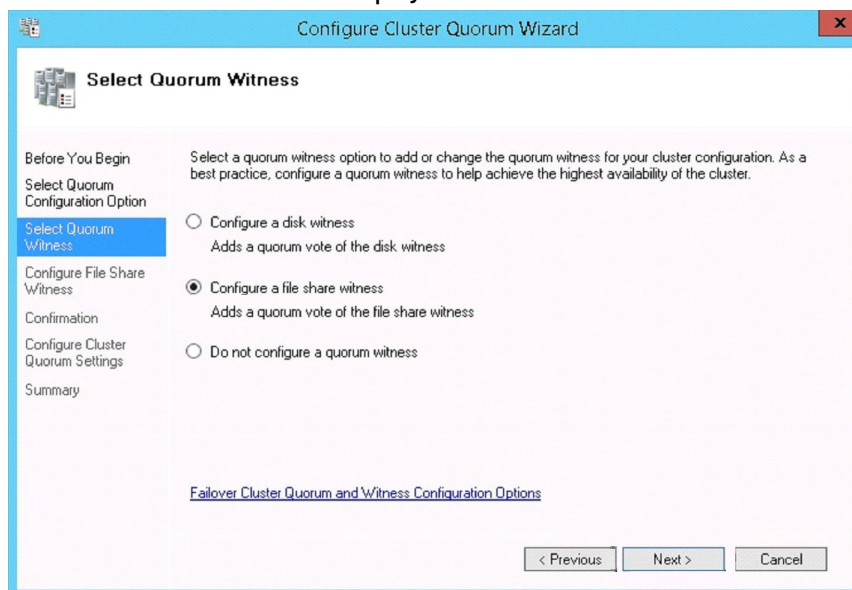
2. Click **Next**.

The **Select Quorum Configuration Option** window displays:



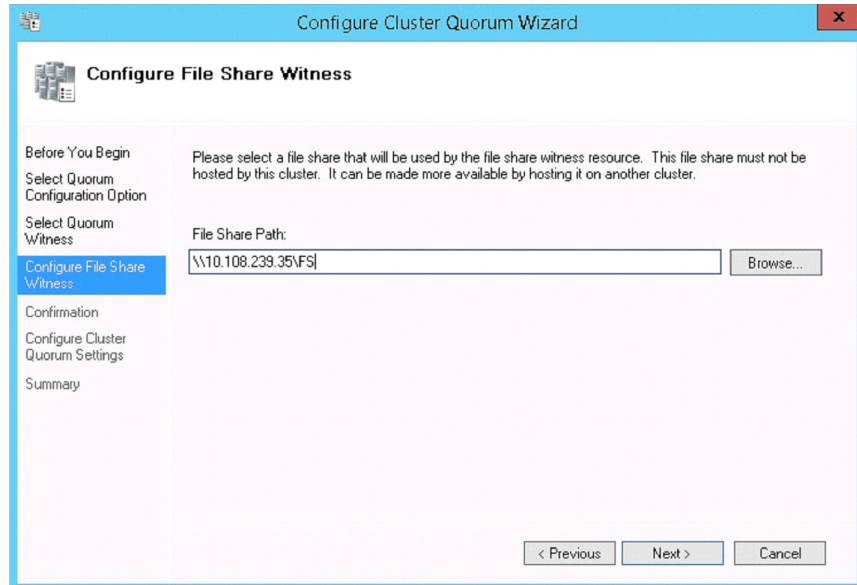
3. Click the **Select Quorum Witness** option and click **Next**.

The **Select Quorum Witness** window displays:



4. Select **Configure a file share witness** and click **Next**.

The **Configure a File Share Witness** window displays:



For the server hosting the file share, follow these requirements and recommendations:

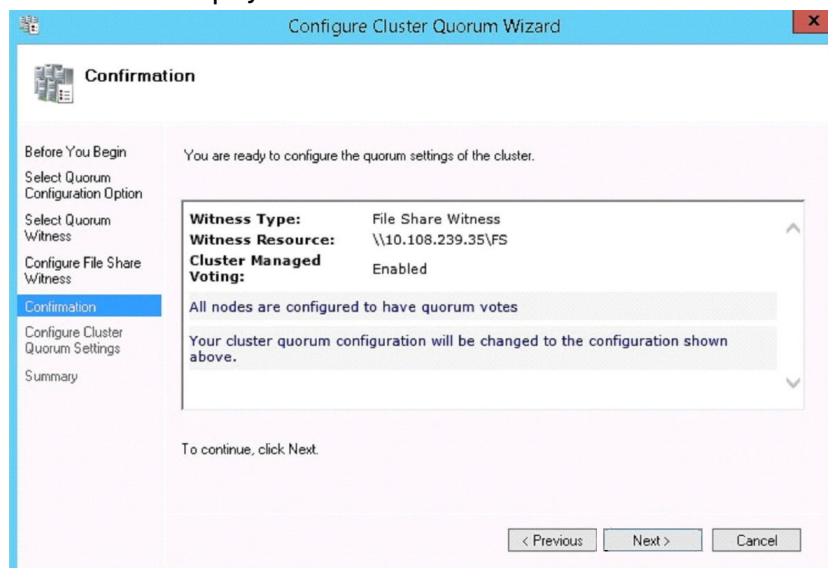
- Must have a minimum of 5 MB of free space
- Must be dedicated to the single cluster and not used to store user or application data
- Must have write permissions enabled for the computer object for the cluster name

The following are additional considerations for a file server that hosts the file share witness:

- A single file server can be configured with file share witnesses for multiple clusters.
- The file server must be on a site that is separate from the cluster workload. This enables equal opportunity for any cluster site to survive if site-to-site network communication is lost. If the file server is on the same site, that site becomes the primary site, and it is the only site that can reach the file share
- The file server can run on a virtual machine if the virtual machine is not hosted on the same cluster that uses the file share witness.
- For high availability, the file server can be configured on a separate failover cluster.

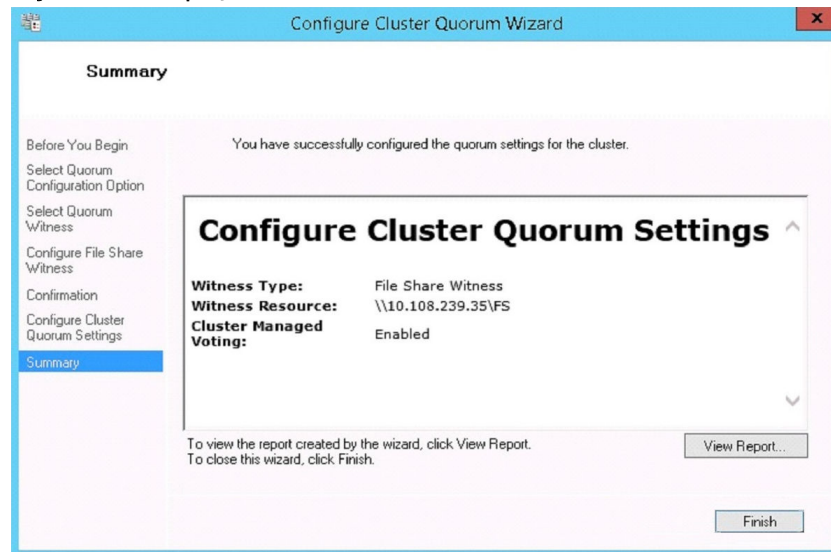
5. Click Next

The **Confirmation** screen displays:



- Verify the settings and click **Next**.

The **Summary** window displays:




- You can view this report or click **Finish** to complete the file share witness configuration.


## Configuring quorum on Windows 2008/2008 R2 Failover Cluster for VPLEX Metro clusters

### VPLEX Metro cluster configuration

Two VPLEX Metro clusters, connected within metro (synchronous) distances of approximately 60 miles (100 kilometers), form a Metro-Plex cluster. Following figure shows an example of a VPLEX Metro cluster configuration. VPLEX Geo cluster configuration is the same and adds the ability to dynamically move applications and data across different compute and storage installations across even greater distances.

**Figure 23** VPLEX Metro cluster configuration example




 **Note:** All connections shown in previous figure are Fiber Channel, except the network connections.

The environment in previous figure consists of the following:

- Node-1: Windows 2008 or Windows 2008 R2 Server connected to the VPLEX instance over Fiber Channel.
- Node-2: Windows 2008 or Windows 2008 R2 Server connected to the VPLEX instance over Fiber Channel.
- VPLEX instance: One or more engine VPLEX having a connection through the L2 switch to back-end and front-end devices.

## Prerequisites

Ensure the following before configuring the VPLEX Metro or Geo cluster:

- VPLEX firmware is installed properly and the minimum configuration is created.
- All volumes to be used during the cluster test should have multiple back-end and front-end paths.
-  **Note:** See the Implementation and Planning Best Practices for Dell EMC VPLEX Technical Notes, available on [Dell EMC Online Support](#), for best practices for the number of paths for back-end and front-end paths.
- All hosts/servers/nodes of the same configuration, version, and service pack of the operating system are installed.
- All nodes are part of the same domain and can communicate with each other before installing Windows Failover Clustering.
- One free IP address is available for cluster IP in the network.
- PowerPath or MPIO is installed and enabled on all the cluster hosts.
- The hosts are registered to the appropriate View and visible to VPLEX.
- All volumes to be used during cluster test should be shared by all nodes and accessible from all nodes.
- A network fileshare is required for cluster quorum.

## Setting up quorum on a Windows 2008/2008R2 Failover Cluster for VPLEX Metro clusters

### About this task

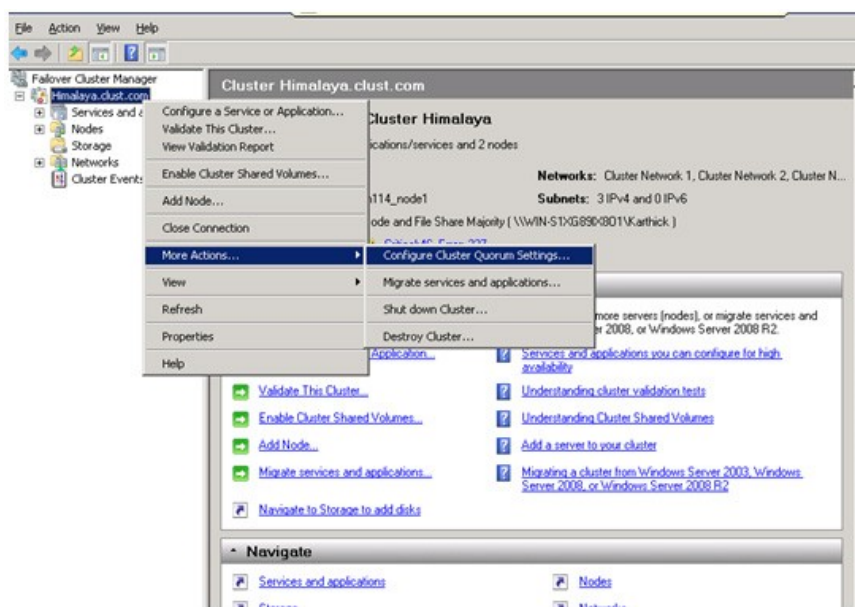
To set up a quorum on VPLEX Metro clusters for Windows Failover Cluster, complete the following steps.

### Procedure

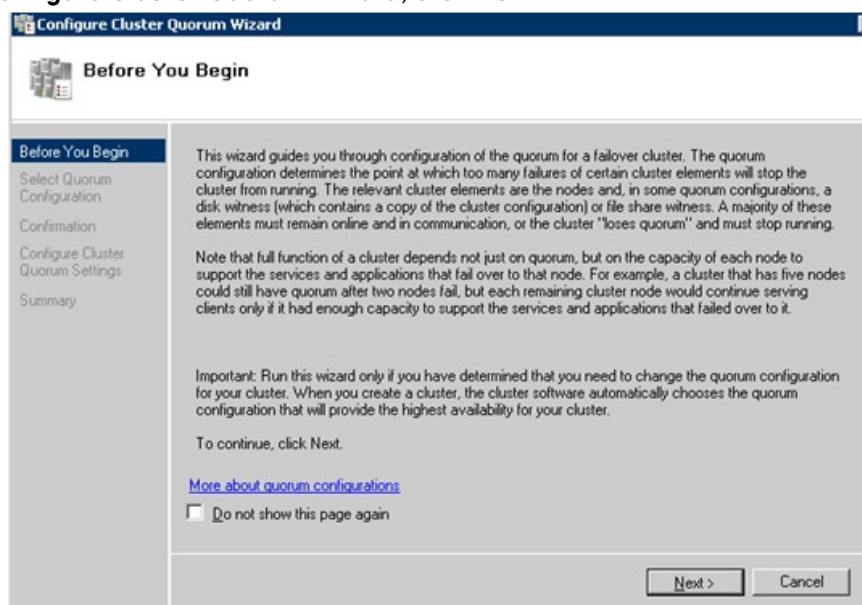
1. Select the quorum settings. In the Failover Cluster Manager, right-click on the cluster name and select **More Actions > Configure Cluster Quorum Settings > Node and File Share Majority**.

The Node and File Share Majority model is recommended for VPLEX Metro and Geo environments.

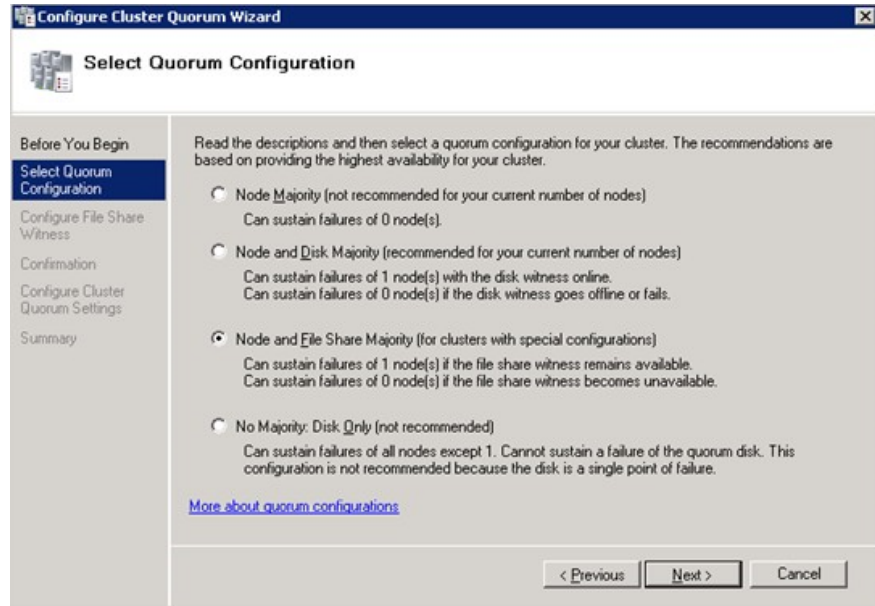




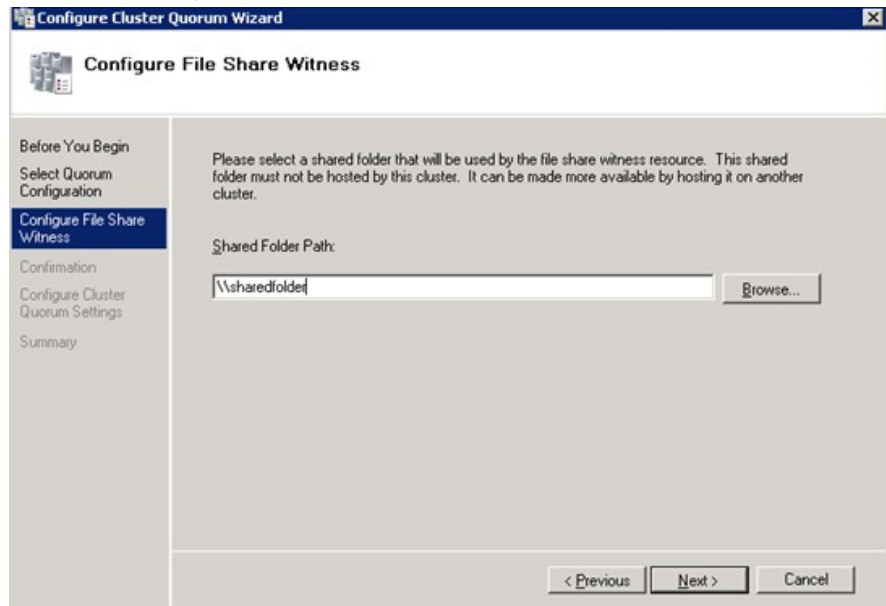
2. In the **Configure Cluster Quorum Wizard**, click **Next**:



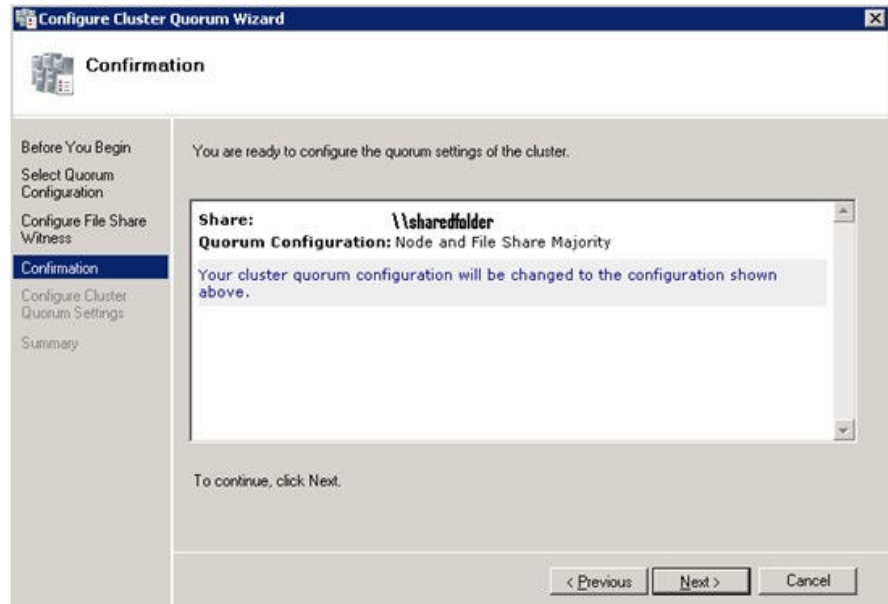
3. In the **Select Quorum Configuration** window, ensure that the **Node and File Share Majority** radio button is selected, and then click **Next**.



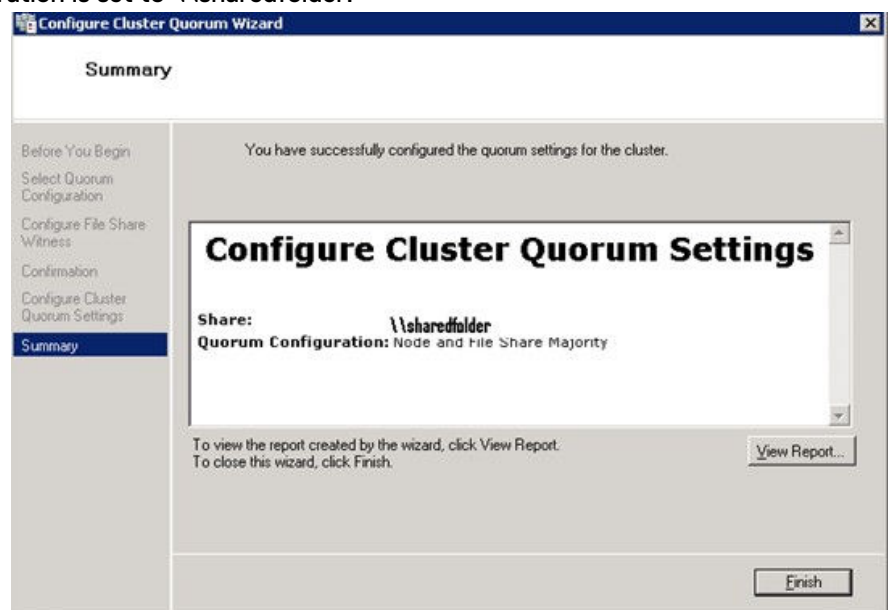
4. 1. In the **Configure File Share Witness** window, ensure that the \\sharedfolder from any Windows host in a domain other than the configured Windows Failover Cluster nodes is in the **Shared Folder Path**, and then click **Next**.



5. In the **Confirmation** window, click **Next** to confirm the details.



6. In the **Summary** window, go to the Failover Cluster Manager and verify that the quorum configuration is set to \\.sharedfolder.



7. Click **Finish**.

## XtremIO

### Best practices for zoning and subnetting

This section describes the best practices for allowing a host and the XtremIO cluster to communicate using 4, 8 or 16 paths per device.



## Recommended configuration values summary

The following table summarizes all used and recommended variables and their values for zoning and subnetting:

Validation	Impact	Severity
Multipathing: Max number of paths shall not exceed 16.	Performance	Warning
Multipathing: Recommended number of paths is 8.	Performance	Warning
Multipathing: Link speed should be consistent across all paths to the XtremIO cluster.	Performance	Warning
Multipathing: Duplex setting should be consistent across all paths to the XtremIO cluster.	Performance	Warning
Balance the hosts between the Storage Controllers to provide a distributed load across all target ports.	Performance	Recommendation
FC-configured zoning: Each host Initiator Group should be at least one path for two Storage Controllers belonging to the same X-Brick.	Stability	Mandatory
iSCSI-configured subnetting: Each host Initiator Group should be at least one path for two Storage Controllers belonging to the same X-Brick.	Stability	Mandatory
iSCSI subnetting: Configuration should not allow traffic between different iSCSI IP subnets.	Performance	Normal
iSCSI MTU value: If jumbo frames are required for iSCSI traffic, all ports (server, switches, and storage) must be configured with the correct MTU value.	Performance	Warning
iSCSI Flow control features: Flow control must be disabled on all ports (server, switches, and storage).	Performance	Warning
iSCSI TCP Offloading: Enable the TCP Offloading Engine (TOE) on the host iSCSI interfaces.	Performance	Warning
iSCSI NIC configuration: Use a dedicated NIC or iSCSI HBA for XtremIO iSCSI and do not partition the iSCSI interface (that is disable NIC Partitioning - NPAR).  Dell EMC recommends using interfaces individually rather than using NIC Teaming (Link Aggregation).	Performance	Recommendation

## General guidelines

- The optimal number of paths depends on the operating system and server information. To avoid multipathing performance degradation, do not use more than 16 paths per device. Dell EMC recommends using eight paths.  
 ⓘ **Note:** This recommendation is not applicable to Linux hosts connected to XtremIO. On such hosts, more than 16 paths per device can be used (if required).
- Balance the hosts between the Storage Controllers to provide a distributed load across all target ports.
- Host I/O latency can be severely affected by SAN congestion. Minimize the use of ISLs by placing the host and storage ports on the same physical switch. When this is not possible, ensure that there is sufficient ISL bandwidth and that both the Host and XtremIO interfaces are separated by no more than two ISL hops. For more information about proper SAN design, see the [Networked Storage Concepts and Protocols Techbook](#).
- Keep a consistent link speed and duplex across all paths between the host and the XtremIO cluster.
- To ensure continuous access to XtremIO storage during cluster software upgrade, verify that a minimum I/O timeout of 30 seconds is set on the HBAs of all hosts that are connected to the affected XtremIO cluster. Similarly, verify that a minimum timeout of 30 seconds is set for all applications that are using storage from the XtremIO cluster.  
 ⓘ **Note:** See the Dell EMC KB article 167514 for references to *Dell EMC Host Connectivity Guides*. These guides provide the procedures that are required for adjusting the HBA minimum I/O timeout.

## Minimal zoning/subnetting configuration

To prevent a host path from going down when two Storage Controllers (from separate X-Brick blocks) are down, while the XtremIO cluster remains active (because the failed Storage Controllers are not in the same X-Brick), follow these guidelines:

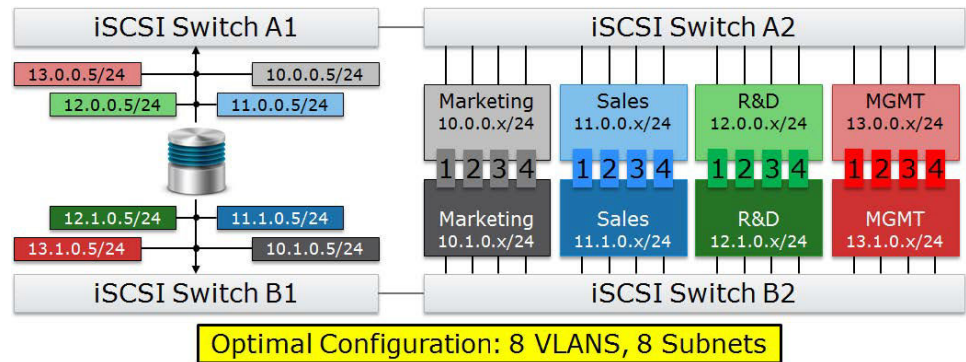
- When configuring zoning/subnetting from the host to the XtremIO cluster, the minimal zoning/subnetting configuration for each host Initiator Group should be at least one path for two Storage Controllers belonging to the same X-Brick.
- A host port must be zoned to at least two Storage Controllers ports from the same X-Brick. This zone can be expanded to additional Storage Controllers from other X-Brick blocks. Moreover, other host ports can be zoned to Storage Controllers from other X-Brick blocks.  
 ⓘ **Note:** The diagrams throughout this chapter illustrate possible implementations of these guidelines. Other possible implementations exist, but are not illustrated.

## iSCSI SAN guidelines

When setting up an iSCSI SAN infrastructure, follow these guidelines:

- If jumbo frames are required for iSCSI traffic, ensure that all ports (server, switches, and storage) are configured with the correct MTU value.  
 ⓘ **Note:** When using Jumbo Frames with VMware ESX, the correct MTU size must be set on the virtual switch as well.
- Disable flow control features on the server, switches, and array ports.
- Make sure that the different iSCSI IP subnets cannot transmit traffic between them.
- Use Virtual LAN (VLAN) technology to partition a single iSCSI network link into multiple distinct domains. If possible, use a dedicated VLAN for XtremIO iSCSI traffic and a dedicated IP subnet for each iSCSI fabric. Do not configure iSCSI routes between the different subnets.

The following figure shows optimal VLAN and IP Subnetting configuration:



- Enable the TCP Offloading Engine (TOE) on the host iSCSI interfaces, to offload the TCP packet encapsulation from the CPU of the Host to the NIC or iSCSI HBA, and free up CPU cycles.
- Dell EMC recommends using a dedicated NIC or iSCSI HBA for XtremIO iSCSI and not to partition the iSCSI interface (in other words, disable NIC Partitioning - NPAR).
- When using XtremIO iSCSI, Dell EMC recommends using interfaces individually rather than using NIC Teaming (Link Aggregation), to combine multiple interfaces into a single virtual interface.

**Note:** See the user manual of the FC/iSCSI switch for instructions about real implementations.

## Fibre Channel SAN guidelines

When setting up a Fibre Channel (FC) SAN infrastructure, follow these guidelines:


- Use a single-initiator per single-target (1:1) zoning scheme. If the FC switch zone count limitation has been reached, it is also possible to use single-initiator per multiple-target (1: many) zoning scheme.

**Note:** See the user manual of the FC/iSCSI switch for instructions about real implementations.

## 10 TB starter X-Brick (5 TB) and single X-Brick cluster

In a 10 TB starter X-Brick (5 TB) or a single X1 X-Brick configuration, a host may have up to four paths per device. On an X2 X-Brick configuration with all iSCSI or FC, a host may have up to eight paths per device.

The following figure shows the logical connection topology for four paths. This topology applies to both dual and quad HBA/NIC host architecture:




## Dual X-Brick clusters


In a dual X-Brick configuration, a host may have up to eight paths per device.

**Note:** When partially zoning/subnetting multiple servers to the XtremIO array, ensure that the I/O load of the server is distributed equally across all X-Brick. For minimal zoning/subnetting configuration guidelines, see [Minimal zoning/subnetting configuration](#).

The following figure shows the logical connection topology for four paths. This topology applies to a dual HBA/NIC host architecture:



The following figure shows the logical connection topology for eight paths. This topology applies to both dual and quad HBA/NIC host architecture:

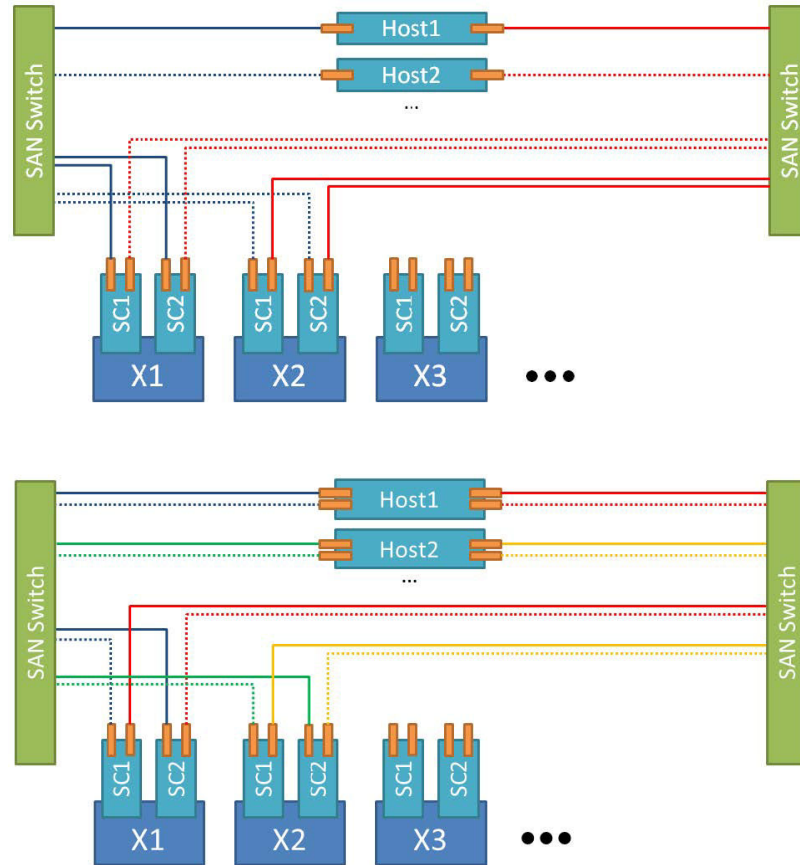


## Clusters with multiple X-Brick blocks(three or more)

In a multiple X-Brick configuration (three or more), a host may have up to 16 paths per device.

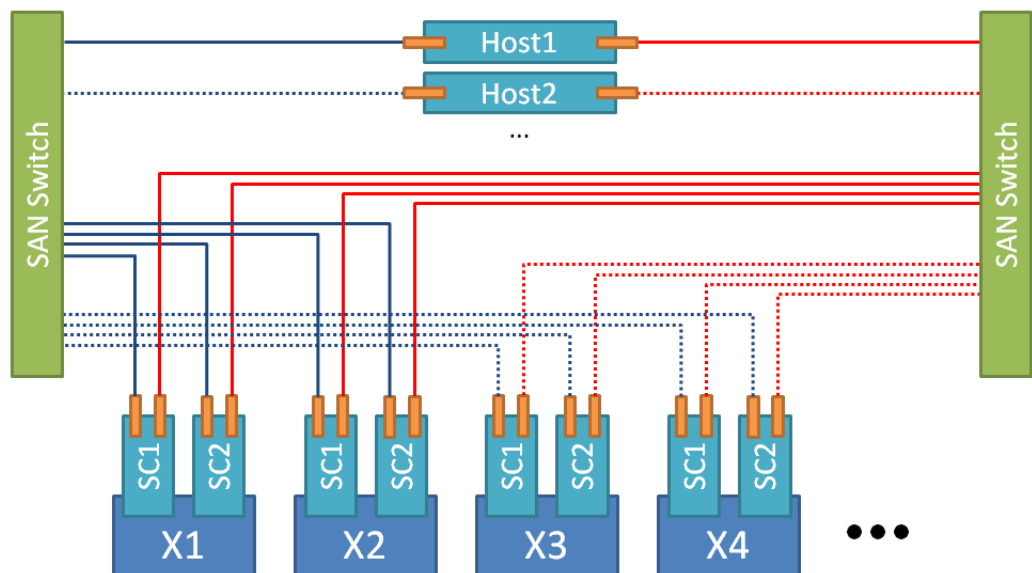
**Note:** When partially zoning/subnetting multiple servers to the XtremIO array, ensure that the I/O load of the server is distributed equally across all X-Brick blocks. For minimal zoning/subnetting configuration guidelines, see [Minimal zoning/subnetting configuration](#).

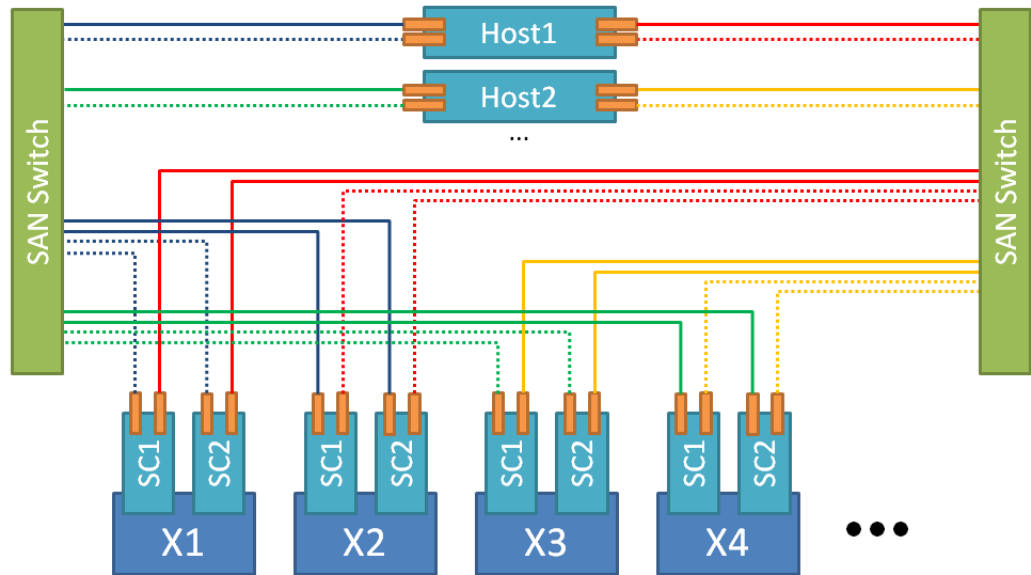
The following figures show the logical connection topology for four paths. This topology applies to both dual and quad HBA/NIC host architecture:



The following figures show the logical connection topology for eight paths. This topology applies to both dual and quad HBA/NIC host architecture.


- Note:** For clusters with an odd number of X-Brick blocks, change these examples to accommodate the cluster configuration and try to balance the host load among the X-Brick blocks of the cluster. The following figures show an eight paths connection topology with four X-Brick blocks (or more):

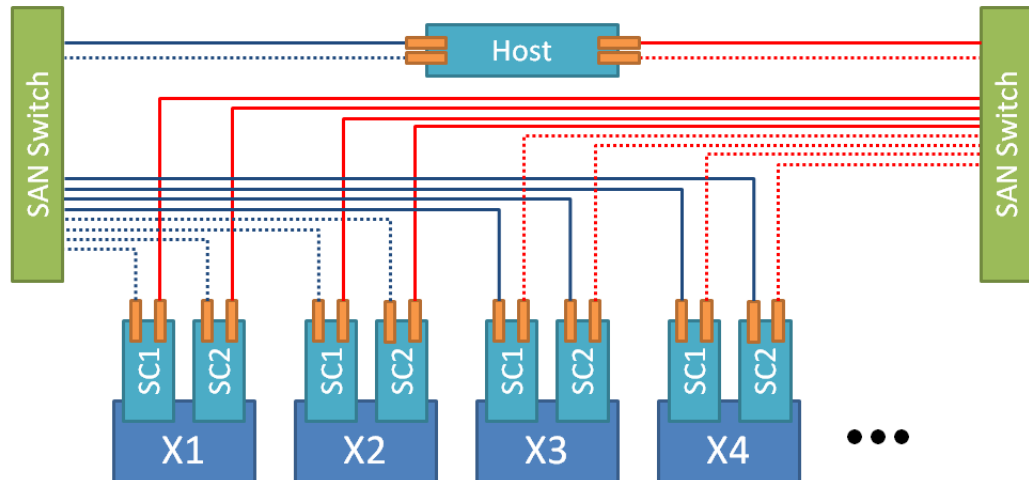




The following figures show the logical connection topology for 16 paths. This topology applies to both dual and quad HBA/NIC host architecture.

**Note:** For clusters with an odd number of X-Brick blocks, change these examples to accommodate the cluster configuration and try to balance the host load among the X-Brick blocks of the cluster. The following figure shows a 16-path connection topology with four X-Brick blocks (or more):





## Host Configuration for Microsoft Windows

### iSCSI configuration

**Note:** This section applies only for iSCSI. If you are using only Fibre Channel with Windows and XtremIO, skip to [Fibre Channel HBA configuration](#).

**Note:** Review [iSCSI SAN Guidelines](#) before you proceed.

This section describes the issues that you must address when using iSCSI with XtremIO, for optimal performance (with or without an iSCSI HBA).

#### Pre-requisites

Follow the VMware recommendations for installation and setup of the appropriate NIC/iSCSI HBA for your system. It is recommended to install the latest driver version (patch), as described in the VMware support site for each specific NIC/iSCSI HBA.

Refer to the [E-Lab Interoperability Navigator](#) for supported NIC/iSCSI HBA models and drivers.

#### Configuring the XtremIO disk device with iSCSI

To configure the XtremIO disk device make the following Windows registry settings:

```
Key:
HKLM\CurrentControlSet\Control\Class{4D36E97B-E325-11CE-BFC1-0
8002BE10318}\<Instance Number>\Parameters
FirstBurstLength (default 64K)
MaxBurstLength (default 256K)
```

Dell EMC recommends to set the FirstBurstLength and MaxBurstLength parameters to 512KB.

**Note:** For more details about configuring iSCSI with Windows, see [Microsoft](#) website.

### Fibre Channel HBA configuration

This section details Fibre Channel (FC) HBA related issues that must be addressed when using FC with XtremIO.

## Pre-requisite

To install one or more Dell EMC-approved Host Bus Adapters (HBAs) into a Windows host, follow one of these documents according to the FC HBA type:

**For Qlogic HBAs** - *Dell EMC Host Connectivity with QLogic Fibre Channel and iSCSI Host Bus Adapters (HBAs) and Converged Network Adapters (CNAs) in the Windows Environment document*

**For Emulex HBAs** - *Dell EMC Host Connectivity with Emulex Fibre Channel and iSCSI Host Bus Adapters (HBAs) and Converged Network Adapters (CNAs) in the Windows Environment document*

These documents provide guidance on configuring the host for connection to the Dell EMC Storage Arrays over a FC including any needed HBA BIOS settings. Both documents are available in the Dell EMC OEM section of the [Qlogic](#) and [Emulex](#) websites. They can also be found on <http://support.EMC.com>.

See [E-Lab Interoperability Navigator](#) for supported FC HBA models and drivers.

## Queue Depth

**Note:** The FC HBA recommendations in this section are applicable to the following FC HBAs:

- Qlogic - adapters with names that start with **ql**
- Emulex - adapters with names that start with **lpfc**

See [E-Lab Interoperability Navigator](#) for all supported FC HBA models and drivers.

**Note:** Changing queue depth settings is designed for advanced users. Increasing the queue depth may cause hosts to over-stress other arrays connected to the Windows host, resulting in performance degradation while communicating with them. To avoid this, in mixed environments with multiple array types connected to the Windows host, compare XtremIO recommendations for queue depth with those of other platform before applying them.

For optimal operation with XtremIO storage, consider adjusting the queue depth of the FC HBA.

Queue depth is the amount of SCSI commands (including I/O requests) that can be handled by storage device at a given time. Queue depth can be set on either of the following:

- Initiator level - HBA queue depth
- LUN level - LUN queue depth

The HBA queue depth (also referred to as execution throttle) setting controls the amount of outstanding I/O requests per HBA port. The HBA queue depth should be set to the maximum value. This can be done on the HBA firmware level using the HBA BIOS or CLI utility provided by the HBA vendor as follows:

- Qlogic - Execution Throttle - Change the default value (32) to 65535.
- Emulex - lpfc\_hba\_queue\_depth - No need to change the default (and maximum) value (8192).

**Note:** HBA queue depth (execution throttle) does not apply to QLE2600 and QLE8300 Series Qlogic adapters, and is read only for 10GbE adapters. For more information, see [Qlogic](#) website.

The LUN queue depth setting controls the amount of outstanding I/O requests per a single path. On Windows, the LUN queue depth can be adjusted via the Windows Registry.

**Note:** When the HBA queue depth is set to a value lower than the LUN queue depth, it may limit the LUN queue depth to a lower value than set.

The following table summarizes the default and recommended queue depth settings for Windows.



	LUN Queue Depth	HBA Queue Depth
Default value	Qlogic: 32 Emulex: 30	Qlogic: 32 Emulex: 8192
Recommended value	Qlogic: 256 Emulex: 128	Qlogic: 65535 (maximum) Emulex: 8192 (maximum)

The following procedures detail setting the LUN queue depth for Qlogic and Emulex HBAs as follows:


- Qlogic - Set the Qlogic HBA adapter LUN queue depth in Windows StorPort driver to 256 (maximum value).
- Emulex - Set the Emulex HBA adapter LUN queue depth in Windows to 128 (maximum value).

Follow the appropriate procedure according to the HBA type.

### Setting the LUN Queue Depth for Qlogic HBA


#### Procedure

1. On the desktop, click **Start**, select **Run** and open the **REGEDIT** (Registry Editor).

 **Note:** Some driver versions do not create the registry by default. In such cases the user needs to manually create the registry.

2. Select **HKEY\_LOCAL\_MACHINE** and follow the tree structure down to the Qlogic driver as follows:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\services\ql2300\
Parameters\Device
```


 **Note:** In some cases the Windows host detects the Qlogic HBA as ql2300i (instead of ql2300). In such cases, the following registry tree structure should be used instead:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\services\ql2300i
\Parameters\Device
```

3. Double-click **DriverParameter:REG\_SZ:qd=32**.
4. Change the value of **qd** to **256**.
5. If the string "qd=" does not exist, append the following text to the end of the string using a semicolon (";"):

```
qd=256
```

6. Click **OK**.
7. Exit the Registry Editor and reboot the Windows host.

 **Note:** Setting the queue-depth per this procedure is not disruptive.

### Setting the LUN queue depth for Emulex FC HBA

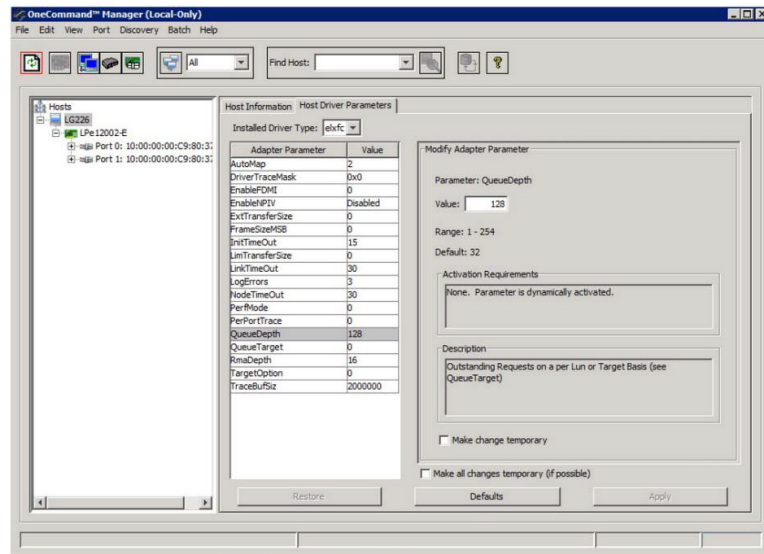
#### Before you begin

#### About this task

Setting the LUN queue depth on the Emulex FC HBA includes using Emulex UI (OneCommand) to set the LUN queue depth of the Emulex driver HBA. OneCommand detects the active Emulex driver and enables changing the relevant driver's parameters, specifically LUN queue depth.

## Procedure

1. Install OneCommand.
2. Launch the OneCommand UI.



3. Select the relevant hostname from the Hosts list.
4. Select the Host Driver Parameters tab.
5. On the Adapter Parameter list, locate the QueueDepth parameter and set its value to 128.
6. Click Apply.

## Windows Offloaded Data Transfer (ODX) Functionality

Windows Offloaded Data Transfer (ODX) functionality in Windows maximizes the enterprise's investment by enabling a storage array to directly transfer data within or between compatible storage devices, bypassing the Windows host. By offloading the file transfer to the storage array, ODX minimizes latencies, maximizes array throughput and reduces resource usage such as CPU and network consumption on the host computer.

XtremIO version 4.0 (and later) supports the Windows ODX functionality.

## Checking ODX Status

### About this task

Although ODX in Windows is enabled by default, it can be disabled. Therefore, before using Windows with XtremIO, confirm that ODX is enabled.

To verify that ODX is enabled:

### Procedure

1. Run the following command:

```
Get-ItemProperty
hkln:\system\currentcontrolset\control\filesystem -Name
"FilterSupportedFeaturesMode"
```

2. Verify that the FilterSupportedFeaturesMode value in the Windows registry is 0.

## Enabling ODX

### Before you begin

If ODX is disabled on Windows, perform the following steps to enable it.

### Procedure

1. Open a Windows PowerShell session as an administrator.
2. Run the following command:

```
Set-ItemProperty
hkml:\system\currentcontrolset\control\filesystem -Name
"FilterSupportedFeaturesMode" -Value 0
```


## Disabling ODX

Dell EMC recommends to enable ODX functionality, when using Windows with XtremIO version 4.0 (or later), for optimal performance. However, in some cases (mainly for testing purposes), it is necessary to disable ODX.

Before disabling ODX, on each Windows host that uses ODX, list the file system filter drivers attached to the volume on which you want to disable ODX. Make sure that this list is empty.

To list the file system drivers attached to a volume for ODX:

1. Open a Windows PowerShell session as an administrator.
2. Run the following command: `Fltmc instances -v <volume>`

 **Note:** <volume> refers to the drive letter of the volume.

The example below shows the expected Fltmc command output prior to disabling ODX:

```
PS C:\Users\Administrator> fltmc instances -v d:\
No instances for this volume
```

To disable the ODX perform the following steps:

1. Open a Windows PowerShell session as an administrator.
2. Check whether ODX is currently enabled, see [Checking ODX status](#).
3. If ODX is enabled on Windows, run the following command to disable ODX:

```
Set-ItemProperty
hkml:\system\currentcontrolset\control\filesystem -Name
"FilterSupportedFeaturesMode" -Value 1
```

## Multipathing software configuration

XtremIO supports native multipathing using Microsoft Multipath I/O (MPIO) with Windows Server 2008 and later, or multipathing using PowerPath.

### Configuring Multipathing using MPIO

#### Before you begin

### About this task

For optimal operation with XtremIO storage, configure the Least Queue Depth policy for MPIO for devices presented from XtremIO. Using this policy, I/O is sent down the path with the fewest currently outstanding I/O requests.

Before configuring the native multipathing, you should enable MPIO on the server by adding the MPIO feature to Windows.

To enable MPIO on the server, perform the following steps:

#### Procedure

1. On the desktop, click **Start**, right-click **Computer**, and select the **Manage** option to display the Server Manager window.
2. On the Server Manager window, click **Features** to switch to the Features section, and then click **Add Feature** to launch the Add Feature Wizard.
3. From the list of possible features select **Multipath I/O** and click **Next**.
4. Click **Install** to install the MPIO on Windows
5. If requested, reboot the host upon conclusion of the MPIO feature installation.

**To configure native multipathing on XtremIO volumes presented to the Windows host:**

#### Procedure

1. Run the following MPCLAIM command to install MPIO optional component and add multipath support for XtremIO device:

```
mpclaim -r -i -d "XtremIO XtremApp "
```

**Note:** There should be eight spaces between 'XtremApp' and the closing quotation mark.

**Note:** The host reboots immediately after running this command.

After reboot, use the `mpclaim -h` command to display the hardware IDs that are already managed by MPIO (the list should include XtremIO).

2. Run the following mpclaim command, to set the load-balancing algorithm to Least Queue Depth for all XtremIO volumes:

```
mpclaim -l -t "XtremIO XtremApp " 4
```

**Note:** There should be eight spaces between 'XtremApp' and the closing quotation mark.

**Note:** This command does not affect non-XtremIO volumes presented to the Windows host.

**Note:** Use the `mpclaim -s -t` command to check the default load-balancing settings for XtremIO devices. Use the `mpclaim -s -d` command to list all disks currently claimed by MPIO and their load-balancing settings.

### PowerPath Multipathing with XtremIO

**Note:** For the most updated information on PowerPath support with XtremIO storage, see the *XtremIO Simple Support Matrix* on [Dell EMC E-Lab Navigator](#).

**Note:** When using PowerPath for Windows, see *ETA 491197: PowerPath for Windows*: Microsoft Windows 2012 R2 server generates a blue screen event which may result in a system reboot,

potential file loss and a bugcheck message that identifies Dell EMC PowerPath driver EMCMPX.SYS (<https://support.emc.com/kb/491197>).

XtremIO supports multipathing using Dell EMC PowerPath on Windows. PowerPath versions 5.7 SP2 and above provide Loadable Array Module (LAM) for XtremIO Array devices. With this support, XtremIO devices running versions 2.2 and above are managed under the XtremIO class.

PowerPath provides enhanced path management capabilities for up to 32 paths per logical device, as well as intelligent dynamic I/O load-balancing functionalities specifically designed to work within the Microsoft Multipathing I/O (MPIO) framework. Having multiple paths enables the host to access a storage device even if a specific path is unavailable. Multiple paths share the I/O traffic to a storage device, using intelligent load-balancing policies which enhance I/O performance and increase application availability. Dell EMC PowerPath is the recommended multipathing choice.

PowerPath features include:

- Multiple paths - provides higher availability and I/O performance
  - Includes the support on Server Core and Hyper-V (available in Windows Server 2008 and later).
- Running PowerPath in Hyper-V VMs (guest operating systems), PowerPath supports:
  - iSCSI through software initiator
  - Virtual Fibre Channel for Hyper-V (available in Windows Server 2012 and above) that provides the guest operating system with unmediated access to a SAN through vHBA
- Path management insight capabilities - PowerPath characterizes I/O patterns and aides in diagnosing I/O problems due to flaky paths or unexpected latency values. Metrics are provided on:
  - Read and write - in MB/seconds per LUN
  - Latency distribution - the high and low watermarks per path
  - Retries - the number of failed I/Os on a specific path
- Autostandby - automatically detects intermittent I/O failures and places paths into autostandby (also known as flaky paths)
- PowerPath Migration Enabler - is a host-based migration tool that allows migrating data between storage systems and supports migration in an MSCS environment (for Windows 2008 and later). PowerPath Migration Enabler works in conjunction with the host operating system (also called Host Copy) and other underlying technologies such as Open Replicator (OR).
- Remote monitoring and management
  - PowerPath Management Appliance 2.2 (PPMA 2.2)
  - Systems Management Server (SMS)
  - Microsoft Operations Manager
  - SNMP management daemon

Further PowerPath related information:

- For details on the PowerPath releases supported for your Windows host, see the *XtremIO Simple Support Matrix*.
- For details on class support with XtremIO for your host, see the *Dell EMC PowerPath release notes* for the PowerPath version you are installing.
- For details on installing and configuring PowerPath with XtremIO class support on your host, see the *Dell EMC PowerPath on Windows Installation and Administration Guide* for the PowerPath version you are installing. This guide provides the required information for placing XtremIO volumes under PowerPath control.

**Note:** The PowerPath with XtremIO class support installation procedure is fully storage-aware. All required PowerPath settings with XtremIO storage are automatically done when PowerPath is installed on your host. This includes settings such as the PowerPath multipathing policy that does not require manual setting.

## Post configuration steps - Using the XtremIO storage

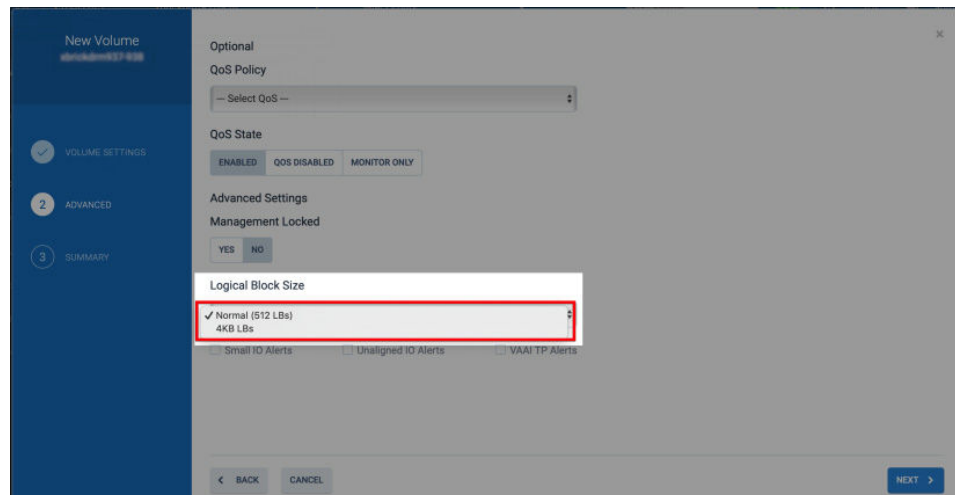
After the host configuration is completed, you can use the XtremIO storage from the host. For details on creating, presenting and managing volumes accessed from the host via either the GUI or the CLI, see the *XtremIO Storage Array User Guide* that matches the version running on your XtremIO cluster.

### Disk formatting

When creating volumes in XtremIO for a Windows host, the following considerations should be made:

- **Disk Logical Block Size**—A 512B logical block size must be used for a new XtremIO volume.

The following figure demonstrates formatting an XtremIO Volume using the WebUI.



For details on formatting a newly created Volume (using either the WebUI or the GUI interfaces), see the *XtremIO Storage Array User Guide* that matches the version running on your XtremIO cluster.

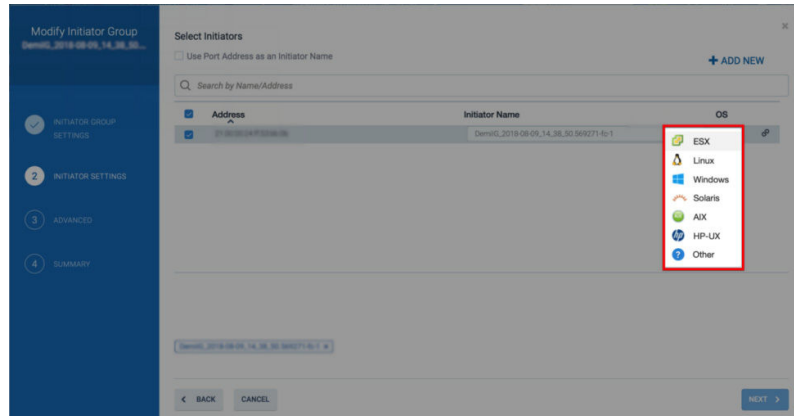
### Presenting XtremIO volumes to the Windows host

**Note:** It is recommended to avoid using host-level software for deduplication, compression and encryption of data stored in XtremIO. These operations are performed natively on the XtremIO cluster, thus releasing host computing resources. Furthermore, using these operations may severely impact data efficiency on the XtremIO cluster, as they can lead to a higher physical or logical capacity utilization.

**Note:** The following information in this section applies only to XtremIO version 4.0 and later.

When adding Initiator Groups and Initiators to allow Windows hosts to access XtremIO volumes, specify **Windows** as the operating system for the newly-created Initiators.

The following figure demonstrates setting the Operating System field for an Initiator using the WebUI.



- Note:** Setting the Initiator's Operating System is required for optimal interoperability and stability of the host with XtremIO storage. You can adjust the setting while the host is online and connected to the XtremIO cluster with no I/O impact.
- Note:** See the *XtremIO Storage Array User Guide* that matches the version running on your XtremIO cluster.

Following a cluster upgrade from XtremIO version 3.0.x to version 4.0 (or later), make sure to modify the operating system for each initiator that is connected to a Windows host.

## Creating a file system

File system configuration and management are out of the scope of this document.

- Note:** Some file systems may require you to properly align the file system on the XtremIO volume. It is recommended to use specified tools to optimally match your host with application requirements.

## Space reclamation

This section provides a comprehensive list of capacity management steps for achieving optimal capacity utilization on the XtremIO array, when connected to a Windows host.

Data space reclamation helps to achieve optimal XtremIO capacity utilization. Space reclamation is a Windows operating system function that enables to reclaim used space by sending zeros to a specific address of the volume after being notified by the file system that the address space was deleted. While some Windows operating systems can perform this action automatically, others require a user-initiated operation.

The following sections detail steps for performing space reclamation with:

- NTFS file systems
- Hyper-V

### NTFS File Systems

**Automatic space reclamation**—On Windows server 2012 and above, NTFS supports automatic space reclamation when a file is deleted from the file system.

Delete notification (also referred to as TRIM or UNMAP) is a feature that notifies the underlying storage device of clusters that have been freed, following a file delete operation. The `DisableDeleteNotify` parameter is used to disable (using the value 1) or enable (using the value 0) delete notifications for all volume. This parameter was introduced in Windows Server 2008 R2 and Windows 7.

The following example shows querying the delete notification status.

```
PS C:\> fsutil behavior query disableddeletenotify
DisableDeleteNotify = 0
```

The following example shows enabling delete notification (TRIM).

```
PS C:\> fsutil behavior set disableddeletenotify 0
```

The following example shows disabling delete notification (TRIM).

```
PS C:\> fsutil behavior set disableddeletenotify 1
```

**Note:** Disabling delete notifications disables all TRIM commands from the Windows Host.

- Hard disk drives and SANs that do not support TRIM, will not receive TRIM notifications.
- TRIM is enabled by default and can be disabled by the administrator.
- Enabling or disabling delete notifications does not require a restart.
- TRIM is effective when the next UNMAP command is issued.
- Existing inflight I/Os are not impacted by the registry change.

**Note:** In Windows Server 2012, when a large file is deleted, the file system performs space reclamation according to the storage array setting. Large file deletions can potentially affect the performance of the regular I/O. To avoid this, delete notification should be disabled. For further details, See [Microsoft](#) website.

**Note:** All files must be permanently deleted (i.e. removed from the Recycle Bin) for the automatic deletion to be performed.

**Manual space reclamation**-You can perform manual space reclamation, using the following options:

- Windows Optimizer (Optimize-Volume cmdlet)
- SDelete utility
- PowerShell script

**Windows Optimizer (Optimize-Volume cmdlet)**-Starting from Windows 2012, Windows introduced the option to reclaim space on a TRIM-enabled array.

The `delete notifications` option must be enabled for running manual space reclamation, using Windows Optimizer or the Optimize-Volume command.

**Note:** It is suggested to enable the `delete notifications` parameter for using Windows space reclamation tools. When space reclamation is complete, disable this parameter to avoid automatic space reclamation.

**Note:** In versions prior to Windows 2012R2, the entire file system is locked during every file deletion. For further details, contact Microsoft.

To run TRIM on a volume, use the following command: `Optimize-Volume`

The following example show running TRIM on volume E.

```
PS C:\> Optimize-Volume -DriveLetter E -ReTrim -Verbose
VERBOSE: Invoking retrim on New Volume (E:)...
VERBOSE: Performing pass 1:
VERBOSE: Retrim: 0% complete...
VERBOSE: Retrim: 10% complete...
VERBOSE: Retrim: 20% complete...
VERBOSE: Retrim: 30% complete...
```



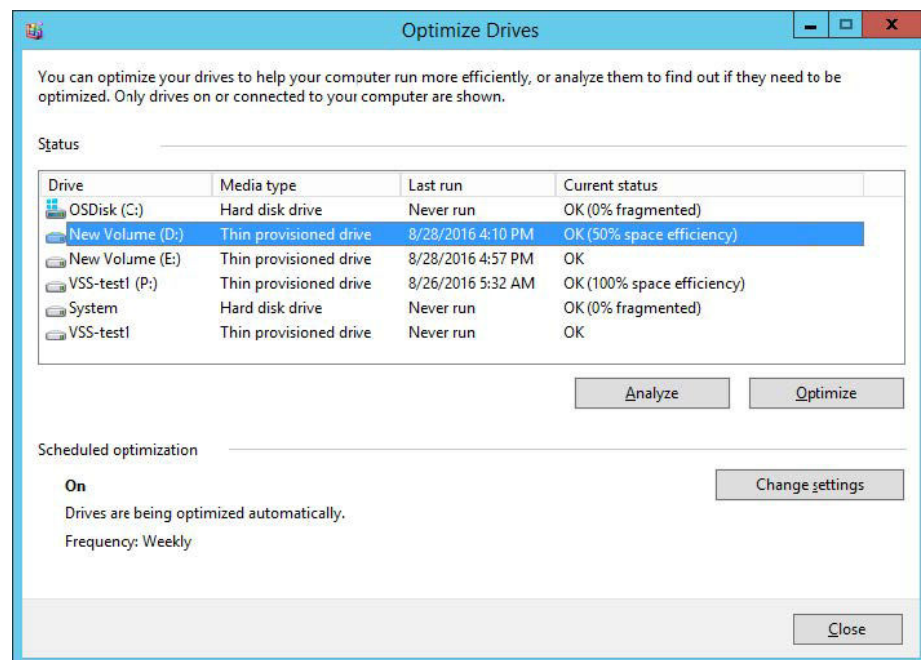
```

VERBOSE: Retrim: 40% complete...
VERBOSE: Retrim: 50% complete...
VERBOSE: Retrim: 60% complete...
VERBOSE: Retrim: 70% complete...
VERBOSE: Retrim: 80% complete...
VERBOSE: Retrim: 90% complete...
VERBOSE: Retrim: 100% complete.
VERBOSE:
Post Defragmentation Report:
VERBOSE:
Volume Information:
VERBOSE: Volume size                = 1023.99 GB
VERBOSE: Cluster size              = 4 KB
VERBOSE: Used space                 = 207.73 MB
VERBOSE: Free space                 = 1023.79 GB
VERBOSE:
Retrim:
VERBOSE: Backed allocations          = 1023
VERBOSE: Allocations trimmed        = 1022
VERBOSE: Total space trimmed         = 1020.87 GB

```

When using Windows Optimizer, see [Microsoft Support](#) to update Windows.

Another option is running the tool from Windows GUI.



When running Optimizer, using the Windows GUI, the Optimize button may be grayed-out in some cases, due to an issue currently under investigation with Microsoft. In this case, use the command line to run Optimizer.

**SDelete utility**—Is supported starting from Windows 2008. This utility was originally designed to provide an option for securely deleting data on magnetic disks by overwriting on-disk data, using various techniques to ensure that disk data is unrecoverable.

For more information and for downloading the SDelete utility, go to the [Microsoft](#) website.

SDelete allocates the largest file possible, using non-cached file I/O, to prevent the contents of the NT file system cache from being replaced with useless data, associated with SDelete's space-hogging file. Since non-cached file I/O must be sector (512-byte) aligned, there may be non-allocated space left even when SDelete cannot increase the file size. To solve this, SDelete allocates the largest cache file possible.

The SDelete option has the following drawbacks:

- Slow performance
- No control on the balloon file. This may lead to a 100% full file system during the reclamation process, generating an I/O error.

The following example shows running SDelete on s:\ mount point.

```
PS C:\Users\Administrator> sdelete.exe -z s:
SDelete - Secure Delete v1.61
Copyright (C) 1999-2012 Mark Russinovich
Sysinternals - www.sysinternals.com
SDelete is set for 1 pass.
Free space cleaned on S:\
1 drives zapped
```

**Note:** To ensure optimal manual space reclamation operation when using SDelete, use the SDelete parameter corresponding to the Zero free space option. In some SDelete versions (e.g. version 1.51), the Zero free space option corresponds to `-c` and not to `-z`.

**Note:** During space reclamation, the drive capacity is reduced to minimum because of the created balloon file.

If the server is running custom made PS scripts for the first time, See Example 2 in [Windows Powershell Reclaim Script](#).

### HYPER-V space reclamation

UNMAP requests from the Hyper-V guest operating system - During the virtual machine (VM) creation, a Hyper-V host inquires whether the storage device, holding the virtual hard disk (VHD), supports UNMAP or TRIM commands. When a large file is deleted from the file system of a VM guest operating system, the guest operating system sends a file delete request to the virtual machine's virtual hard disk (VHD) or VHD file. The VM's VHD or VHD file tunnels the SCSI UNMAP request to the class driver stack of the Windows Hyper-V host.

Microsoft hypervisor passes T10 commands. Therefore, if the guest operating system file system supports online space reclamation, no additional task is required.

To run manual space reclamation in the guest OS level, refer to the relevant OS chapter.

**Note:** If delete notification is disabled on the hypervisor level, all guest VMs cannot utilize space reclamation features until delete notification is re-enabled.

The following example describes a Windows Powershell reclaim script.

```
param(
    [Parameter(Mandatory=$true, ValueFromPipelineByPropertyName=$true)]
    [ValidateNotNullOrEmpty()]
    [Alias("name")]
    $Root,
    [Parameter(Mandatory=$false)]
    [ValidateRange(0,1)]
    $PercentFree = .05
)
process{
    #Convert the #Root value to a valid WMI filter string
    $FixedRoot = ($Root.Trim("\") -replace "\\","\\") + "\\"
    $FileName = "ThinSAN.tmp"
    $FilePath = Join-Path $Root $FileName

    if( (TestPath $FilePath) ) {
        Write-Error -Message "The file $FilePath already exists, please
```

```

delete the file and try again"
} else {
$Volume = gwmi win32_volume -filter "name='$FixedRoot'"
if($Volume) {
    #Change block size to control performance impact
    $ArraySize = 64kb
    $SpaceToLeave = $Volume.Capacity * $PercentFree
    $FileSize = $Volume.FreeSpace - $SpaceToLeave
    $ZeroArray = new-object byte[] ($ArraySize)

    $Stream = [io.File]::OpenWrite($FilePath)
    try {
        $CurFileSize = 0
        while($CurFileSize -lt $FileSize) {
            $Stream.Write($ZeroArray,0, $ZeroArray.Length)
            $CurFileSize += $ZeroArray.Length
        }
    } finally {
        #always close our file stream, even if an exception occurred
        if($Stream) {
            $Stream.Close()
        }
        #always delete the file if we created it, even if an exception occurred
        if( (Test-Path $filePath) ) {
            del $FilePath
        }
    }
} else {
    Write-Error "Unable to locate a volume mounted at $Root"
}
}
}

```

In case the error message, as displayed in **Example 1**, appears while running the script for the first time, update the execution policy as displayed in **Example 2**.

- **Example 1**

```

PS C:\Users\Administrator> Write-ZeroesToFreeSpace.ps1 -Root "m:\"
-PercentFree .1
File C:\Windows\system32\Write-ZeroesToFreeSpace.ps1 cannot be
loaded because the execution of scripts is disabled on this system.
Please see "get-help about_signing" for more details.
At line:1 char:28
+ Write-ZeroesToFreeSpace.ps1 <<<< -Root "m:\" -PercentFree .1
+ ~~~~~ CategoryInfo          : NotSpecified: (:) [],
PSSecurityException
+ ~~~~~ FullyQualifiedErrorId : RunTimeException

```

- **Example 2**

```

PS C:\Users\Administrator> Set-ExecutionPolicy unrestricted
Execution Policy Change
The execution policy helps protect you from scripts that you do not
trust. Changing the execution policy might expose
you to the security risks described in the about_Execution_Policies
help topic. Do you want to change the execution
policy?
[Y] Yes [N] No [S] Suspend [?] Help (default is "Y"): Y

```

## ReFS File Systems

ReFS (Resilient File System) is a new file system introduced by Microsoft on Windows 2012. While Windows 2016 enables the unmap feature on NTFS by default, for ReFS unmap is not set.

Example:

```
C:\Users\Administrator>fsutil behavior query disableddeletenotify
NTFS DisableDeleteNotify = 0
ReFS DisableDeleteNotify is not currently setY
```

To enable the unmap feature on ReFS, run the following command:

```
fsutil behavior set disableddeletenotify refs 0
```

**Note:** If the command does not specify refs, the unmap feature only works on NTFS (by default).

To query the Windows host about the current DisableDeleteNotify setting for ReFS, specify the file system type as well. If the file system type is not specified, both values are listed (as shown in the previous example).

## Disabling the UNMAP Feature

When using space reclamation with Windows 2012 or Windows 2012 R2, two problems may arise:

- Quick formatting of a volume can take a long time to complete since a large number of TRIM commands is sent to the array.
- Due to a limitation in Windows 2012 (solved in Windows 2012 R2), asynchronous TRIM command locks the entire disk (MS PhysicalDisk) for an I/O until all of its required space is reclaimed.

It is possible to avoid these situations by disabling the UNMAP feature, and then re-enabling it when the relevant operation (i.e. format or delete) is completed.

### Disable the UNMAP feature:

Issue the following command to disable the sending of TRIM commands to the array:

```
fsutil behavior set disableddeletenotify 1
```

**Note:** This setting applies to all volumes in the host and therefore needs to be re-issued for each disk to format.

While UNMAP is disabled, automatic TRIM commands are not sent to XtremIO. Space reclamation can only be carried out manually using the Windows SDelete utility. For details about using the SDelete utility, see: [Microsoft](#) website.

## Configuring the volume snapshot service using XtremIO VSS hardware provider

Volume snapshot service (VSS) is a set of Component Object Model (COM) application programming interfaces (APIs), that provides standardized interfaces. VSS enables third-party backup and restoration software to centrally manage the backup and restore operations on a variety of applications. VSS also implements a framework that enables volume backups to be performed while applications on the system continue to write to the volumes.

VSS includes the following components:

- **Requestor** - Application that requests the creation of a shadow copy
- **Provider** - Interface that provides the functionality to perform the shadow copy
- **Writer** - Application-specific software, that ensures application data is ready for shadow copy creation

The Requestor initiates the backup and restore processes, and the Provider controls the processes and instructs the Writer to prepare a dataset for backup. When these steps are completed, the Requestor instructs the Provider to create a shadow copy.

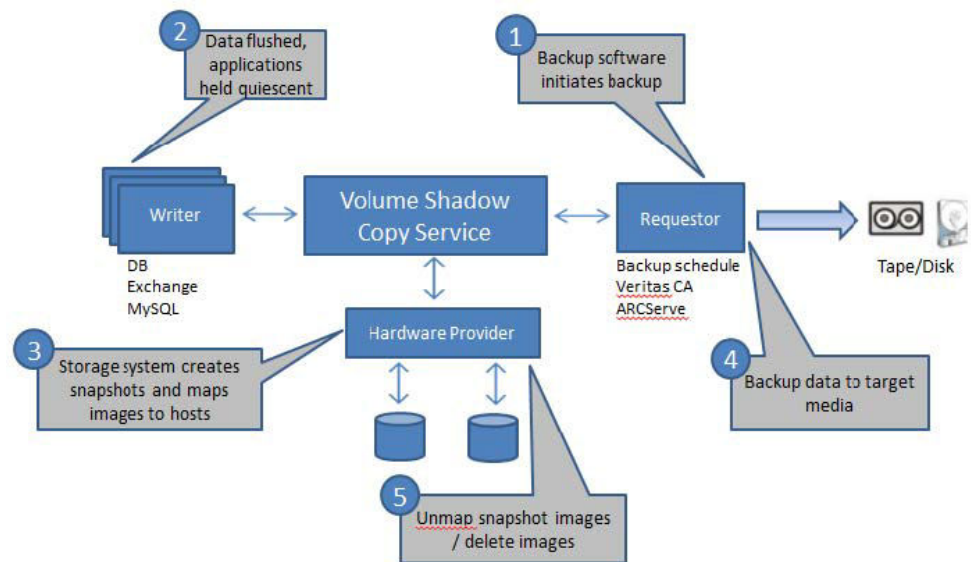
### Basic VSS Scenario

#### About this task

The following steps describe a basic VSS scenario, in which a backup application (Requestor) requests to back up a volume while the Exchange server (Writer) is up and running:

#### Procedure

1. The Requestor interacts with the VSS service to start the backup.
2. VSS service sends a Freeze request to the Writer. The Writer freezes a stable state of the volume (after closing relevant files, locking additional requests etc.).  
**Note:** Each Writer provides its own implementation of the freeze procedure.
3. The Requestor queries the Writer Metadata Document xml document to acquire information about files that require backup.
4. The Provider creates a snapshot of the file-system, maps a new volume with the snapshot content (mapping is performed on-demand when the snapshot is exposed as a volume), and returns a response to the VSS Service.
5. The VSS service requests the Writer to unfreeze the volume, by sending a Thaw event.
6. The Requestor uses the snapshot to back up the content while having a read-only/complete state of that volume



#### Requirements for configuring VSS with XtremIO VSS hardware provider

- The Windows host should be installed with Windows Server 2008 or above.
- The Windows host should have access to the XMS RESTful API.
- XtremIO volume(s) should be mapped to the Windows host.
- XMS access level should be Configuration (or above).
- XtremIO VSS provider should be installed.
- VSS service should be up and running.

## Installing the VSS Provider on the Windows Host

### Before you begin

#### About this task

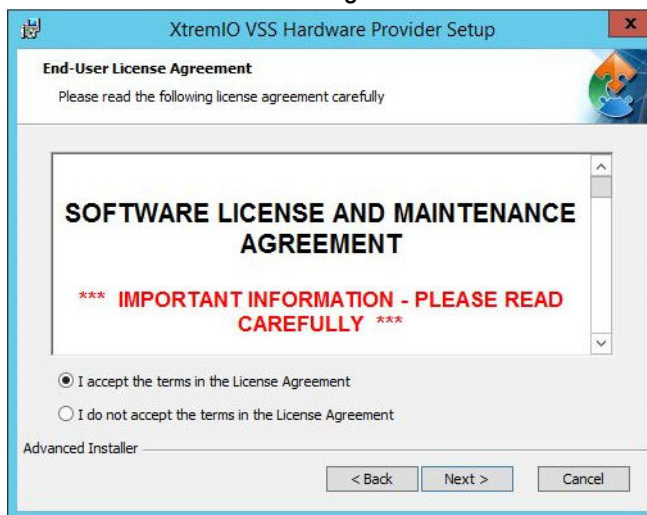
Perform the following steps to locate the VSS Provider package on the support page and to install the package on the Windows host:

#### Procedure


1. Access the Dell EMC Support page for XtremIO in [support.emc.com](http://support.emc.com).
2. Download the XtremIO VSS Provider package to the Windows host. For details on which XtremIO VSS Provider package to download, refer to the Release Notes of the version you are installing.
3. Run the XtremIOVSSProvider.msi from an elevated command prompt or PowerShell prompt (i.e. the shell should be run as an administrator).



4. Accept the software license and maintenance agreement and click **Next**.




5. Select the installation folder.



**Note:** The XtremIO Hardware Provider should be installed on the local disk.

6. Click **Install**.



7. When installation is complete, verify that the Launch XtremIO VSS Hardware Provider option is selected, and click **Finish**.





8. In the opened XtremIO VSS Provider Control Panel window, provide the required details.

The XtremIO VSS Provider Control Panel window is displayed. It contains the following fields and controls:

- XMS Server IP Address:** A text input field.
- Username:** A text input field.
- Password:** A password input field with a masked character (•).
- Buttons:** OK and Cancel buttons at the bottom right.

Text inside the window: "The XtremIO VSS Provider Control Panel would configure the VSS Hardware provider. Please insert the XMS address and credentials to allow the Hardware Provider to communicate with the cluster."

- Note:** Verify that the typed IP address and the credentials are correct before clicking OK.

9. Click OK; a confirmation message is displayed.



10. In the PowerShell, run the `vssadmin list providers` command and verify that the XtremIO VSS provider appears in the providers list and is properly installed.

- Note:** If XtremIO VSS Hardware Provider fails to successfully register as a COM+ application on the machine it is installed on, it may not function properly.

### XtremIO VSS Provider Components

- Windows Installer (MSI)
  - Installs and uninstalls the provider service and the control panel Applet.
  - Installs VC2012 Runtime environment if required.
  - Supports silent mode.
- Control Panel GUI Applet
  - The control panel displays an XtremIO provider icon.
  - Running the applet enables configuring the provider's server address and username/password.
- VSS Hardware Provider service
  - VSS COM service is registered as a hardware provider and communicates with the VSS service.
  - The XMS communicates with the VSS on the Windows host via RESTful API.

### XtremIO VSS Features

#### VSS Functionality

- AreLunsSupported** - Used to identify LUNs managed by XtremIO Storage system (using SCSI inquiry Vendor-id and Product-Id) and request VSS to flag these volumes, to be managed by XtremIO provider



- [GetTargetLuns](#) - Used to provide VSS with the LUN(s) information for the snapshot volumes.
- [LocateLuns](#) - Used to map LUNs requested by the VSS service.
- [OnLunEmpty](#) - A VSS request used to delete and unmap a LUN.
- [BeginPrepare Snapshot](#) - A VSS request used to prepare for snapshotting.
- [FillInLunInfo](#) - A VSS request used to confirm and fix LUN(s) structure.
- [CommitSnapshots](#) - Used to commit the snapshots, using a naming convention of (Original volume name\_-Year-Month-Day-Hour-Minute-Sec).
- [AbortSnapshots](#) - Used to abort snapshot creation.
- [RegisterProvider](#) - Used to register the hardware provider.
- [UnregisterProvider](#) - Used to unregister the hardware provider.
- [OnLunStateChange](#) - A VSS request used to unmap a LUN.

#### Extra VSS Provider Features

- Multi-clusters support
- Reconfiguration (changing configuration via the control panel)
- Fibre Channel and iSCSI interfaces support
- On Fibre Channel, map to all ports with different Initiator Groups

#### Tools

- Vshadow
  - A command-line tool used for creating and managing volume shadow copies
  - Acts as a requestor
  - Provided with Microsoft SDK 7+
    - ⓘ **Note:** Verify that you are running the x64 version, by checking the path C:\Program Files (x86)\Microsoft SDKs\Windows\v7.1A\Bin\x64.
  - Vshadow Syntax examples:  
Create a persistent shadow copy set from the d: volume (d: must be an XtremIO volume):  
`Vshadow.exe -p d:`  
  
Create a persistent shadow copy set from multiple volumes x: y: z: (must be XtremIO volumes).  
`Vshadow.exe -p x: y: z:`  
  
Delete a specific shadow copy.  
`Vshadow.exe -ds=ShadowCopyId`  
  
Delete all shadow copies.  
`Vshadow.exe -da`  
  
Expose a shadow copy to the X: drive (read-only).  
`Vshadow.exe -er=ShadowCopyId, x:`

For more syntax examples, see [VShadow Tool and Sample](#).

### Configuring boot-from-SAN with an XtremIO volume

For guidelines on configuring the host to boot from an XtremIO volume, see the *Dell EMC XtremIO Storage Array Host Configuration Guide* on [Dell EMC Online Support](#).

## **Executing the Windows host validation script on Windows hosts**

After all Windows hosts connected to XtremIO are configured according to the recommendations provided in this document, execute the Windows Host Validation Script (HVS) on these hosts to validate that the host configuration is aligned with these recommendations. For details on using the Windows HVS, see the Dell EMC [KB# 504234](#).

# CHAPTER 5

## Operating System Features

This chapter contains the following topics:

• <a href="#">General Windows information</a> .....	112
• <a href="#">Booting Windows from external storage</a> .....	112
• <a href="#">Microsoft Windows failover clustering</a> .....	114
• <a href="#">Operating system characteristics</a> .....	115
• <a href="#">Hyper-V</a> .....	116

## General Windows information

This section provides information that is common to all supported versions of Windows. Read the entire section before proceeding to the rest of the chapter.

## Booting Windows from external storage

Currently, it is not possible to boot a Windows system using an iSCSI disk volume that is provided by the Microsoft software iSCSI Initiator kernel mode driver. However, it is possible to boot a Windows system using an iSCSI disk volume that provided by an iSCSI host bus adapter (HBA). The only currently supported method for booting a Windows system is using an iSCSI disk volume is through supported HBA.

Windows hosts have been qualified for booting from Dell EMC array devices interfaced through Fibre Channel as described under Boot Device Support in the [Dell EMC Simple Support Matrix](#). See the appropriate Windows HBA guide, available on [Dell EMC Online Support](#), for information on configuring your HBA and installing the Windows operating system on an external storage array.

For details about Broadcom (Emulex) HBAs and converged network adapters (CNAs), see the [Broadcom](#) website.

For Glogic driver downloads, see the [Cavium](#) website.

## Boot-from-SAN

Although Windows servers typically boot the operating system from a local, internal disk, many customers want to use the features of PowerMax, VMAX, Unity, Unified VNX, and XtremIO storage to store and protect their boot disks and data. Boot-from-SAN enables PowerMax, VMAX3, Unity, Unified VNX, and XtremIO to be used as the boot disk for your server instead of a directly attached (or internal) hard disk. You can configure a server to use a LUN that is presented from the array as its boot disk using a properly configured Fibre Channel (FC) host bus adapter (HBA), Fibre Channel over Ethernet (FCoE), converged network adapter (CNA), or blade server mezzanine adapter that is connected and zoned to the same switch or fabric as the storage array.

### Benefits of boot-from-SAN

Boot-from-SAN can simplify management in the data center. Separating the boot image from each server enables administrators to take advantage of their investments in Dell EMC storage arrays to achieve high availability, better data integrity, and more efficient storage management. Other benefits can include:


- Improved disaster tolerance
- Reduced total cost through diskless servers
- High-availability storage
- Rapid server repurposing
- Consolidation of image management

### Boot-from-SAN configuration restrictions

See the [Dell EMC Simple Support Matrix](#) for any specific boot-from-SAN restrictions. This guide no longer contains restriction information. The information in the [Dell EMC Simple Support Matrix](#) supersedes any restriction information that is in previous HBA installation guides.


## Risks of booting from the storage array

When you use the storage array as a boot disk, Dell EMC recommends that you shut down the host server during any maintenance procedures that could make the boot disk unavailable to the host.

 **CAUTION** Microsoft Windows operating systems use virtual memory paging files that reside on the boot disk. If the paging file becomes unavailable to the memory management system when it is needed, the operating system will fail and produce a blue-screen error message.

Any of the following events could cause a system that is booting from the storage array to fail:

- Lost connection to array (pulled or damaged cable connection)
- Array service and upgrade procedures, such as online microcode upgrade or configuration change, or both
- Array failure, including failed lasers on Fibre Channel (FC) ports
- Array power failure
- Storage array network (SAN) failure such as FC switches, switch components, or switch power
- SAN service and upgrade procedure, such as a firmware upgrade or hardware replacement

 **Note:** Dell EMC recommends moving the Windows virtual memory paging file to a local disk when booting from the storage array. See the Windows manual for instructions on how to move the paging file.

## How to determine I/O latency and load on the boot LUN

The current restrictions for boot-from-array configurations that are listed in the [Dell EMC Simple Support Matrix](#) represent the maximum configuration that is allowed using typical configurations. In some cases, your applications, host, array, or SAN might already be used to a point at which these maximum values might not be achieved. Under these conditions, for improved performance and functionality, you might want to reduce the configuration from the maximums that are listed in the [Dell EMC Simple Support Matrix](#).

You can use the following guidelines to determine if your environment might not support the maximum allowed boot-from-array configurations:

- Using the Windows Performance Monitor, capture and analyze the Physical Disk and Paging File counters for your boot LUN. If response time (sec/operation) or disk queue depth seem to be increasing over time, review any additional loading that might be affecting the boot LUN performance (HBA/SAN saturation, failovers, ISL usage, and so forth).
- Use available array performance management tools to determine if the array configuration, LUN configuration, and access are configured optimally for each host.

Possible ways to reduce the load on the boot LUN include:

- Move application data away from the boot LUN.
- Reduce the number of LUNs that are bound to the same physical disks.
- Select a RAID type that provides improved performance.
- Contact your Dell EMC support representative for additional information.

## Configuring Unity and VNX series systems for boot-from-SAN

Unity series and Unified VNX storage systems are configured with all the settings that a Windows server requires for a successful boot-from-SAN as default. Unity and Unified VNX series storage systems have two storage processors (SPs), which allow for highly available data access even if a

single hardware fault has occurred. For a host to be properly configured for high availability with boot-from-SAN, the HBA BIOS must have connections to both SPs on the Unity and Unified VNX systems.

At the start of the Windows boot procedure, no failover software is running. The HBA BIOS, with a primary path and one or more secondary paths that are properly configured (with access to both SPs), will provide high availability while booting from SAN with a single hardware fault.

**Note:** It is recommended to use failover mode 4 (ALUA active/active), when it is supported. ALUA allows I/O access to the boot LUN from either SP, regardless of which SP currently owns the boot LUN.

Failover mode 1 is an active/passive failover mode. I/O can only be completed successfully if it is directed to the SP that currently owns the boot LUN. If the HBA BIOS attempts to boot from a passive path, the BIOS will have to time out before attempting a secondary path to the active (owning) SP, which can cause delays at boot time. Using ALUA failover mode whenever possible will avoid these delays.

To configure a host to boot from SAN, the server must have a boot LUN presented to it from the array, which requires registration of the WWN of the HBAs or converged network adapters (CNAs), or the iSCSI Qualified Name (IQN) of an iSCSI host.

In configurations with Windows servers attached to VNX arrays, and the Dell EMC Unisphere or Navisphere Agent installed on the server, the agent automatically registers the server's HBA WWNs on the array.

Unity series or Unified arrays do not use an agent to perform the registration. You must manually register of the HBA WWNs to present a LUN to the server for boot.

**Note:** In boot-from-SAN configurations, where the operating system is installed on the Unity series or Unified VNX series, no agent is available to perform the registration. You must manually register the HBA WWNs to present a LUN to the server for boot.

## Microsoft Windows failover clustering

Failover clustering is a Windows Server feature that enables you to group multiple servers into a fault-tolerant cluster. It provides new and improved features for software-defined data center customers and for many other workloads that run clusters on physical hardware or virtual machines.

A failover cluster is a group of independent systems that work together to increase the availability and scalability of clustered roles (formerly called clustered applications and services). The clustered servers (called nodes) are connected by physical cables and by software. If one or more of the cluster nodes fail, other nodes begin to provide service (a process known as failover). In addition, the clustered roles are proactively monitored to verify that they are working. If they are not working, they are restarted or moved to another node.

Failover clusters also provide Cluster Shared Volume (CSV) functionality that provides a consistent, distributed namespace that clustered roles can use to access shared storage from all nodes. With Windows Server failover clustering, users experience a minimum of disruptions in service.

Failover clustering has many practical applications, including:

- Highly available or continuously available file share storage for applications such as Microsoft SQL Server and Hyper-V virtual machines.
- Highly available clustered roles that run on physical servers or on virtual machines that are installed on servers running Hyper-V.

For more information, see [Failover Clustering in Windows Server](#) in the Microsoft Windows IT Center.

To guaranty Quorum disk functionality, it is recommended to reset Cluster FailoverThreshold count to 5 to 10. This FailoverThresholdSetting value 4294967295 (FFFF FFFF) is a default value, which is 1 by default. This might lead to Quorum disk failure during migration.

You can modify this value for cluster group by running the following command (run as admin):

```
(Get-ClusterGroup -Name "Cluster Group").FailoverThreshold=5
```

You can verify the change by running the following PS command:

```
get-clustergroup -Name "Cluster Group" | fl *
```

A message similar to the following example displays:

```
PS C:\Users\administrator.WINELAB> get-clustergroup -name "Cluster Group" | fl *

AntiAffinityClassNames : {}
AutoFailbackType       : 0
ColdStartSetting       : 0
Cluster                : win2019_FC
DefaultOwner           : 4294967295
Description            :
GroupType              : Cluster
FailoverPeriod         : 6
FailoverThreshold      : 5
FailbackWindowEnd      : 4294967295
FailbackWindowStart    : 4294967295
FaultDomain            : 0
IsCoreGroup            : True
Name                   : Cluster Group
OwnerNode              : E2E-L4-9750
PersistentState        : 1
PlacementOptions       : 0
PreferredSite          : {}
Priority                : 13000
ResiliencyPeriod       : 0
State                  : Online
StatusInformation      : 0
UpdateDomain           : 0
Id                     : a4d75d0b-4e0b-4dca-8f88-9722038a1730
```

## Operating system characteristics

Most host applications will behave in a similar manner in comparison to the normal devices when writing to thin devices. This same behavior can also be observed as long as the thin device written capacity is less than thin device subscribed capacity. However, issues can arise when the application writes beyond the provisioned boundaries. With the current behavior of the Windows operating system, the exhaustion of the thin pool causes the following undesired results:

- Logical Volume Manager software SVM and VxVM cannot write to any volumes that are built on the exhausted pool.
- In the Windows NTFS file system:
  - The host reports File System is full errors to the Windows system event log. The larger the data file size that is being written to the thin device, the more File System is full error messages will be reported.
  - The writing data file has corrupted data.
  - Cannot create a file system on the exhausted pool.
  - Cannot write a data file to the exhausted pool.

if the host is exposed to preprovisioned thin devices that had not been bound to the thin pool, the host might take a bit longer to boot.

- Windows Resilient File System (ReFS) support starts with Windows Server 2012 and later. ReFS is designed to scale past what NTFS supports.

File system	Maximum file size	Maximum volume size
NTFS	256 TB	256 TB
ReFS	35 PB	35 PB

- Offloaded Data Transfer (ODX) is not supported with ReFS.
- Thin provisioning is supported in Storage Spaces only, and not supported in SAN arrays.

For more information on ReFS, see [Resilient File System overview](#) in the Windows IT Pro Center.

## Hyper-V

Hyper-V in Windows Server enables you to create a virtualized server computing environment. You can improve the efficiency of your hardware resources by using Hyper-V to create and manage virtual machines and their resources. Each virtual machine is a self-contained virtualized computer system that operates in an isolated environment. This enables multiple operating systems to run simultaneously on one physical computer. Hyper-V is an available role in Windows Server 2008 and later.

For information on Hyper-V, including features, benefits, and installation procedures, see [Microsoft](#) website and [Microsoft TechNet](#).