

**EMC[®] VNX[™] Series MPFS over FC and iSCSI
v6.0 Windows Clients**

Product Guide

**P/N 300-012-183
REV A01**

EMC Corporation
Corporate Headquarters:
Hopkinton, MA 01748-9103
1-508-435-1000
www.EMC.com

Copyright © 2007-2011 EMC Corporation. All rights reserved.

Published February 2011

EMC 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." EMC CORPORATION 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 EMC software described in this publication requires an applicable software license.

For the most up-to-date regulatory document for your product line, go to the Technical Documentation and Advisories section on the EMC Online Support website at Support.EMC.com.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com.

All other trademarks used herein are the property of their respective owners.

Preface

Chapter 1 Introducing EMC VNX MPFS over FC and iSCSI

Overview of MPFS over FC and iSCSI	18
EMC VNX MPFS architectures	19
VNX MPFS over Fibre Channel.....	19
VNX MPFS over iSCSI	21
VNX MPFS over iSCSI/FC.....	22
How VNX MPFS works	24

Chapter 2 EMC VNX MPFS Environment Configuration

Configuration roadmap	26
Implementation guidelines.....	28
VNX with MPFS recommendations.....	28
Storage configuration recommendations	29
MPFS feature configurations.....	30
MPFS installation and configuration process	32
Configuration planning checklist.....	33
Verifying system components.....	35
Required hardware components	35
Required software components	37
Verifying configuration.....	37
Verifying system requirements	38
Verifying the FC switch requirements (FC configuration) ..	39
Verifying the IP-SAN VNX for block requirements.....	40
Setting up the VNX for file	41
Setting up the file system.....	42
File system prerequisites	42

Creating a file system on a VNX for file.....	43
Enabling MPFS for the VNX for file.....	53
Configuring the VNX for block by using CLI commands.....	54
Best practices for VNX for block and VNX VG2/VG8 configurations	54
Configuring the SAN and storage.....	55
Installing the FC switch (FC configuration)	55
Zoning the SAN switch (FC configuration).....	55
Creating a security file on the VNX for file	56
Configuring the VNX for block iSCSI port	57
Access Logix configuration.....	60
Configuring and accessing storage	64
Installing the FC driver (FC configuration)	64
Adding hosts to the storage group (FC configuration).....	65
Installing and configuring the Microsoft iSCSI initiator (iSCSI configuration).....	67
Adding initiators to the storage group (FC configuration)	74
Adding initiators to the storage group (iSCSI configuration).....	76
Mapping an MPFS share to a network drive.....	79
Chapter 3	Installing, Upgrading, or Uninstalling VNX MPFS Software
Installing the MPFS software	86
Before installing.....	86
Installing the MPFS software	87
Operating MPFS through a firewall	90
IPV6 supported Microsoft Windows platforms.....	91
Upgrading the MPFS software	92
Attempting to install over existing MPFS software	92
Installing a later version of MPFS software.....	93
Uninstalling the MPFS software.....	94
Disabling an active MPFS software session.....	97

Chapter 4	EMC VNX MPFS Command Line Interface	
	Overview of the command line interface	100
	mpfscopy	101
	mpfsctl -?	103
	mpfsctl config	104
	mpfsctl disable.....	104
	mpfsctl diskspeed	105
	mpfsctl enable.....	105
	mpfsctl inq.....	106
	mpfsctl list	107
	mpfsctl reset	107
	mpfsctl stats	108
	mpfsctl version	109
	mpfsinfo	110
	mpfsinq.....	112
Appendix A	Error Messages and Troubleshooting	
	EMC VNX MPFS messages sent to the event log	114
	Messages displayed in error message boxes	115
	Using Windows Event Viewer with MPFS	116
	Log type: Warning	116
	Log type: Error	116
	Log type: Information	116
	Troubleshooting MPFS software	121
	Confirming MPFS software installation	121
	Installing difficulties.....	122
	Uninstalling difficulties.....	122
	Known problems and limitations	122
	Glossary	
	Index	

	Title	Page
1	VNX over FC.....	20
2	VNX VG2/VG8 over FC.....	20
3	VNX over iSCSI	21
4	VNX VG2/VG8 over iSCSI.....	22
5	VNX over iSCSI/FC.....	23
6	VNX VG2/VG8 over iSCSI/FC.....	23
7	Configuration roadmap.....	27
8	Log On to Target dialog box.....	72
9	Target Properties dialog box	73
10	Opening a VNX for file to view possible shares	79
11	List of Available Shares	80
12	Selecting the MPFS Volume Properties.....	80
13	MPFS Properties window	81
14	Advanced Option Dialog window	82
15	Global Options window	83
16	Information window before Install.....	88
17	EMC MPFS Installer dialog box.....	88
18	EMC MPFS is Installed dialog box	89
19	Previous installation detected window.....	92
20	EMC MPFS Installer window with Exit button active.....	92
21	Opening the Control Panel	94
22	Control Panel window.....	95
23	Add or Remove Programs window.....	95
24	EMC MPFS Uninstall completed message window.....	96
25	MPFS Properties window	97
26	Global Options window	104
27	Opening the Windows Event Viewer.....	117
28	Event Viewer window	118
29	Selecting an Error message	118
30	Event Properties window.....	119

31	Up and down arrows on Event Properties window	120
32	Using mpfsctl version to verify MPFS software is installed.....	121

	Title	Page
1	Prefetch and read cache requirements	29
2	Arraycompath and failovermode settings	64
3	MPFS executable files with Windows OS versions	87
4	Windows server firewall ports.....	90
5	Windows command line interface commands	100
6	MPFS iSCSI error messages	114
7	MPFS FC error messages.....	114
8	Windows messages displayed in message boxes.....	115

As part of an effort to improve and enhance the performance and capabilities of its product lines, EMC periodically releases versions of its hardware and software. Therefore, some functions described in this document may not be supported by all versions of the software or hardware currently in use. For the most up-to-date information on product features, refer to your product release notes.

If a product does not function properly or does not function as described in this document, contact your EMC representative.

Review the EMC Online Support website, <http://Support.EMC.com>, to ensure that you have the latest versions of the MPFS software and documentation.

For software, open Support > Software Downloads and Licensing > Downloads V > and then select the necessary software for VNX MPFS from the menu.

For documentation, open Support > Technical Documentation > Hardware/Platforms > VNX Series.

Note: Only registered EMC Online Support users can download the MPFS software.

Audience This document is part of the EMC VNX MPFS documentation set, and is intended for use by Windows system administrators responsible for installing and maintaining VNX Windows servers.

Readers of this document are expected to be familiar with the following topics:

- ◆ VNX for block or EMC Symmetrix
- ◆ VNX for file
- ◆ CIFS protocol
- ◆ Windows operating system
- ◆ Operating environments to install the Microsoft Windows Server include:
 - Windows Server 2003 Standard/Enterprise
 - Windows Server 2003 R2 Standard/Enterprise
 - Windows Server 2008 Standard/Enterprise
 - Windows Server 2008 R2 Standard/Enterprise (64-bit only)
 - Windows Server XP Professional
 - Windows Vista Enterprise
 - Windows 7 Professional/Enterprise

Note: In this document, *Windows* refers to Microsoft Windows Server 2003, Microsoft Windows Server 2008, Microsoft Windows Server XP, Microsoft Windows Vista, or Microsoft Windows 7. The *EMC VNX MPFS for Windows Clients Release Notes* provide a complete listing of specific release information.

Related documentation

Related documents include:

- ◆ *EMC VNX MPFS for Windows Clients Release Notes*
- ◆ *EMC Host Connectivity Guide for VMWare ESX Server*
- ◆ EMC documentation for HBAs

VNX for block

- ◆ *Removing ATF or CDE Software before Installing other Failover Software*
- ◆ Unisphere online help

Symmetrix

- ◆ Symmetrix product manual

VNX for file

- ◆ *EMC VNX Documentation*
- ◆ *Using VNX Multi-Path File System*

All of these publications are found on the EMC Online Support website.

EMC Online Support

The EMC Online Support website provides the most up-to-date information on documentation, downloads, interoperability, product lifecycle, target revisions, and bug fixes. As a registered EMC Online Support user, you can subscribe to receive notifications when updates occur.

E-Lab Interoperability Navigator

The EMC E-Lab Interoperability Navigator tool provides access to EMC interoperability support matrices. After logging in to EMC Online Support, go to Support > Interoperability and Product Lifecycle Information > E-Lab Interoperability Navigator.

Conventions used in this document

EMC uses the following conventions for special notices.

Note: A note presents information that is important, but not hazard-related.



CAUTION

A caution contains information essential to avoid data loss or damage to the system or equipment.



IMPORTANT

An important notice contains information essential to operation of the software.



WARNING

A warning contains information essential to avoid a hazard that can cause severe personal injury, death, or substantial property damage if you ignore the warning.



DANGER

A danger notice contains information essential to avoid a hazard that will cause severe personal injury, death, or substantial property damage if you ignore the message.

Typographical conventions

EMC uses the following type style conventions in this document:

Normal	Used in running (nonprocedural) text for: <ul style="list-style-type: none"> Names of interface elements (such as names of windows, dialog boxes, buttons, fields, and menus) Names of resources, attributes, pools, Boolean expressions, buttons, DQL statements, keywords, clauses, environment variables, functions, utilities URLs, pathnames, filenames, directory names, computer names, links, groups, service keys, file systems, notifications
Bold:	Used in running (nonprocedural) text for: <ul style="list-style-type: none"> Names of commands, daemons, options, programs, processes, services, applications, utilities, kernels, notifications, system call, man pages Used in procedures for: <ul style="list-style-type: none"> Names of interface elements (such as names of windows, dialog boxes, buttons, fields, and menus) What user specifically selects, clicks, presses, or types
<i>Italic:</i>	Used in all text (including procedures) for: <ul style="list-style-type: none"> Full titles of publications referenced in text Emphasis (for example a new term) Variables
<code>Courier:</code>	Used for: <ul style="list-style-type: none"> System output, such as an error message or script URLs, complete paths, filenames, prompts, and syntax when shown outside of running text
Courier bold:	Used for: <ul style="list-style-type: none"> Specific user input (such as commands)
<i>Courier italic:</i>	Used in procedures for: <ul style="list-style-type: none"> Variables on command line User input variables
< >	Angle brackets enclose parameter or variable values supplied by the user
[]	Square brackets enclose optional values
	Vertical bar indicates alternate selections - the bar means "or"
{ }	Braces indicate content that you must specify (that is, x or y or z)
...	Ellipses indicate nonessential information omitted from the example

- Where to get help** EMC support, product, and licensing information is obtained as follows.
- Product information** — For documentation, release notes, software updates, or for information about EMC products, licensing, and service, go to the EMC Online Support website (registration required) at:
- <http://Support.EMC.com>
- Technical support** — For technical support, go to EMC Customer Service on EMC Online Support. To open a service request through EMC Online Support, you must have a valid support agreement. Contact your EMC Customer Support Representative for details about obtaining a valid support agreement or to answer any 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 opinion of this document to:
- techpubcomments@EMC.com

Introducing EMC VNX MPFS over FC and iSCSI

This chapter provides an overview of EMC VNX MPFS over FC and iSCSI and its architecture. The chapter includes the following topics:

- ◆ Overview of MPFS over FC and iSCSI..... 18
- ◆ EMC VNX MPFS architectures 19
- ◆ How VNX MPFS works 24

Overview of MPFS over FC and iSCSI

EMC® VNX™ series Multi-Path File System (MPFS) over Fibre Channel (FC) lets Linux, Windows, UNIX, AIX, or Solaris servers access shared data concurrently over FC connections, whereas VNX MPFS over Internet small computer system interface (iSCSI) lets servers access shared data concurrently over an iSCSI connection.

MPFS uses common Internet Protocol local area network (IP LAN) topology to transport data and metadata to and from the servers.

Without the MPFS file system, servers can access shared data by using standard network file system (NFS) or Common Internet File System (CIFS) protocols; The MPFS file system accelerates data access by providing separate transports for file data (file content) and metadata (control data).

For an FC-enabled server, data is transferred directly between the Windows server and storage over an FC storage area network (SAN).

For an iSCSI-enabled server, data is transferred over the IP LAN between the Windows server and storage for a VNX or VNX VG2/VG8 configuration.

Metadata passes through the VNX for file (and the IP network), which includes the network-attached storage (NAS) portion of the configuration.

EMC VNX MPFS architectures

Three basic VNX MPFS architectures are available:

- ◆ VNX MPFS over FC
- ◆ VNX MPFS over iSCSI
- ◆ VNX MPFS over iSCSI/FC

The FC architecture consists of the following configurations:

- EMC VNX5300, 5500, 5700, or 7500 with FC
- VNX VG2/VG8 with FC

The iSCSI architecture consists of the following configurations:

- VNX5300, 5500, 5700, or 7500 with iSCSI
- VNX VG2/VG8 with iSCSI

The iSCSI/FC architecture consists of the following configurations:

- VNX with iSCSI/FC
- VNX VG2/VG8 with iSCSI/FC

VNX MPFS over Fibre Channel

The VNX MPFS over FC architecture consists of the following:

- ◆ VNX with MPFS — A NAS device configured with a VNX and MPFS software
- ◆ VNX for block or Symmetrix®
- ◆ Windows servers with MPFS software connected to a VNX through the IP LAN, VNX for block, or Symmetrix by using FC architecture

[Figure 1 on page 20](#) shows the VNX over FC configuration where the Windows servers are connected to a VNX Series (VNX for file and VNX for block in a single cabinet) by using an IP switch and one or more FC or FC over Ethernet (FCoE) switches. In a smaller configuration of one or two servers, the servers are connected directly to the VNX Series without the use of FC or FCoE switches.




Figure 1 VNX over FC

Figure 2 on page 20 shows the VNX VG2/VG8 over FC configuration. In this figure, the Windows servers are connected to a VNX for block or Symmetrix by using a VNX for file, IP switch and optional FC switch or FCoE switch.




Figure 2 VNX VG2/VG8 over FC

VNX MPFS over iSCSI

The VNX MPFS over iSCSI architecture consists of the following:

- ◆ VNX with MPFS — A NAS device that is configured with a VNX and MPFS software
- ◆ VNX for block or Symmetrix
- ◆ Windows server with MPFS software connected to a VNX through the IP LAN, VNX for block, or Symmetrix by using iSCSI architecture

Figure 3 on page 21 shows the VNX over iSCSI configuration where the Windows servers are connected to a VNX Series by using one or more IP switches.




Figure 3 VNX over iSCSI

Figure 4 on page 22 shows the VNX VG2/VG8 over iSCSI configuration where the Windows servers are connected to a VNX for block or Symmetrix with a VNX for file by using one or more IP switches.




Figure 4 VNX VG2/VG8 over iSCSI

VNX MPFS over iSCSI/FC

The VNX MPFS over iSCSI/FC architecture consists of the following:

- ◆ VNX with MPFS — A NAS device that is configured with a VNX and MPFS software
- ◆ VNX for block or Symmetrix
- ◆ Windows server with MPFS software connected to a VNX through the IP LAN, VNX for block, or Symmetrix by using iSCSI/FC architecture

Figure 5 on page 23 shows the VNX over iSCSI/FC configuration where the Windows servers are connected to a VNX for file by using one or more IP switches and an FC switch or FCoE switch.




Figure 5 VNX over iSCSI/FC

Figure 6 on page 23 shows the VNX VG2/VG8 over iSCSI/FC configuration where the Windows servers are connected to a VNX for block or Symmetrix with a VNX for file by using one or more IP switches and an FC switch or FCoE switch.




Figure 6 VNX VG2/VG8 over iSCSI/FC

How VNX MPFS works

Although called a file system, the VNX MPFS is neither a new nor a modified format for storing files. Instead, the MPFS file system interoperates and uses the standard NFS and CIFS protocols to enforce access permissions. The MPFS file system uses a protocol called File Mapping Protocol (FMP) to exchange metadata between the Windows server and the VNX for file.

All requests unrelated to file I/O pass directly to the NFS/CIFS layer. The MPFS layer intercepts only the open, close, read, and write system calls.

When a Windows server intercepts a file-read call, it sends a request to the VNX for file asking for the file's location. The VNX for file responds with a list of file extents, which the Windows server then uses to read the file data directly from the disk.

When a Windows server intercepts a file-write call, it asks the VNX for file to allocate blocks on disk for the file. The VNX for file allocates the space in contiguous extents and sends the extent list to the Windows server. The Windows server then writes data directly to disk, informing the VNX for file when finished, so that the VNX for file can permit other Windows servers to access the file.

The remaining chapters describe how to install, manage, and tune Windows servers. *Using VNX Multi-Path File System* technical module, available on EMC Online Support at <http://Support.EMC.com>, provides information on the MPFS commands.

This chapter presents a high-level overview of configuring and installing the EMC VNX MPFS.

Topics include:

◆ Configuration roadmap	26
◆ Implementation guidelines.....	28
◆ MPFS installation and configuration process	32
◆ Verifying system components	35
◆ Setting up the VNX for file	41
◆ Setting up the file system.....	42
◆ Enabling MPFS for the VNX for file.....	53
◆ Configuring the VNX for block by using CLI commands	54
◆ Configuring the SAN and storage.....	55
◆ Configuring and accessing storage	64
◆ Mapping an MPFS share to a network drive.....	79

Configuration roadmap

[Figure 7 on page 27](#) shows the roadmap for configuring and installing the EMC VNX MPFS over FC and iSCSI architecture for both FC and iSCSI environments. The roadmap contains the topics representing sequential phases of the configuration and installation process. The descriptions of each phase, which follow, contain an overview of the tasks required to complete the process, and a list of related documents for more information.




Figure 7 Configuration roadmap

Implementation guidelines

The following MPFS implementation guidelines are valid for all MPFS installations.

VNX with MPFS recommendations

The following recommendations are described in detail in the *EMC VNX MPFS over iSCSI Applied Best Practices Guide* and the *VNX Best Practices for Performance*, which are found on EMC Online Support at <http://Support.EMC.com>:

- ◆ MPFS is optimized for large I/O transfers and is useful for workloads with average I/O sizes as small as 16 KB. However, MPFS has been shown conclusively to improve performance for I/O sizes of 128 KB and greater.
- ◆ For best MPFS performance, in most cases, configure the VNX for file volumes by using a volume stripe size of 256 KB.
- ◆ EMC PowerPath® is supported, but is not recommended. Path failover is built into the Windows server. When using PowerPath, the performance of the MPFS system is lower. Knowledgebase article emc 165953 provides details on using PowerPath and MPFS.
- ◆ When MPFS is started, 16 threads are run, which is the default number of MPFS threads. The maximum number of threads is 128. If system performance is slow, gradually increase the number of threads allotted for the Data Mover to improve system performance. Add threads conservatively, as the Data Mover allocates 16 KB of memory to accommodate each new thread. The optimal number of threads depends on the network configuration, the number of Windows servers, and the workload.

Using VNX Multi-Path File System provides the procedures necessary to adjust the thread count and is available with the *EMC Documentation* on EMC Online Support.

Data Mover capacity

The *EMC Support Matrix* provides Data Mover capacity guidelines. After logging in to EMC Online Support, go to Support > Interoperability and Product Lifecycle Information > Interoperability Matrices.

Windows server configuration

All Windows servers using the MPFS software require:

- ◆ At least one FC connection or an iSCSI initiator connected to a SAN switch, or directly to a VNX for block or Symmetrix
- ◆ Network connections to the Data Mover

Note: When deploying MPFS over iSCSI on a VNX5300, VNX5500, VNX5700, VNX7500 or a VNX VG2/VG8 configuration based on the iSCSI-enabled VNX for block, the VNX for block iSCSI target is used.

Storage configuration recommendations

Windows servers read and write directly from a VNX for block. This has several implications:

- ◆ Use the VNX Operating Environment (VNX OE) for best performance in new MPFS configurations. The *VNX Best Practices for Fibre Channel Storage: VNX OE Firmware Update* provides more details.
- ◆ Unmount all mounted MPFS file systems from the Windows server before changing any storage device or switch configuration.

[Table 1 on page 29](#) lists the prefetch and read cache requirements.

Table 1

Prefetch and read cache requirements

Prefetch requirements	Read cache	Notes
Modest	50–100 MB	80% of the systems fall under this category.
Heavy	250 MB	Requests greater than 64 KB and sequential reads from many LUNs expected over 300 MB/s.
Extremely heavy	1 GB	120 or more drives reading in parallel.

MPFS feature configurations

iSCSI CHAP authentication

The following sections describe the configurations for MPFS features.

The Windows server with MPFS software and the VNX for block support the Challenge Handshake Authentication Protocol (CHAP) for iSCSI network security.

CHAP provides a method for the Windows server and VNX for block to authenticate each other through an exchange of a shared secret (a security key that is similar to a password), which is typically a string of 12 to 16 bytes.



CAUTION

If CHAP security is not configured for the VNX for block, any computer connected to the same IP network as the VNX for block iSCSI ports can read from or write to the VNX for block.

CHAP has two variants — One-way and reverse CHAP authentication:

- ◆ In one-way CHAP authentication, CHAP sets up the accounts that the Windows server uses to connect to the VNX for block. The VNX for block authenticates the Windows server.
- ◆ In reverse CHAP authentication, the VNX for block authenticates the Windows server and the Windows server also authenticates the VNX for block.

Because CHAP secrets are shared between the Windows server and VNX for block, the CHAP secrets are configured the same on both the Windows server and VNX for block.

The *CX-Series iSCSI Security Setup Guide* provides detailed information regarding CHAP and is found on the EMC Online Support website.

**VMware ESX
(optional)**

VMware is a software suite for optimizing and managing IT environments through virtualization technology. MPFS supports the Windows server guest operating systems running on a VMware ESX server.

The VMware ESX server is a robust, production-proven virtualization layer that abstracts processor, memory, storage, and networking resources into multiple virtual machines (software representation of a physical machine) running side-by-side on the same server.

VMware is not tied to any operating system, giving customers a bias-free choice of operating systems and software applications. All operating systems supported by VMware are supported with both iSCSI and NFS protocols for basic connectivity. This allows several instances of similar and different guest operating systems to run as virtual machines on one physical machine.

To run a Windows server guest operating system on a VMware ESX server, the configuration must meet the following requirements:

- ◆ Run a supported version of the Windows operating system.
- ◆ Have the VNX for block supported HBA hardware and driver installed.
- ◆ Connect to each SP in each VNX for block directly or through a switch. Each SP must have an IP connection.
- ◆ Connect to a TCP/IP network with both SPs in the VNX for block.

Currently, the VMware ESX server has the following limitations:

- ◆ Booting the guest Windows server off iSCSI is not supported.
- ◆ PowerPath is not supported.
- ◆ Virtual machines that run the Windows server guest operating system must use iSCSI to access the VNX for block.
- ◆ Store the virtual machine on a VMware datastore (VNX for block or Symmetrix) and access it by the VMware ESX server by using either FC (ESX server versions 3.0.1, 3.0.2, or 3.5.1) or iSCSI (ESX server version 3.5.1).

The *EMC Host Connectivity Guide for VMWare ESX Server* provides information on how to configure iSCSI initiator ports and how VMware operates in a Windows environment. The VMware website, at <http://www.vmware.com>, provides more information.

MPFS installation and configuration process

The MPFS configuration process involves performing tasks on various system components in a specific order.

Note: This document provides guidelines for installing and configuring MPFS with several options. Disregard steps that do not pertain to your environment.

To manually install and configure MPFS:

1. Collect installation and configuration planning information and complete the checklist:
 - a. Collect the IP network addresses, FC port addresses, and VNX for block or Symmetrix information.
 - b. Map the Ethernet and TCP/IP network topology.
 - c. Map the FC zoning topology.
 - d. Map the virtual storage area network (VSAN) topology.
2. Install the MPFS software manually (on a native or VMware¹ hosted Windows operating system):
 - a. Install the HBA driver (for FC configuration).
 - b. Install and configure the iSCSI (for iSCSI configuration).²
 - c. Start the iSCSI service (for iSCSI configuration).
 - d. Install the MPFS software.
 - e. Verify the MPFS software configuration.

1. “VMware ESX (optional)” on page 31 provides information.
2. *Installing VNX iSCSI Host Components* provides details.

Configuration planning checklist

Collect the following information before beginning the MPFS installation and configuration process.

Note: The use of dynamic host configuration protocol (DHCP) to obtain IP addresses is not supported. Use static IP addresses.

For an FC and iSCSI configuration:

- SP A IP address
- SP A login name.....
- SP A password
- SP B IP address.....
- SP B login name
- SP B password.....
- Zoning for Data Movers
- First Data Mover LAN blade IP address or Data Mover IP address.....
- Second Data Mover LAN blade IP address or Data Mover IP address.....
- Control Station IP address or CS address
- LAN IP address (same as LAN Data Movers).....
- Windows server IP address on LAN
- VSAN name
- VSAN number (ensure that it is not in use).....

For an FC configuration:

- SP A FC port assignment or FC ports.....
- SP B FC port assignment or FC ports.....
- FC switch name.....
- FC switch password
- FC switch port IP address.....
- Zoning for each FC HBA port.....
- Zoning for each FC director

For an iSCSI configuration:

- VNX with MPFS target IP address
- VNX for block or Symmetrix target IP address
- Windows server IP address for iSCSI Gigabit connection
- Initiator and Target Challenge Handshake Authentication Protocol (CHAP) Password (optional).....

Verifying system components

MPFS environments require standard VNX for file hardware and software, with the addition of a few components that are specific to either FC or iSCSI configurations. This involves setting up an MPFS environment to verify that each of the previously mentioned components is in place and functioning normally. Each hardware and software component is discussed in the following sections.

Required hardware components

VNX over FC configuration

This section lists the MPFS configurations with the required hardware components.

The hardware components for a VNX over FC configuration are:

- ◆ A VNX Series connected to an FC network and SAN
- ◆ An IP switch connecting the VNX Series to the servers
- ◆ An FC switch or FCoE switch with an HBA for each Windows server

[“VNX MPFS over Fibre Channel” on page 19](#) provides more information.

VNX VG2/VG8 over FC configuration

The hardware components for a VNX VG2/VG8 over FC configuration are:

- ◆ A VNX for file connected to an FC network and SAN
- ◆ A fabric-connected VNX for block or Symmetrix with available LUNs
- ◆ An IP switch connecting the VNX for file to the servers
- ◆ An FC switch or FCoE switch with an HBA for each Windows server

[“VNX MPFS over Fibre Channel” on page 19](#) provides more information.

VNX over iSCSI configuration

The hardware components for an VNX over iSCSI configuration are:

- ◆ A VNX Series
- ◆ One or two IP switches connecting the VNX Series to the servers

[“VNX MPFS over iSCSI” on page 21](#) provides more information.

VNX VG2/VG8 over iSCSI configuration

The hardware components for a VNX VG2/VG8 over iSCSI configuration are:

- ◆ A VNX for file connected to an FC network and SAN
- ◆ A fabric-connected VNX for block or Symmetrix with available LUNs
- ◆ One or two IP switches connecting the VNX for file and the VNX for block or Symmetrix to the servers

“[VNX MPFS over iSCSI](#)” on page 21 provides more information.

VNX over iSCSI/FC configuration

The hardware components for a VNX over iSCSI/FC configuration are:

- ◆ A VNX for file
- ◆ One or two IP switches and an FC switch or FCoE switch connecting the VNX for file to the servers

“[VNX MPFS over iSCSI/FC](#)” on page 22 provides more information.

VNX VG2/VG8 over iSCSI/FC configuration

The hardware components for a VNX VG2/VG8 over iSCSI/FC configuration are:

- ◆ A VNX for file connected to an FC network and SAN
- ◆ A fabric-connected VNX for block or Symmetrix with available LUNs
- ◆ One or two IP switches and an FC switch or FCoE switch connecting the VNX for file and the VNX for block or Symmetrix to the servers

“[VNX MPFS over iSCSI/FC](#)” on page 22 provides more information.

Configuring Gigabit Ethernet ports

Two Gigabit Ethernet NICs, or a multiport NIC with two available ports, connected to isolated IP networks or subnets are recommended for each Windows server for iSCSI. For each Windows server for FC, one NIC is required for CIFS and FMP traffic. For maximum performance, use:

- ◆ One port for the connection between the Windows server and the Data Mover for MPFS metadata transfer and CIFS traffic

Note: The same NIC port contains the CIFS and FMP traffic. This NIC is first in the binding order as well.

- ◆ One port for the connection between the Windows server and the same subnet as the iSCSI discovery address dedicated to data transfer

Note: The second NIC for iSCSI must be on the same subnet as the discovery address.

Configuring and Managing EMC VNX Networking provides detailed information for setting up network connections and is available on the EMC Online Support website.

Required software components

The following software components are required for an MPFS configuration:

Note: The *EMC VNX MPFS for Windows Clients Release Notes* provide a complete list of EMC supported operating system versions.

- ◆ NAS software version that supports either FC or iSCSI on Windows platforms
- ◆ Windows operating system version that supports HBAs or an iSCSI initiator

Note: The EMC E-Lab™ Interoperability Matrix lists the latest versions of Windows operating systems.

- ◆ MPFS software version 5.0 or later
- ◆ Windows iSCSI initiator 2.03 or later
- ◆ Java Runtime Environment 1.5.0.11 or later to run the console applications

Verifying configuration

The next step in setting up MPFS is to verify whether each of the previously mentioned components is in place and functioning normally. If each of these components is operational, the [“MPFS installation and configuration process” on page 32](#) provides more information.

Configure CIFS and start the services on the VNX for file that is used for MPFS connectivity.

Related documentation

The following documents, available on the EMC Online Support website, provide additional information:

- ◆ *Configuring and Managing EMC VNX Networking*
- ◆ *Managing VNX Volumes and File Systems Manually*
- ◆ *Configuring Standbys on VNX*
- ◆ *Configuring CIFS on VNX*

Verifying system requirements

This section describes system requirements for an MPFS environment. The documents listed in “[Related documentation](#)” on [page 39](#) detail system setup information.

**CAUTION**

Ensure that the systems used for MPFS do not contain both VNX for block and Symmetrix LUNs. MPFS does not support a mixed storage environment.

System requirements

VNX for file and VNX VG2/VG8 used within an MPFS environment must meet these requirements:

- ◆ Use only the VNX for file designed for MPFS file systems. The following models are supported:
 - VNX5300
 - VNX5500
 - VNX5700
 - VNX7500
- ◆ Ensure that all MPFS system environments have file systems built on disks from only one type: either all FC, SAS or ATA drives, not a mixture of FC, SAS, or ATA drives.
- ◆ Ensure that MPFS does not use a file system spanning across two different system enclosures.
- ◆ Build LUNs by using RAID 1, RAID 3, RAID 5, or RAID 6 only.
- ◆ Build Management LUNs by using 4+1 RAID 5 only.
- ◆ Enable write cache.
- ◆ Use EMC Access Logix™.
- ◆ Run VNX OE with NAS 7.0.x or later.

All Symmetrix storage used within an MPFS environment must meet these requirements:

- ◆ Use only the following Symmetrix models designed for MPFS file systems:
 - Symmetrix DMX™ series Enterprise Storage Platform (ESP)
 - Symmetrix VMAX™ series
 - Symmetrix 8000 series
- ◆ Ensure that the correct version of the microcode is used. Contact your EMC Customer Support Representative or review the EMC E-Lab Interoperability Navigator for microcode release updates.
- ◆ Ensure that the Symmetrix FC/SCSI port flags are properly configured for the MPFS file system. Set the Avoid_Reset_Broadcast (ARB) for each port connected to a Windows server.
- ◆ Ensure that MPFS does not use a file system spanning across two different system enclosures.

Related
documentation

Rails and enclosures documentation available on the EMC Online Support website provides additional information.

Verifying the FC
switch requirements
(FC configuration)

Ensure that the following to set up the FC switch:

- ◆ Install the FC switch.
- ◆ Verify that the host bus adapter (HBA) driver is loaded.
To verify that the HBA driver is loaded, select Start > Run and type compmgmt.msc in the window. In the Explorer window, select Device Manager > Disk drives.
- ◆ Connect cables from each HBA FC port to a switch port.
- ◆ Verify the HBA connection to the switch by checking LEDs for the switch port connected to the HBA port.
- ◆ Configure zoning for the switch as described in [“Zoning the SAN switch \(FC configuration\)” on page 55](#).

Note: Configure zoning as *single initiator*, meaning that each HBA port has its own zone in which it is the only HBA port.

Related
documentation

The documentation that ships with the FC switch provides more information about the switch.

Verifying the IP-SAN VNX for block requirements

The VNX MPFS over FC and iSCSI environment with VNX for block configurations requires the following:

- ◆ For a VNX for file configuration, a VNX5300, VNX5500, VNX5700, or VNX7500.
- ◆ For a VNX VG2/VG8 configuration:
 - VNX5300, VNX5500, VNX5700, or VNX7500.
 - VNX for block, Symmetrix DMX, Symmetrix VMAX, or Symmetrix 8000.
 - Cabled as any shared VNX for block.
 - Access Logix LUN masking by using iSCSI to present all managed LUNs to the Windows servers.
- ◆ Windows server configuration is the same as a standard Windows server connection to an iSCSI connection.
- ◆ Windows servers are load balanced across VNX for block iSCSI ports for performance improvement and protection against single-port and Ethernet cable problems.
- ◆ Port 0 iSCSI through port 3 iSCSI on each storage processor is connected to the iSCSI network.

Setting up the VNX for file

The *VNX System Software Installation Guide* provides information on how to set up the VNX for file, which is available on the EMC Online Support website.

Setting up the file system

This section describes the prerequisites for file systems and the procedure for creating a file system.

File system prerequisites

File system prerequisites are guidelines to met before building a file system. A properly built file system must:

- ◆ Use disk volumes from the same VNX for block.

Note: Do not use a file system spanning across two system enclosures. A file system spanning multiple systems is not supported even if the multiple systems are of the same type, such as VNX for block or Symmetrix.

- ◆ Use disk volumes from the same disk type, all FC, SAS, or ATA, not a mixture of FC, SAS, or ATA.
- ◆ For best MPFS performance, in most cases, configure the volumes by using a volume stripe size of 256 KB. The *EMC VNX MPFS over iSCSI Applied Best Practices Guide* provides detailed performance related information.
- ◆ In a Symmetrix environment, ensure that the Symmetrix FC/SCSI port flag settings are properly configured for the MPFS file system; in particular, set the ARB flag. The EMC Customer Support Representative configures these settings.

Creating a file system on a VNX for file

This section describes how to configure, create, mount, and export file systems.

Ensure that systems are created optimally for MPFS. All LUNs must:

- ◆ Use the same RAID type
- ◆ Have the same number of spindles in each RAID group
- ◆ Contain spindles of the same type and speed

In addition, ensure that all LUNs do not share spindles with:

- ◆ Other LUNs in the same file system
- ◆ Another file system heavily utilized by high-I/O applications

Before creating the LUNs, ensure that the total usable capacity of all the LUNs within a single file system does not exceed 16 TB. The maximum number of LUNs tested that are supported in MPFS configurations per file system is 256. Ensure that the LUNs are accessible by the Data Movers through LUN masking, switch zoning, and VSAN settings.

Use this procedure to build or mount the MPFS file system on the VNX for file:

1. Log in to the Control Station as NAS administrator.
2. Before building the file system, type the **nas_disk** command to return a list of unused disks by using this command syntax:

```
$ nas_disk -list |grep n | more
```

For example, type:

```
$ nas_disk -list |grep n | more
```

The output shows all disks not in use:

id	inuse	sizeMB	storageID-devID	type	name	servers
7	n	466747	APM00065101342-0010	CLSTD	d7	1,2
8	n	466747	APM00065101342-0011	CLSTD	d8	1,2
9	n	549623	APM00065101342-0012	CLSTD	d9	1,2
10	n	549623	APM00065101342-0014	CLSTD	d10	1,2
11	n	549623	APM00065101342-0016	CLSTD	d11	1,2
12	n	549623	APM00065101342-0018	CLSTD	d12	1,2
13	n	549623	APM00065101342-0013	CLSTD	d13	1,2
14	n	549623	APM00065101342-0015	CLSTD	d14	1,2
15	n	549623	APM00065101342-0017	CLSTD	d15	1,2
16	n	549623	APM00065101342-0019	CLSTD	d16	1,2
17	n	549623	APM00065101342-001A	CLSTD	d17	1,2
18	n	549623	APM00065101342-001B	CLSTD	d18	1,2
19	n	549623	APM00065101342-001C	CLSTD	d19	1,2
20	n	549623	APM00065101342-001E	CLSTD	d20	1,2
21	n	549623	APM00065101342-0020	CLSTD	d21	1,2
22	n	549623	APM00065101342-001D	CLSTD	d22	1,2
23	n	549623	APM00065101342-001F	CLSTD	d23	1,2
24	n	549623	APM00065101342-0021	CLSTD	d24	1,2
25	n	549623	APM00065101342-0022	CLSTD	d25	1,2
26	n	549623	APM00065101342-0024	CLSTD	d26	1,2
27	n	549623	APM00065101342-0026	CLSTD	d27	1,2
28	n	549623	APM00065101342-0023	CLSTD	d28	1,2
29	n	549623	APM00065101342-0025	CLSTD	d29	1,2
30	n	549623	APM00065101342-0027	CLSTD	d30	1,2

3. Display all disks by using this command syntax:

```
$ nas_disk -list
```

For example, type:

```
$ nas_disk -list
```

Output:

id	inuse	sizeMB	storageID-devID	type	name	servers
1	y	11263	APM00065101342-0000	CLSTD	root_disk	1,2
2	y	11263	APM00065101342-0001	CLSTD	root_disk	1,2
3	y	2047	APM00065101342-0002	CLSTD	d3	1,2
4	y	2047	APM00065101342-0003	CLSTD	d4	1,2
5	y	2047	APM00065101342-0004	CLSTD	d5	1,2
6	y	2047	APM00065101342-0005	CLSTD	d6	1,2
7	n	466747	APM00065101342-0010	CLSTD	d7	1,2
8	n	466747	APM00065101342-0011	CLSTD	d8	1,2
9	n	549623	APM00065101342-0012	CLSTD	d9	1,2
10	n	549623	APM00065101342-0014	CLSTD	d10	1,2
11	n	549623	APM00065101342-0016	CLSTD	d11	1,2
12	n	549623	APM00065101342-0018	CLSTD	d12	1,2
13	n	549623	APM00065101342-0013	CLSTD	d13	1,2
14	n	549623	APM00065101342-0015	CLSTD	d14	1,2
15	n	549623	APM00065101342-0017	CLSTD	d15	1,2
16	n	549623	APM00065101342-0019	CLSTD	d16	1,2
17	n	549623	APM00065101342-001A	CLSTD	d17	1,2
18	n	549623	APM00065101342-001B	CLSTD	d18	1,2
19	n	549623	APM00065101342-001C	CLSTD	d19	1,2
20	n	549623	APM00065101342-001E	CLSTD	d20	1,2
21	n	549623	APM00065101342-0020	CLSTD	d21	1,2
22	n	549623	APM00065101342-001D	CLSTD	d22	1,2
23	n	549623	APM00065101342-001F	CLSTD	d23	1,2
24	n	549623	APM00065101342-0021	CLSTD	d24	1,2
25	n	549623	APM00065101342-0022	CLSTD	d25	1,2
26	n	549623	APM00065101342-0024	CLSTD	d26	1,2
27	n	549623	APM00065101342-0026	CLSTD	d27	1,2
28	n	549623	APM00065101342-0023	CLSTD	d28	1,2
29	n	549623	APM00065101342-0025	CLSTD	d29	1,2
30	n	549623	APM00065101342-0027	CLSTD	d30	1,2

The first stripe alternate SP ownership A,B,A,B,A,B is displayed in bold text and the second stripe alternate SP ownership B,A,B,A,B,A is displayed in a shaded background. The two different stripes (A, B, A) and (B, A, B) are both in RAID group X, Y, and Z.

Note: Use Navicli or EMC Navisphere® Manager to determine which LUNs are on SP A and SP B.

4. Find the names of file systems mounted on all servers by using this command syntax:

```
$ server_df ALL
```

For example, type:

```
$ server_df ALL
```

Output:

```
server_2 :
Filesystem      kbytes      used      avail      capacity  Mounted on
S2_Shgvdms_FS1  831372216   565300   825719208   1%        /root_vdm_5/S2_Shgvdms_FS1
root_fs_vdm_vdm01 114592      7992     106600      7%        /root_vdm_5/.etc
S2_Shg_FS2      831372216  19175496  812196720   2%        /S2_Shg_mnt2
S2_Shg_FS1      1662746472 25312984 1637433488   2%        /S2_Shg_mnt1
root_fs_common   153         5280     10088       34%       /.etc_common
root_fs_2        2581        80496    177632      31%       /
```

```
server_3 :
Filesystem      kbytes      used      avail      capacity  Mounted on
root_fs_vdm_vdm02 114592      7992     106600      7%        /root_vdm_6/.etc
S3_Shgvdms_FS1  831372216  4304736  827067480   1%        /root_vdm_6/S3_Shgvdms_FS1
S3_Shg_FS1      831373240  11675136  819698104   1%        /S3_Shg_mnt1
S3_Shg_FS2      831373240  4204960  827168280   1%        /S3_Shg_mnt2
root_fs_commo   15368       5280     10088       34%       /.etc_common
root_fs_3        258128      8400     249728      3%        /
```

```
vdm01 :
Filesystem      kbytes      used      avail      capacity  Mounted on
S2_Shgvdms_FS1  831372216   5653008   825719208   1%        /S2_Shgvdms_FS1
```

```
vdm02 :
Filesystem      kbytes      used      avail      capacity  Mounted on
S3_Shgvdms_FS1  831372216  4304736   827067480   1%        /S3_Shgvdms_FS1
```

Find the names of file systems mounted on a specific server by using this command syntax:

```
$ server_df <server_name>
```

where:

<server_name> = name of the Windows server

For example, type:

```
$ server_df vdm02
```

Output:

```
vdm02 :
Filesystem      kbytes    used      avail  capacity  Mounted on
S3_Shgvdms_FS1 831372216 4304736   827067480 1%        /S3_Shgvdms_FS1
```

- Find the names of existing file systems that are not mounted by using this command syntax:

```
$ nas_fs -list
```

For example, type:

```
$ nas_fs -list
```

Output:

id	inuse	type	acl	volume	name	server
1	n	1	0	10	root_fs_1	
2	y	1	0	12	root_fs_2	2
3	n	1	0	14	root_fs_3	
4	n	1	0	16	root_fs_4	
5	n	1	0	18	root_fs_5	
6	n	1	0	20	root_fs_6	
7	n	1	0	22	root_fs_7	
8	n	1	0	24	root_fs_8	
9	n	1	0	26	root_fs_9	
10	n	1	0	28	root_fs_10	
11	n	1	0	30	root_fs_11	
12	n	1	0	32	root_fs_12	
13	n	1	0	34	root_fs_13	
14	n	1	0	36	root_fs_14	
15	n	1	0	38	root_fs_15	
16	y	1	0	40	root_fs_common	2
17	n	5	0	73	root_fs_ufslog	
18	n	5	0	76	root_panic_reserve	
19	n	5	0	77	root_fs_d3	
20	n	5	0	78	root_fs_d4	
21	n	5	0	79	root_fs_d5	
22	n	5	0	80	root_fs_d6	
25	y	1	0	116	S2_Shg_FS2	2
221	y	1	0	112	S2_Shg_FS1	2
222	n	1	0	1536	S3_Shg_FS1	
223	n	1	0	1537	S3_Shg_FS2	
384	y	1	0	3026	testdoc_fs2	2

6. Find the names of volumes already mounted by using this command syntax:

```
$ nas_volume -list
```

For example, type:

```
$ nas_volume -list
```

Part of the output is similar to this:

id	inuse	type	acl	name	cltype	clid
1	y	4	0	root_disk	0	1-34,52
2	y	4	0	root_ldisk	0	35-51
3	y	4	0	d3	1	77
4	y	4	0	d4	1	78
5	y	4	0	d5	1	79
6	y	4	0	d6	1	80
7	n	1	0	root_dos	0	
8	n	1	0	root_layout	0	
9	y	1	0	root_slice_1	1	10
10	y	3	0	root_volume_1	2	1
11	y	1	0	root_slice_2	1	12
12	y	3	0	root_volume_2	2	2
13	y	1	0	root_slice_3	1	14
14	y	3	0	root_volume_3	2	3
15	y	1	0	root_slice_4	1	16
16	y	3	0	root_volume_4	2	4
.
.
.
1518	y	3	0	Meta_S2vdm_FS1	2	229
1527	y	3	0	Meta_S2_FS1	2	235

7. Create the first stripe by using this command syntax:

```
$ nas_volume -name <name> -create -Stripe
<stripe_size> <volume_set>,...
```

where:

<name> = name of new stripe pair

<stripe_size> = size of the stripe

<volume_set> = set of disks

For example, to create a stripe pair named s2_stripe1 and a depth of 262144 bytes (256 KB) by using disks d9, d14, d11, d16, d17, and d22, type:

```
$ nas_volume -name s2_stripe1 -create -Stripe 262144
d9,d14,d11,d16,d17,d22
```

Output:

```

id           = 135
name        = s2_stripe1
acl         = 0
in_use     = False
type       = stripe
stripe_size = 262144
volume_set  = d9,d14,d11,d16,d17,d22
disks      = d9,d14,d11,d16,d17,d22

```

Note: For best MPFS performance, in most cases, configure your file volumes by using a volume stripe size of 256 KB. Detailed performance-related information is available in the *EMC VNX MPFS over iSCSI Applied Best Practices Guide*.

8. Create the second stripe by using this command syntax:

```

$ nas_volume -name <name> -create -Stripe
<stripe_size> <volume_set>, ...

```

where:

<name> = name of new stripe pair

<stripe_size> = size of the stripe

<volume_set> = set of disks

For example, to create a stripe pair named `s2_stripe2` and a depth of 262144 bytes (256 KB) by using disks `d13`, `d10`, `d15`, `d12`, `d18`, and `d19`, type:

```

$ nas_volume -name s2_stripe2 -create -Stripe 262144
d13,d10,d15,d12,d18,d19

```

Output:

```

id           = 136
name        = s2_stripe2
acl         = 0
in_use     = False
type       = stripe
stripe_size = 262144
volume_set  = d13,d10,d15,d12,d18,d19
disks      = d13,d10,d15,d12,d18,d19

```

9. Create the metavolume by using this command syntax:

```
$ nas_volume -name <name> -create -Meta <volume_name>
```

where:

<name> = name of the new meta volume

<volume_name> = names of the volumes

For example, to create a meta volume s2_meta1 with volumes s2_stripe1 and s2_stripe2, type:

```
$ nas_volume -name s2_meta1 -create -Meta s2_stripe1,
s2_stripe2
```

Output:

```
id           = 137
name         = s2_meta1
acl          = 0
in_use      = False
type        = meta
volume_set  = s2_stripe1, s2_stripe2
disks       =
            d9, d14, d11, d16, d17, d22, d13, d10, d15, d12, d18, d19
```

10. Create the file system by using this command syntax:

```
$ nas_fs -name <name> -create <volume_name>
```

where:

<name> = name of the new file system

<volume_name> = name of the meta volume

For example, to create a file system s2fs1 with a meta volume s2_meta1, type:

```
$ nas_fs -name s2fs1 -create s2_meta1
```

Output:

```
id           = 33
name         = s2fs1
acl          = 0
in_use      = False
type        = uxf
worm         = compliance
worm_clock  = Thu Mar 6 16:26:09 EST 2008
worm Max Retention Date = Fri April 18 12:30:40 EST 2008
volume      = s2_meta1
pool        =
rw_servers  =
ro_servers  =
rw_vdms     =
```

```

ro_vdms      =
auto_ext    = no, virtual_provision=no
stor_devs   =
    APM00065101342-0012,APM00065101342-0015,APM00065101
    342-0016,APM00065101342-0019,APM00065101342-001A,AP
    M00065101342-001D,APM00065101342-0013,APM0006510134
    2-0014,APM00065101342-0017,APM00065101342-0018,APM0
    0065101342-001B,APM00065101342-001C
disks       =
    d9,d14,d11,d16,d17,d22,d13,d10,d15,d12,d18,d19

```

11. Create the mount point by using this command syntax:

```
$ server_mountpoint <movername> -create <pathname>
```

where:

<movername> = name of the Data Mover

<pathname> = path of the new mount point

For example, to create a mount point on Data Mover server_2 with a path of /s2fs1, type:

```
$ server_mountpoint server_2 -c /s2fs1
```

Output:

```
server_2 : done
```

12. Mount the file system by using this command syntax:

```
$ server_mount <movername> <fs_name> <mount_point>
```

where:

<movername> = name of the Data Mover

<fs_name> = name of the file system to mount

<mount_point> = name of the mount point

For example, to mount a file system on Data Mover server_2 with file system s2fs1 and mount point /s2fs1, type:

```
$ server_mount server_2 s2fs1 /s2fs1
```

Output:

```
server_2 : done
```

13. Export the file system by using this command syntax:

```
$ server_export <mover_name> -Protocol nfs -name
<name> -option <options> <pathname>
```

where:

<mover_name> = name of the Data Mover

<name> = name of the alias for the <pathname>

<options> = options to include

<pathname> = path of the mount point created

For example, to export a file system on Data Mover server_2 with a pathname alias of ufs1 and mount point path /ufs1, type:

```
$ server_export server_2 -P nfs -name ufs1 /ufs1
```

Output:

```
server_2 : done
```

Related documentation

The following documents provide more information on building the MPFS file system and are available on the EMC Online Support website:

- ◆ *VNX Command Reference Manual*
- ◆ *Configuring and Managing VNX Networking*
- ◆ *Managing VNX Volumes and File Systems Manually*
- ◆ *Using VNX Multi-Path File System*

Enabling MPFS for the VNX for file

Start MPFS on the VNX for file. Use this command syntax:

```
$ server_setup <movername> -Protocol mpfs -option  
<options>
```

where:

<movername> = name of the Data Mover

<options> = options to include

For example, to start MPFS on Data Mover server_2, type:

```
$ server_setup server_2 -Protocol mpfs -option start
```

Output:

```
server_2 : done
```

Note: Start MPFS on the same Data Mover on which the file system was exported by using CIFS.

Configuring the VNX for block by using CLI commands

This section presents an overview of configuring the VNX for block array ports mounted on the VNX for block in VNX VG2/VG8 configurations. Use site-specific parameters for these steps.

Use VNX CLI for block commands to configure the VNX for block array ports for a VNX VG2/VG8 configuration.

Best practices for VNX for block and VNX VG2/VG8 configurations

To simplify the configuration and management of the Windows server, EMC recommends that the discovery addresses (IP addresses) and enabled targets for each Windows server be configured so that all the iSCSI target ports on the system are equally balanced to achieve maximum performance and availability. Balancing the load across all ports enables speeds up to 4 x 10 Gb/s per storage processor. If one of the iSCSI target ports fails, the other three remain operational, so one-fourth of the Windows servers fail over to the native NFS or CIFS protocol, but three-fourths of the Windows servers continue operating at higher speeds attainable through iSCSI.

VNX for block discovery sessions reveal paths to all four iSCSI ports on each storage processor. The ports are described to the iSCSI initiators as individual targets. Each of these connections creates another session. The maximum number of initiator sessions or hosts per storage processor is dependent on the VNX for block configuration. To increase the number of achievable Windows servers for a VNX VG2/VG8 configuration, disable access on each Windows server to as many as three out of four iSCSI targets per storage processor. Ensure that the enabled iSCSI targets (VNX for block iSCSI ports) match the storage group definition.

For VNX VG2/VG8 configurations, Access Logix LUN masking by using iSCSI is used to present all VNX for file managed LUNs to the Windows servers. The non-VNX for file LUNs are protected from the iSCSI initiators. A separate storage group is created for MPFS initiators and all VNX for file LUNs that are not Control LUNs are added to this group. Enable at least one port from each SP for each Windows server in this group.

In a VNX VG2/VG8 environment, iSCSI initiator names are used in providing the path in the storage group for the Windows server to access the iSCSI targets. Unique, known iSCSI names are required by using Access Logix software.

Configuring the SAN and storage

This section describes how to configure the SAN switch along with specific configuration information for VNX for block and Symmetrix.

Installing the FC switch (FC configuration)

To set up the FC switch, complete these tasks:

1. Install the FC switch (if not already installed).
2. Connect cables from each HBA FC port to a switch port.
3. Verify the HBA connection to the switch by checking the LEDs for the switch port connected to the HBA port.

Note: Configure zoning as *single initiator*, meaning that each HBA port has its own zone in which it is the only HBA port.

Zoning the SAN switch (FC configuration)

This section presents an overview of configuring and zoning a FC switch:

1. Record all attached port WWNs.
2. Create a zone for each FC HBA port and its associated FC Target.

Related documentation

The documentation that ships with the FC switch provides additional information on installing or zoning.

Note: Configure the VNX for block so that each target is zoned to an SP A and SP B port. Configure the Symmetrix so that it is zoned to a single FC Director (FA).

Creating a security file on the VNX for file

A VNX for block does not accept a Secure CLI command unless the user who issues the command has a valid user account on the VNX for block. Configure a Navisphere 6.X security file to issue Secure CLI commands on the server. Secure CLI commands require the servers (or the password prompt) in each command line; they are not needed in the command line if a security file is created.

To create a security file:

1. Log in to the Control Station as NAS administrator.
2. Create a security file by using the following **naviseccli** command syntax:

```
$ /nas/sbin/naviseccli -h <hostname:IP address>
-AddUserSecurity -scope 0 -user nasadmin -password
nasadmin
```

where:

<hostname:IP address> = name of the VNX for file or IP address of the VNX for block

For example, type:

```
$ /nas/sbin/naviseccli -h 172.24.107.242
-AddUserSecurity -scope 0 -user nasadmin -password
nasadmin
```

Output:

This command produces no system response. When the command has finished executing, only the command line prompt is returned.

3. Verify that the security file was created correctly by using the following command syntax:

```
$ /nas/sbin/naviseccli -h <hostname:IP address>
getagent
```

where:

<hostname:IP address> = name of the VNX for file or IP address of the VNX for block

For example, type:

```
$ /nas/sbin/naviseccli -h 172.24.107.242 getagent
```

Output:

```
Agent Rev:          7.30.0 (4.93)
Name:              K10
Desc:
Node:              A-FNM00084100106
Physical Node:     K10
Signature:         2280155
Peer Signature:    2259565
Revision:          04.30.000.5.004
SCSI Id:           0
Model:             CX4-960
Model Type:        Rackmount
Prom Rev:          6.60.00
SP Memory:         16352
Serial No.         FNM00084100106
SP Identifier:     A
Cabinet:           SPE4
```

If the security file was not created correctly or is not found, an error message is displayed:

```
Security file not found. Already removed or check
-secfilepath option.
```

4. If an error message is displayed, repeat [step 2](#) and [step 3](#) to create the security file.

Configuring the VNX for block iSCSI port

This section describes how to set up the VNX for block in an iSCSI configuration:

Note: The IP addresses of all systems *<hostname:IP address>* are located in the */etc/hosts* file on the Control Station. If multiple systems are used, EMC recommends registering them in the */etc/hosts* file.

1. Configure iSCSI target hostname SP A and port IP address 0 on the system by using the following **naviseccli** command syntax:

```
$ /nas/sbin/naviseccli -h <hostname:IP address>
connection -setport -sp a -portid 0 -address <port IP
address> -subnetmask <subnet mask> -gateway <gateway
IP address>
```

where:

<hostname:IP address> = name of the VNX for file or IP address of the VNX for block.

<port IP address> = IP address of a named logical element mapped to a port on a Data Mover. Each interface assigns an IP address to the port.

<subnet mask> = 32-bit address mask used in IP to identify the bits of an IP address used for the subnet address.

<gateway IP address> = IP address of the machine through which network traffic is routed.

For example, type:

```
$ /nas/sbin/naviseccli -h 172.24.107.242 connection
-setport -sp a -portid 0 -address 172.241.107.1
-subnetmask 255.255.255.0 -gateway 172.241.107.2
```

Output:

It is recommended that you consult with your Network Manager to determine the correct settings before applying these changes. Changing the port properties may disrupt iSCSI traffic to all ports on this SP. Initiator configuration changes may be necessary to regain connections. Do you really want to perform this action (y/n)? **y**

```
SP: A
Port ID: 0
Port WWN: iqn.1992-04.com.emc:cx.apm00065101342.a0
iSCSI Alias: 2147.a0
IP Address: 172.24.107.242
Subnet Mask: 255.255.255.0
gateway Address: 172.241.107.2
Initiator Authentication: false
```

Note: If the iSCSI target is not configured (by replying with **n**), the command line prompt is returned.

2. Continue for SP A ports 1–3 and SP B ports 0–3 by using the following command syntax:

```
$ /nas/sbin/naviseccli -h <hostname:IP address>
connection -setport -sp a -portid 1 -address <port IP
address> -subnetmask <subnet mask> -gateway <gateway
IP address>
```

```
$ /nas/sbin/naviseccli -h <hostname:IP address>
connection -setport -sp a -portid 2 -address <port IP
address> -subnetmask <subnet mask> -gateway <gateway
IP address>
```

```
$ /nas/sbin/naviseccli -h <hostname:IP address>
connection -setport -sp a -portid 3 -address <port IP
address> -subnetmask <subnet mask> -gateway <gateway
IP address>
```

```
$ /nas/sbin/naviseccli -h <hostname:IP address>
connection -setport -sp b -portid 0 -address <port IP
address> -subnetmask <subnet mask> -gateway <gateway
IP address>
```

```
$ /nas/sbin/naviseccli -h <hostname:IP address>
connection -setport -sp b -portid 1 -address <port IP
address> -subnetmask <subnet mask> -gateway <gateway
IP address>
```

```
$ /nas/sbin/naviseccli -h <hostname:IP address>
connection -setport -sp b -portid 2 -address <port IP
address> -subnetmask <subnet mask> -gateway <gateway
IP address>
```

```
$ /nas/sbin/naviseccli -h <hostname:IP address>
connection -setport -sp b -portid 3 -address <port IP
address> -subnetmask <subnet mask> -gateway <gateway
IP address>
```

The outputs for SP A ports 1–3 and SP B ports 0–3 are the same as SP A port 0 with specific port information for each port.

Note: Depending on the system configuration, additional storage processors (SP C, SP D, and so on) each containing ports 0–3 can exist.

Access Logix configuration

Setting failovermode and the arraycommpath using Access Logix

This section describes how to set up an Access Logix configuration, create storage groups, add LUNs, and set failovermode and the arraycommpath.

The `naviseccli failovermode` command enables or disables the type of trespass needed for the failover software. This method of setting failovermode works for VNX for block with Access Logix only.

The `naviseccli arraycommpath` command enables or disables a communication path from the VNX for file to the VNX for block. This command is needed to configure a VNX for block when LUN 0 is not configured. This method of setting arraycommpath works for VNX for block with Access Logix only.



CAUTION

Changing the failovermode setting can force the VNX for block to reboot. Changing the failovermode to the wrong value makes the storage group inaccessible to any connected server.

Note: It is suggested that failovermode and arraycommpath are both set to 1 for MPFS. If EMC PowerPath is enabled, set failovermode to 1.

To set and verify failovermode and arraycommpath settings:

1. Set failovermode to 1 (VNX for file only) by using the following `naviseccli` command syntax:

```
$ /nas/sbin/naviseccli -h <hostname:IP address> -scope 0 -user nasadmin -password nasadmin failovermode 1
```

where:

`<hostname:IP address>` = name of the VNX for file or IP address of the VNX for block

For example, type:

```
$ /nas/sbin/naviseccli -h 172.24.107.242 -scope 0 -user nasadmin -password nasadmin failovermode 1
```

Output:

```
WARNING: Previous Failovermode setting will be lost!
DO YOU WISH TO CONTINUE (y/n)? y
```

Note: Setting or not setting failovermode produces no system response. When the command has finished executing, only the command line prompt is returned.

2. Verify the failovermode setting (VNX for file only) by using the following **naviseccli** command syntax:

```
$ /nas/sbin/naviseccli -h <hostname:IP address> -scope
0 -user nasadmin -password nasadmin failovermode
```

For example, type:

```
$ /nas/sbin/naviseccli -h 172.24.107.242 -scope 0 -user
nasadmin -password nasadmin failovermode
```

Output:

```
Current failovermode setting is: 1
```

3. Set arraycommpath to 1 (VNX for file only) by using the following **naviseccli** command syntax:

```
$ /nas/sbin/naviseccli -h <hostname:IP address> -scope
0 -user nasadmin -password nasadmin arraycommpath 1
```

where:

<hostname:IP address> = name of the VNX for file or IP address of the VNX for block

For example, type:

```
$ /nas/sbin/naviseccli -h 172.24.107.242 -scope 0 -user
nasadmin -password nasadmin arraycommpath 1
```

Output:

```
WARNING: Previous arraycommpath setting will be lost!
DO YOU WISH TO CONTINUE (y/n)? y
```

Note: Setting or not setting arraycommpath produces no system response. When the command has finished executing, only the command line prompt is returned.

4. Verify the arraycommpath setting (VNX for file only) by using the following **naviseccli** command syntax:

```
$ /nas/sbin/naviseccli -h <hostname:IP address> -scope
0 -user nasadmin -password nasadmin arraycommpath
```

For example, type:

```
$ /nas/sbin/naviseccli -h 172.24.107.242 -scope 0 -user
nasadmin -password nasadmin arraycommpath
```

Output:

```
Current arraycommpath setting is: 1
```

To discover the current settings of failovermode or the arraycommpath, also use the **port -list -failovermode** or **port -list -arraycommpath** commands.

Note: The outputs of these commands provide more detail than just the failovermode and arraycommpath settings and are multiple pages in length.

Creating storage groups and adding LUNs

This section describes how to create storage groups, add LUNs to the storage groups, and configure the storage groups.

The IP addresses of all systems *<hostname:IP address>* are located in the */etc/hosts* file on the Control Station. If multiple arrays are used, EMC recommends registering them in the */etc/hosts* file:

Note: Specify the hostname as the name of the VNX for file, for example *Server_2*.

1. Create the storage group by using the following **navicli** command syntax:

```
$ /nas/sbin/navicli -h <hostname:IP address>
storagegroup -create -gname MPFS_Clients
```

where:

<hostname:IP address> = name or IP address of the VNX for file

For example, type:

```
$ /nas/sbin/navicli -h 172.24.107.242
storagegroup -create -gname MPFS_Clients
```

Output:

This command produces no system response. When the command has finished executing, only the command line prompt is returned.

2. Add LUNs to the storage group by using the following **navicli** command syntax:

```
$ /nas/sbin/navicli -h <hostname:IP address>  
storagegroup -addhlu -gname MPFS_Clients -hlu 0 -alu 16
```

where:

<hostname:IP address> = name or IP address of the VNX for file

For example, type:

```
$ /nas/sbin/navicli -h 172.24.107.242 storagegroup  
-addhlu -gname MPFS_Clients -hlu 0 -alu 16
```

Output:

This command produces no system response. When the command has finished executing, only the command line prompt is returned.

3. Continue adding LUNs to the rest of the storage group:

```
$ /nas/sbin/navicli -h <hostname:IP address>  
storagegroup -addhlu -gname MPFS_Clients -hlu 0 -alu 17
```

where:

<hostname:IP address> = name or IP address of the VNX for file

For example, type:

```
$ /nas/sbin/navicli -h 172.24.107.242 storagegroup  
-addhlu -gname MPFS_Clients -hlu 0 -alu 17
```

Output:

This command produces no system response. When the command has finished executing, only the command line prompt is returned.

Configuring and accessing storage

This section describes how to install the FC driver, add hosts to storage groups, install and configure the Microsoft iSCSI initiator, and add initiators to the storage group.

Note: EMC recommends to view the iSCSI disks by using the Windows Device Manager. This method of viewing iSCSI disks does not expose the VNX for file disks to Windows operations.

The arraycommpath and failovermode settings are used to see both active and passive paths concurrently. For a LUN failover, LUNs are presented from active to passive path or passive to active path. Use the arraycommpath and failovermode settings as described in [Table 2 on page 64](#).

Table 2 Arraycommpath and failovermode settings

		Default	VNX for file ports	MPFS clients
Access Logix units	arraycommpath	0	0	1
	failovermode	0	0	1
Storage systems	arraycommpath	0	n/a	n/a
	failovermode	1	n/a	n/a

Any MPFS server that is connected and logged in to a storage group should have the arraycommpath and failovermode set to 1. For any VNX for file port connected to a storage group, these settings are 0. These settings are on an individual server/port basis and override the global settings on the system default of 0.

When using the VNX for block in a VNX VG2/VG8 iSCSI configuration, the iSCSI initiator name, or IQN, is used to define the server, not a WWN.

Installing the FC driver (FC configuration)

Install the FC driver on the Windows server. The latest driver and qualification information is available on the FC manufacturer's website, the EMC E-Lab Interoperability Navigator, or the documentation provided with the FC driver.

Adding hosts to the storage group (FC configuration)

Use the following steps to view hosts in the storage group and add hosts to the storage group for SP A and SP B:

1. List the hosts in the storage group by using the following **navicli** command syntax:

```
$ /nas/sbin/navicli -h <hostname:IP address> port
-list |grep "HBA UID:"
```

where:

<hostname:IP address> = name or IP address of the VNX for file

For example, type:

```
$ /nas/sbin/navicli -h 172.24.107.242 port -list |grep
"HBA UID:"
```

Output:

```
HBA UID: 20:01:00:1B:32:20:D1:3A:21:01:00:1B:32:20:D1:3A
HBA UID: 20:00:00:1B:32:00:D1:3A:21:00:00:1B:32:00:D1:3A
HBA UID: 20:01:00:1B:32:20:B5:35:21:01:00:1B:32:20:B5:35
HBA UID: 20:00:00:1B:32:00:B5:35:21:00:00:1B:32:00:B5:35
```

2. Add hosts to the storage group by using the following **navicli** command syntax:

```
$ /nas/sbin/navicli -h <hostname:IP address>
storagegroup -setpath -gname <gname> -hbauid <hbauid>
-sp <sp> -spport <spport> -failovermode <failovermode>
-arraycommpath <arraycommpath>
```

where:

<hostname:IP address>	= name or IP address of the VNX for file
<gname>	= storage group name
<hbauid>	= WWN of proxy initiator
<sp>	= storage processor
<spport>	= port on SP
<failovermode>	= enables or disables the type of trespass needed for failover software (1 = enable, 0 = disable)
<arraycommpath>	= creates or removes a communication path between the server and the VNX for block (1 = enable, 0 = disable)

Examples of adding hosts to storage groups are shown in [step 3](#) and [step 4](#).

3. Add hosts to storage group A:

```
$ /nas/sbin/navicli -h 172.24.107.242
storagegroup -setpath -gname MPFS_Clients -hbauid
20:0a:00:0d:ec:01:53:82:20:09:00:0d:ec:01:53:82 -sp a
-spport 0 -failovermode 1 -arraycomppath 1
```

Note: The IP addresses of all systems <hostname:IP address> are located in the /etc/hosts file on the Control Station. If multiple arrays are used, EMC recommends registering them in the /etc/hosts file.

Output:

The recommended configuration is to have all HBAs on one host mapped to the same storage group.

```
Set Path to storage group MPFS_Clients (y/n)? y
```

```
WARNING: Changing configuration options may cause the
array to stop functioning
correctly. Failover-related Initiator settings for a
single host MUST BE CONSISTENT
for all paths from the host to the storage system.
Please verify after reconnect.
Do you wish to continue (y/n)? y
```

Note: This command produces no system response. When the command has finished executing, only the command line prompt is returned.

4. Add hosts to storage group B:

```
$ /nas/sbin/navicli -h 172.24.107.242
storagegroup -setpath -gname MPFS_Clients -hbauid
20:0a:00:0d:ec:01:53:82:20:09:00:0d:ec:01:53:82 -sp b
-spport 0 -failovermode 1 -arraycomppath 1
```

Output:

The recommended configuration is to have all HBAs on one host mapped to the same storage group.

```
Set Path to storage group MPFS_Clients (y/n)? y
```

```

WARNING: Changing configuration options may cause the
array to stop functioning
correctly. Failover-related Initiator settings for a
single host MUST BE CONSISTENT
for all paths from the host to the storage system.
Please verify after reconnect.
Do you wish to continue (y/n)? y

```

Note: This command produces no system response. When the command has finished executing, only the command line prompt is returned.

Installing and configuring the Microsoft iSCSI initiator (iSCSI configuration)

The *Microsoft iSCSI Software Initiator User Guide*, which is installed with the Microsoft iSCSI initiator, provides detailed configuration information. Install and configure the Microsoft iSCSI initiator by performing the following tasks:

- ◆ Install the Microsoft iSCSI initiator.
- ◆ Register the initiator name with the Windows Registry.
- ◆ Configure the CHAP secret for reverse authentication (optional).
- ◆ Configure iSCSI discovery on the initiator.
- ◆ Log in to the iSCSI target.

Installing the Microsoft iSCSI initiator (iSCSI configuration)

To connect to the iSCSI targets on a Windows server, the server requires an iSCSI initiator. VNX supports the Microsoft iSCSI Software initiator. To install the iSCSI initiator:

1. Download the latest Microsoft iSCSI initiator software from the Microsoft website.
2. Run the initiator executable. The **Welcome to Microsoft iSCSI Initiator Setup Wizard** dialog box appears.
3. Click **Next**.
4. In the **Select Installation Folder** dialog box, select **Everyone**, and then do one of the following:
 - To accept the default installation folder, click **Next**.
 - To use another folder for installation, click **Browse** to locate the folder, and then click **Next**.
5. In the **Confirm Installation** dialog box, click **Next**.
6. In the **License Agreement** dialog box, select **I Agree**.
7. Click **Next**.

8. In the **Microsoft iSCSI Installation Program** dialog box, click **OK** to install the complete iSCSI initiator.
9. In the **End User License Agreement** dialog box, click **Agree** to accept the terms of the license agreement.
10. To complete the installation:
 - a. Click **OK** to close the installation confirmation message.
 - b. At the **Microsoft iSCSI Initiator Information** dialog box, click **Next**.
 - c. At the **Installation Complete** dialog box, click **Close**.

Registering the initiator name (iSCSI configuration)

Complete this procedure for each iSCSI initiator that needs to connect to a VNX. If the initiator's iSCSI Qualified Name (IQN) is not written to the Windows Registry, iSCSI-aware applications cannot find the initiator's IQN:

1. Open the Microsoft iSCSI initiator.
2. Click **General** to verify the version shown.
3. In the **Initiator Node Name Change** window, leave the displayed initiator name unchanged. Click **OK**.

This procedure converts uppercase characters in the IQN to lowercase. Only lowercase IQNs are supported.

Configuring the CHAP secret for reverse authentication (optional)

To use reverse authentication for iSCSI sessions, configure the iSCSI initiator with a CHAP secret.

Before configuring the CHAP secret

CHAP provides a method for iSCSI initiators and targets to authenticate one another by exchanging a shared secret or password. Because the secret is shared between the initiator and target, configure the same secret on both. The initiator and target systems maintain databases of CHAP entries:

- ◆ With regular CHAP, the target challenges the initiator for a CHAP secret.
- ◆ With reverse authentication, the initiator challenges the target for a CHAP secret.

The Microsoft iSCSI initiator supports only one secret per initiator.

By default, targets on the server do not require CHAP authentication. Depending on the organization's needs, a CHAP authentication is required. CHAP authentication is enabled at two different points of the iSCSI login:

- ◆ **Discovery authentication:** CHAP authentication is required before the initiator can contact the server to establish a discovery session during which the initiator tries to discover available targets. If discovery authentication is enabled on the iSCSI target, define the CHAP secret when configuring iSCSI discovery.

[“Configuring iSCSI discovery on the initiator \(iSCSI configuration\)” on page 69](#) provides the procedure.

- ◆ **Session authentication:** CHAP authentication is required before the initiator can establish a regular iSCSI session with the target. If session authentication is enabled on the server, the initiator must provide the CHAP secret when logging in to the iSCSI target.

[“Logging in to the iSCSI target \(iSCSI configuration\)” on page 71](#) provides the procedure.

Configuring the CHAP secret

To configure the CHAP secret:

1. In the Microsoft iSCSI initiator, click **General**.
2. Click **Secret**.
3. In the **CHAP Secret Setup** window:
 - Type a CHAP secret of 12 to 16 characters
or
 - To clear the secret, click **Reset**.
4. Click **OK**.

Configuring iSCSI discovery on the initiator (iSCSI configuration)

Before the initiator can establish a session with a target, it must discover where targets are located and the names of the targets available. To obtain this information, the initiator uses the iSCSI discovery process. To configure iSCSI discovery:

1. In the Microsoft iSCSI initiator, click **Discovery**.
2. In the **Target Portals** section, click **Add**.

3. In the **Add Target Portal** dialog box, type the IP address of the target's network portal. If the target uses a port other than 3260 (the default), type that port number in the **Port** field.

To ensure that the network is available, use the ping command to target the IP address before configuring it in the iSCSI initiator. If the server is unavailable, or if you type an invalid IP address, you receive the error Connection Failed.

4. Do you want to use forward CHAP authentication (where the target challenges the initiator)?
 - If YES, continue with [step 5](#).
 - If NO, go to [step 7](#).

CHAP authentication is optional only if the target does not require it. If the target requires authentication and you do not configure a forward CHAP secret on the initiator, the initiator cannot log in to the target. If the target does not require CHAP authentication, but the initiator offers it, the target complies with the initiator's request. *Configuring iSCSI Targets on VNX* provides more information about enabling CHAP authentication.

5. In the **Add Target Portal** screen, click **Advanced**.
6. In the **Advanced Settings** dialog box, do the following to type the CHAP secret:
 - a. Select **CHAP logon information**.
 - b. In the **Target secret** field, type the secret configured for the iSCSI target. Microsoft supports CHAP secrets of 12 to 16 characters only.
 - c. If you also want the initiator to authenticate the target for iSCSI discovery, select **Perform mutual authentication**. The optional task "[Configuring the CHAP secret](#)" on [page 69](#) provides instructions for configuring reverse authentication on the initiator.
 - d. Click **OK**.
7. Click **OK**.

Logging in to the iSCSI target (iSCSI configuration)

After configuring the initiator with the target's network portal, a list of available targets appears on the initiator's Targets tab. To access the target's LUNs, the initiator must log in to the target:

1. In the Microsoft iSCSI initiator, click the **Targets** tab.

The list of available targets appears.

2. Select the target, and then click **Log On**.

The **Log On to Target** dialog box appears.

3. (Optional) If the target requires CHAP authentication, click **Advanced**.

In the **Advanced Settings** dialog box, do the following to type the CHAP secret:

- a. Select **CHAP logon information**.
 - b. In the **Target secret** field, type the secret configured for the iSCSI target. Microsoft supports CHAP secrets of 12 to 16 characters.
 - c. If you want the initiator to authenticate the target for iSCSI discovery, select **Perform mutual authentication**. The optional task "[Configuring the CHAP secret](#)" on page 69 provides instructions for configuring reverse authentication on the initiator.
 - d. Click **OK**.
4. Select **Automatically restore this connection when the system boots** in the **Log On to Target** dialog box as shown in [Figure 8 on page 72](#), and then click **OK**.




Figure 8 Log On to Target dialog box

The initiator connects to the target.

- To get information about the new session, right-click the target, and then select **Details**. The **Target Properties** dialog box, as shown in [Figure 9 on page 73](#), appears.

- Click the **Sessions** tab in the **Target Properties** dialog box as shown in [Figure 9](#) on page 73.




Figure 9 Target Properties dialog box

- To end one or more sessions, select the checkbox for each, and then click **Log off**.

Adding initiators to the storage group (FC configuration)

In an FC configuration, the storage group should contain the HBA UID of the Windows servers.

The IP addresses of all systems *<hostname:IP address>* are located in the `/etc/hosts` file on the Control Station. If multiple arrays are used, EMC recommends registering them in the `/etc/hosts` file.

Use the following steps to add initiators to the storage group for SP A and SP B:

1. Use the following **navicli** command to list hosts in the storage group:

```
$ /nas/sbin/navicli -h <hostname:IP address> port
-list |grep "HBA UID:"
```

For example, type:

```
$ /nas/sbin/navicli -h 172.24.107.242 port
-list |grep "HBA UID:"
```

Output:

```
HBA UID: 20:01:00:1B:32:20:D1:3A:21:01:00:1B:32:20:D1:3A
HBA UID: 20:00:00:1B:32:00:D1:3A:21:00:00:1B:32:00:D1:3A
HBA UID: 20:01:00:1B:32:20:B5:35:21:01:00:1B:32:20:B5:35
HBA UID: 20:00:00:1B:32:00:B5:35:21:00:00:1B:32:00:B5:35
```

2. Use the following **navicli** command to add initiators to the storage group by using this command syntax:

```
$ /nas/sbin/navicli -h <hostname:IP address>
storagegroup -setpath -gname <gname> -hbauid <hbauid>
-sp <sp> -spport <spport> -failovermode <failovermode>
-arraycommpath <arraycommpath>
```

where:

- | | |
|------------------------------|--|
| <i><gname></i> | = storage group name |
| <i><hbauid></i> | = HBA UID of Windows servers |
| <i><sp></i> | = storage processor |
| <i><spport></i> | = port on SP |
| <i><failovermode></i> | = enables or disables the type of trespass needed for failover software (1 = enable, 0 = disable) |
| <i><arraycommpath></i> | = creates or removes a communication path between the server and the VNX for block (1 = enable, 0 = disable) |

Note: Perform this command for each SP.

Examples of adding initiators to storage groups are shown in [step 3](#) and [step 4](#).

3. Add initiators to the storage group for SP A:

```
$ /nas/sbin/navicli -h 172.24.107.242
storagegroup -setpath -gname MPFS_Clients -hbauid
20:01:00:1B:32:20:D1:3A:21:01:00:1B:32:20:D1:3A -sp a
-spport 0 -failovermode 1 -arraycommpath 1
```

Output:

The recommended configuration is to have all HBAs on one host mapped to the same storage group.

```
Set Path to storage group MPFS_Clients (y/n)? y
```

```
WARNING: Changing configuration options may cause the
array to stop functioning
correctly. Failover-related Initiator settings for a
single host MUST BE CONSISTENT
for all paths from the host to the storage system.
Please verify after reconnect.
Do you wish to continue (y/n)? y
```

Note: This command produces no system response. When the command has finished executing, only the command line prompt is returned.

4. Add initiators to the storage group for SP B:

```
$ /nas/sbin/navicli -h 172.24.107.242
storagegroup -setpath -gname MPFS_Clients -hbauid
20:01:00:1B:32:20:D1:3A:21:01:00:1B:32:20:D1:3A -sp b
-spport 0 -failovermode 1 -arraycommpath 1
```

Output:

The recommended configuration is to have all HBAs on one host mapped to the same storage group.

```
Set Path to storage group MPFS_Clients (y/n)? y
```

```
WARNING: Changing configuration options may cause the
array to stop functioning
correctly. Failover-related Initiator settings for a
single host MUST BE CONSISTENT
for all paths from the host to the storage system.
Please verify after reconnect.
```

Do you wish to continue (y/n)? **y**

Note: This command produces no system response. When the command has finished executing, only the command line prompt is returned.

Adding initiators to the storage group (iSCSI configuration)

When using the VNX for block in a VNX VG2/VG8 iSCSI configuration, the iSCSI initiator name, or IQN, is used to define the host, not a WWN.

The IP addresses of all systems *<hostname:IP address>* are located in the `/etc/hosts` file on the Control Station. If multiple arrays are used, EMC recommends registering them in the `/etc/hosts` file.

Use the following steps to add initiators to the storage group for SP A and SP B:

1. Find the IQN used to define the host by using the following **navicli** command syntax:

```
$ /nas/sbin/navicli -h <hostname:IP address> port
-list |grep "HBA UID:" |grep iqn
```

For example, type:

```
$ /nas/sbin/navicli -h 172.24.107.242 port
-list |grep "HBA UID:" |grep iqn
```

Output:

```
InitiatorName=iqn.1991-05.com.microsoft:test99.mpfs1.com
```

2. Use the following **navicli** command to add initiators to the storage group by using this command syntax:

```
$ /nas/sbin/navicli -h <hostname:IP address>
storagegroup -setpath -gname <gname> -hbauid <hbauid>
-sp <sp> -spport <spport> -failovermode <failovermode>
-arraycommpath <arraycommpath>
```

where:

<code><gname></code>	= storage group name
<code><hbauid></code>	= iSCSI initiator name
<code><sp></code>	= storage processor
<code><spport></code>	= port on SP
<code><failovermode></code>	= enables or disables the type of trespass needed for failover software (1 = enable, 0 = disable)
<code><arraycommpath></code>	= creates or removes a communication path between the server and the VNX for block (1 = enable, 0 = disable)

Examples of adding initiators to storage groups are shown in [step 3](#) and [step 4](#).

Note: Perform this command for each iSCSI proxy-initiator.

3. Add initiators to the storage group for SP A:

```
$ /nas/sbin/navicli -h 172.24.107.242
storagegroup -setpath -gname MPFS_Clients -hbauid
iqn.1991-05.com.microsoft:test99.mpfs1.com -sp a
-spport 0 -failovermode 1 -arraycommpath 1
```

Output:

The recommended configuration is to have all HBAs on one host mapped to the same storage group.

```
Set Path to storage group MPFS_Clients (y/n)? y
```

```
WARNING: Changing configuration options may cause the
array to stop functioning
correctly. Failover-related Initiator settings for a
single host MUST BE CONSISTENT
for all paths from the host to the storage system.
Please verify after reconnect.
Do you wish to continue (y/n)? y
```

Note: This command produces no system response. When the command has finished executing, only the command line prompt is returned.

4. Add initiators to the storage group for SP B:

```
$ /nas/sbin/navicli -h 172.24.107.242
storagegroup -setpath -gname MPFS_Clients -hbauid
iqn.1991-05.com.microsoft:test99.mpfs1.com -sp b
-spport 0 -failovermode 1 -arraycomppath 1
```

Output:

The recommended configuration is to have all HBAs on one host mapped to the same storage group.

```
Set Path to storage group MPFS_Clients (y/n)? y
```

```
WARNING: Changing configuration options may cause the
array to stop functioning
correctly. Failover-related Initiator settings for a
single host MUST BE CONSISTENT
for all paths from the host to the storage system.
Please verify after reconnect.
Do you wish to continue (y/n)? y
```

Note: This command produces no system response. When the command has finished executing, only the command line prompt is returned.

Mapping an MPFS share to a network drive

Follow these steps to map an MPFS share to a network drive:

1. Click **Start > Run** to open the **Run** window. Type the VNX for file (Data Mover interface) containing the CIFS shares for mapping preceded by `\\`. [Figure 10 on page 79](#) shows an example of a VNX for file named `wrs2`.



Figure 10 Opening a VNX for file to view possible shares

2. Click **OK** and a **List of Available Shares** window appears as shown in [Figure 11 on page 80](#).




Figure 11 List of Available Shares

3. Right-click the share for mapping and select **MPFS Volume Properties** from the pop-up window as shown in [Figure 12 on page 80](#). The **MPFS Properties** window is displayed in [Figure 13 on page 81](#).




Figure 12 Selecting the MPFS Volume Properties

4. In the **MPFS Properties** window, select **Enable MPFS** to enable MPFS on a share.




Figure 13 MPFS Properties window

5. Optionally, select **Persistent Connection** to allow the driver to enable MPFS automatically when restarting the system and save the newly specified property values across restarts. If **Persistent Connection** is not selected, all property values revert to their default values when the server is restarted.
6. Optionally, select any of the share properties below:
 - Flush policy — By default, the Windows server writes metadata updates asynchronously. Select **Strict Flush Policy** to cause the Windows server to flush after every write. Clear **Strict Flush Policy** to cause asynchronous flushing. Asynchronous flushing (the default) provides superior performance and is preferable in most situations. Nevertheless, in certain situations use strict (synchronous) flushing. For example, select **Strict Flush Policy** when two collaborating applications, on two computers, require a high degree of serialization.

When multiple servers access the same MPFS file, their caches can get out of sync.

Use the Strict Flush Policy (synchronous flushing) if problems occur with multi-Windows server sharing.

- **Advanced** — Click the **Advanced** button in [Figure 13 on page 81](#) to open the **Advanced Option Dialog** window as shown in [Figure 14 on page 82](#).




Figure 14 Advanced Option Dialog window

Select any of the options described below:

Read ahead — The prefetching of data that an application requests so it is read from the disk drive into its read cache before the data is needed. Read ahead improves the performance of a read operation since the read speed increases if the file is read again. Read ahead is enabled by default.

Write behind — Stores data in a buffer that is written to the disk drive and sends the data when the system is idle or after a certain amount of time has elapsed. This operation optimizes how the disk drive saves data from memory and increases the write speed. When write behind is enabled and the Windows server loses the SAN connection, the data no longer passes through to CIFS. Write behind now issues a new FMP request to synchronize data back to the Data Mover. Write behind is enabled by default.

Section size — The size of the data prefetched for a read ahead or stored for a write behind operation. The default section size is 1 MB.

Async Close Policy — Closes the file before sending data and waiting for an acknowledgement before sending more data. Async Close Policy is disabled by default.

Minimum Read IO Size — The minimum size of the amount of data (in KB) read from a file. The default minimum read I/O size is 8 KB.

Global Option — Select **Global Option** in the **Advanced Option** window as shown in [Figure 14 on page 82](#) to open the **Global Options** window as shown in [Figure 15 on page 83](#).



Figure 15 Global Options window

Select the option described below:

Max Memory Driver Uses — Select the maximum amount of memory that the driver uses (in MB) from the list box.

7. Click **OK**.

Installing, Upgrading, or Uninstalling VNX MPFS Software

This chapter describes how to install, upgrade, uninstall, and disable an active session of the EMC VNX MPFS software.

This chapter covers the following topics:

- ◆ [Installing the MPFS software](#) 86
- ◆ [Upgrading the MPFS software](#)..... 92
- ◆ [Uninstalling the MPFS software](#) 94
- ◆ [Disabling an active MPFS software session](#) 97

Installing the MPFS software

This section describes the requirements necessary before installing the MPFS software and how to install the MPFS software from EMC Online Support or the EMC MPFS for Windows Client Software CD.

Before installing

Before installing the MPFS software, read the prerequisites for the Windows server and VNX for block, listed in this section:

- ❑ Verify that the Windows server on which you plan to install the MPFS software meets the MPFS configuration requirements specified in the *EMC VNX MPFS for Windows Clients Release Notes*.
- ❑ Ensure that the Windows server has a network connection to the Data Mover on which the MPFS software resides and that you can contact the Data Mover.
- ❑ Ensure that the Windows server meets the overall system and other configuration requirements specified in the E-Lab Interoperability Navigator.



CAUTION

Install the Windows server before installing the Microsoft iSCSI initiator. This protects disks from being inadvertently overwritten.



CAUTION

Cancel the Initialize and Convert the Disk Wizard when starting the Disk Manager Microsoft Management Console.

Installing the MPFS software

To install the MPFS software:

1. Create a directory on the local machine for the MPFS software files.
2. Locate the Windows executable file on the EMC Online Support website at <http://Support.EMC.com> or on the EMC MPFS for Windows Client Software CD.
3. Download the executable file from EMC Online Support or from the EMC MPFS for Windows Client Software CD to the directory created in step 1.
4. Double-click the appropriate executable file as listed in [Table 3 on page 87](#).

Table 3 MPFS executable files with Windows OS versions

OS	Service pack	Executable file ^a
Windows Server 2003 32 bit	SP2	EMCmpfs.win.<version>.win32.Win2003.exe
Windows Server 2003 64 bit	SP2	EMCmpfs.win.<version>.x64.Win2003.exe
Windows Server 2003 R2 32 bit	SP2	EMCmpfs.win.<version>.win32.Win2003.exe
Windows Server 2003 R2 64 bit	SP2	EMCmpfs.win.<version>.x64.Win2003.exe
Windows XP Professional 32 bit	SP3	EMCmpfs.win.<version>.win32.WinXP.exe
Windows XP Professional 64 bit	SP2	EMCmpfs.win.<version>.x64.WinXP.exe
Windows Vista Enterprise 32 bit	SP2	EMCmpfs.win.<version>.win32.Vista.exe
Windows Vista Enterprise 64 bit	SP2	EMCmpfs.win.<version>.x64.Vista.exe
Windows Server 2008 32 bit	SP2	EMCmpfs.win.<version>.win32.Win2008.exe
Windows Server 2008 64 bit	SP2	EMCmpfs.win.<version>.x64.Win2008.exe
Windows Server 2008 R2 64 bit	N/A	EMCmpfs.win.<version>.x64.Win2008R2.exe
Windows 7 Enterprise/Professional 32 bit	N/A	EMCmpfs.win.<version>.win32.Win7.exe
Windows 7 Enterprise/Professional 64 bit	N/A	EMCmpfs.win.<version>.x64.Win7.exe

a. Where <version> is the MPFS software version number.

5. The **Information** window appears before the installation begins as shown in [Figure 16 on page 88](#). Verify that UDP ports 625–635 are not blocked by a firewall and click **OK** to start the installation.



Figure 16 Information window before Install

6. The **EMC MPFS Installer** dialog box appears as shown in [Figure 17 on page 88](#).




Figure 17 EMC MPFS Installer dialog box

7. In the **EMC MPFS Installer** dialog box, type the install directory, if different from the default directory (C:\Program Files\EMC). Click **Install**.

- When the installation is complete, the message shown in [Figure 18 on page 89](#) appears.



Figure 18 EMC MPFS is Installed dialog box

- Click **Yes** to restart the computer.

Note: Note: Restart the Windows server after the MPFS software is installed.

Operating MPFS through a firewall

For proper MPFS operation, the Windows server and VNX for file (a Data Mover) must contact each other on their File Mapping Protocol (FMP) ports.

For proper MPFS operation, the Windows server must contact the VNX for file (a Data Mover) on its File Mapping Protocol (FMP) port. The VNX for file must also contact the Windows server on its FMP port.

If a firewall resides between the Windows server and the VNX for file, the firewall must allow access to the ports listed in [Table 4 on page 90](#) for the Windows server.

Table 4 Windows server firewall ports

Windows server	Windows server port	VNX for file port
Windows O/S Windows Server 2003 x32 Windows Server 2003 x64 Windows Server 2003 R2 x32 Windows Server 2003 R2 x64 Windows XP Professional x32 Windows XP Professional x64 Windows Vista Enterprise x32 Windows Vista Enterprise x64 Windows Server 2008 x32 Windows Server 2008 x64 Windows Server 2008 R2 x64 Windows 7 Enterprise/Professional x32 Windows 7 Enterprise/Professional x64	6907	4656, 2079, 1234, 111, 625–635

IPv6 supported Microsoft Windows platforms

Internet Protocol version 6 (IPv6) has a new suite of standard protocols for the network layer of the Internet replacing IPv4. IPv6 is built in to later versions of Microsoft Windows®. The supported Microsoft Windows platforms with IPv6 are:

- ◆ Windows Server® 2008 x32
- ◆ Windows Server® 2008 x64
- ◆ Windows Server® 2008 R2 x64
- ◆ Windows® 7 Enterprise/Professional x32
- ◆ Windows® 7 Enterprise/Professional x64

Upgrading the MPFS software

When upgrading the MPFS software, EMC requires uninstalling earlier versions of the MPFS software before installing a later version of MPFS software.

Attempting to install over existing MPFS software

If you attempt to install MPFS software on a system with existing MPFS software installed, a message window appears as shown in [Figure 19 on page 92](#).



Figure 19 Previous installation detected window

The steps to install over existing MPFS software are as follows:

1. Click **OK**. The **EMC MPFS Installer** window remains open with the **Exit** button remaining active as shown in [Figure 20 on page 92](#).
2. Click **Exit**.




Figure 20 EMC MPFS Installer window with Exit button active

3. Remove the existing MPFS software as described in [“Uninstalling the MPFS software” on page 94](#). Then, restart the Windows server before attempting to install new MPFS software.

Installing a later version of MPFS software

To install a later version of MPFS software:

1. Uninstall the previous version of MPFS software as described in [“Uninstalling the MPFS software” on page 94](#).
2. Install the new version of MPFS software as described in [“Installing the MPFS software” on page 86](#).

Each of these steps includes an MPFS software restart as part of the procedure.



IMPORTANT

Restarting is very important, so perform this step when instructed to do so.

Note: If upgrading from FC to iSCSI, complete the environmental changes prior to installing the new MPFS software.

Uninstalling the MPFS software

Use the following procedure to uninstall the MPFS software from a Windows server:

1. Click **Start** > **Settings** > **Control Panel** as shown in [Figure 21 on page 94](#).




Figure 21 Opening the Control Panel

- When the **Control Panel** opens, as shown in [Figure 22 on page 95](#), double-click **Add/Remove Programs**.



Figure 22 Control Panel window

- When the **Add or Remove Programs** window opens, locate **EMC MPFS 6.0.x.x** and click **Change/Remove** as shown in [Figure 23 on page 95](#).

Note: **Note:** If you are uninstalling an early version of MPFS, the name **EMC HighRoad** can appear in the **Add/Remove Programs** list instead of **EMC MPFS**. If this happens, select **HighRoad** and then click **Add/Remove** or **Change/Remove**.




Figure 23 Add or Remove Programs window

4. After the MPFS software program is uninstalled, the **Uninstall completed** message, as shown in [Figure 24 on page 96](#), is displayed.



Figure 24 EMC MPFS Uninstall completed message window

5. Click **Yes** to restart the Windows server.

Disabling an active MPFS software session

To disable an active MPFS software session:

1. From **My Computer**, double-click the **VNX for file**, and then right-click the **share name**.
2. Select **MPFS Volume Properties** from the menu.

The **MPFS Properties** window opens as shown in [Figure 25 on page 97](#).




Figure 25 MPFS Properties window

3. Clear **Enable MPFS** and then click **OK**.

This chapter describes the **mpfs** and **mpfsctl** commands, which are used on a Windows server to establish and terminate MPFS sessions and to set and display MPFS share properties. Topics covered are:

◆ Overview of the command line interface	100
◆ mpfscopy	101
◆ mpfsctl -?	103
◆ mpfsctl config	104
◆ mpfsctl disable	104
◆ mpfsctl diskspeed	105
◆ mpfsctl enable	105
◆ mpfsctl inq	106
◆ mpfsctl list	107
◆ mpfsctl reset	107
◆ mpfsctl stats	108
◆ mpfsctl version	109
◆ mpfsinfo	110
◆ mpfsinq	112

Note: As an alternative to the command line interface, establish and terminate sessions and set MPFS share properties by using the share's Properties dialog box as described in [Chapter 2, "EMC VNX MPFS Environment Configuration."](#)

Overview of the command line interface

The **mpfsctl** and **mpfs** commands are shown in [Table 5 on page 100](#). Type these commands at the Windows command prompt, or include them in a command file to execute at boot or login time.

Table 5 Windows command line interface commands

Command	Description	Page
mpfscopy	Duplicates directories and files.	101
mpfsctl -?	Shows information about the mpfsctl commands.	103
mpfsctl config	Specifies the maximum amount of memory the driver uses.	104
mpfsctl disable	Terminates a session.	104
mpfsctl diskspeed	Shows all LUNs that the MPFS kernel driver can access and tests the raw operating speeds.	105
mpfsctl enable	Establishes a session and optionally sets MPFS share properties.	105
mpfsctl inq	Lists the MPFS disks that the driver sees and groups multiple paths together.	106
mpfsctl list	Lists all enabled MPFS shares and their properties.	107
mpfsctl reset	Clears previously accumulated counter statistics.	107
mpfsctl stats	Shows MPFS driver statistics.	108
mpfsctl version	Shows the MPFS software version number.	109
mpfsinfo	Validates the installation of the Windows server.	110
mpfsinq	Shows MPFS devices and device IDs visible to the Windows server, along with the path where the disks are mapped.	112

mpfscope

The **mpfscope** command is a tool used to duplicate source directories and files to multiple targets specified in the command line. This command uses multiple threads to perform large I/O operations to utilize the I/O characteristics of MPFS.

The **mpfscope** command also acts as a test program to perform I/O operations and measure the total throughput of the file systems.

Syntax

```
mpfscope [-i ii] [-t tt] [-r] [-v] [-h] SOURCE DEST1
DEST2 ...
```

```
mpfscope [-test type] [-i ii] [-s ss] [-v] [-h] FILE1
FILE2 ...
```

where:

`[-i ii]` specifies the size of each I/O operation. If not specified, the default is 1 MB. The size must align to 8K.

`[-t tt]` specifies the number of worker threads. If not specified, the default is eight threads.

`[-r]` copies the directories in order as specified.

`[-v]` runs in verbose mode printing out additional information.

`[-h]` prints out the help information.

`[SOURCE]` specifies the source files or directories for copying. This option accepts the wild card "*" designator.

`[DEST1 DEST2]` specifies the destination where the source files or directories are copied. Specifying multiple destinations is allowed.

`[-test type]` indicates which test is performed.

`[-s ss]` specifies the size of the files in MB. This option is used only for the write test. If not specified, the default is 1 GB.

`[FILE1 FILE2]` specifies the files used for testing.

Example 1 Copy a single file named mpfscopy.exe to the root directory of the D: drive:

```
$ mpfscopy mpfscopy.exe d:\
Enumerating files...
Start...

Done.

Elapsed time 0.168069 s
Average throughput 1.162098 M/s
```

Example 2 Copy all the files and directories starting with "m" to drive G: and enable the verbose mode. The program prints out the filename of each file it is about to copy:

```
$ mpfscopy -r -v m* g:\
Enumerating files...
Start...
mpfscopy.exe
mpfsinfo.exe
MSDOS.SYS

Done.

Elapsed time 0.355408 s
Average throughput 0.645715 M/s
```

Example 3 Start a write test with a 128 MB file and 1 MB of I/O:

```
$ mpfscopy -test w -i 1048576 -s 128 g:\test_file.tmp
Generating file list...
Start...

Done.

Elapsed time 4.343671 s
Average throughput 29.468161 M/s
```

mpfsctl -? The **mpfsctl -?** or **mpfsctl** command shows information about the **mpfsctl** commands.

Syntax mpfsctl -? or mpfsctl

Example The following commands display information about the **mpfsctl** commands:

```
mpfsctl -?
```

to enable MPFS:

```
mpfsctl enable server share [[-]wflush] [[-]persistent] [RA:b_size]
[WB:b_size] [MRIS:b_size] [[-]ACPolicy]
```

to disable MPFS:

```
mpfsctl disable server share
```

where:

server	- server name (DART, 125.127.128.12, etc.)
share	- share name
[RA:b_size]	- Read ahead section size in MB, 0:disable read ahead (Min:1, Max:10, default "RA:1")
[WB:b_size]	- Write behind section size in MB, 0:disable write behind (Min:1, Max:10, default "WB:1")
[MRIS:b_size]	- Minimum read I/O size in KB (Min:1, Max:1024, default "MRIS:8")
[[-]ACPolicy]	- Asynch Close Policy (Default -ACPolicy)
[[-]wflush]	- turn on ("wFlush") or off ("-wFlush", default) flushing after each write
[[-]persistent]	- enable ("persistent") the connection at boot time (default "-persistent")

to list MPFS Shares:

```
mpfsctl list
```

to display driver statistics:

```
mpfsctl stats [ahead]
mpfsctl reset
```

to get the disks that is visible to driver now:

```
mpfsctl inq
```

to config the global options by GUI:

```
mpfsctl config
```

to display the product version:

```
mpfsctl version
```

to test MPFS kernel driver LUN access speed (MB/s):

```
mpfsctl diskspeed
```

mpfsctl config

The **mpfsctl config** command opens the **Global Options** window as shown in [Figure 26 on page 104](#).

Syntax `mpfsctl config`

Example `$ mpfsctl config`




Figure 26 Global Options window

where:

Max Memory Driver Uses: — Select the maximum amount of memory the driver uses (in MB) from the list box. Click **OK** to close the **Global Options** window.

mpfsctl disable

The **mpfsctl disable** command terminates a session.

Syntax `mpfsctl disable <server> <share>`

where:

`<server>` = MPFS VNX

`<share>` = MPFS share on `<server>`

Example The following command terminates the session on the MPFS share named `c2s3fs5` on the server named `c2s3e0`:

```
$ mpfsctl disable c2s3e0 c2s3fs5
```

Request executed.

mpfsctl diskspeed

The **mpfsctl diskspeed** command shows all LUNs that the mpfs kernel driver can access and tests the raw operating speeds.

Syntax `mpfsctl diskspeed`

Example The following command tested three out of three LUNs and shows their raw operating speeds:

```
$ mpfsctl diskspeed
MPFS kernel driver tested 3 out of 3 LUNs:
    0001874307271FA0-0003: 25 MB/s
    0001874307271FA0-0004: 23 MB/s
    0001874307271FA0-0005: 23 MB/s
Request executed.
```

mpfsctl enable

The **mpfsctl enable** command establishes a session and, optionally, sets properties on an MPFS share.

Syntax `mpfsctl enable <server> <share> [[-]wflush] [[-]persistent] [RA:b_size] [WB:b_size] [MRIS:b_size] [[-]ACPolicy]`

where:

`<server>` = VNX MPFS for File

`<share>` = MPFS share on `<server>`

`[wflush]` specifies synchronous flushing where the Windows server flushes metadata updates to the VNX for file after every write.

`[-wflush]` specifies asynchronous flushing which provides superior performance, is preferable in most situations, and is the default. However, in certain situations you can specify synchronous flushing. For example, specify synchronous flushing when two collaborating applications on two computers require a high degree of serialization.

`[persistent]` specifies that the connection to the server is retained across restarts.

`[-persistent]` specifies that the connection is terminated when the server is restarted. This is the default.

`[RA: b_size]` prefetches data that an application requests to improve read operation performance.

`[WB: b_size]` stores data in a buffer that is written to a disk drive and sends the data when the system is idle to improve the write speed.

[MRIS: b_size] is the minimum size of the amount of data read from a file.

[-ACPolicy] closes the file before sending data and waits for an acknowledgement before sending more data.

Example The following command establishes a session on the MPFS share named 172.24.107.242 on the server named s2_sh1, with synchronous flushing:

```
$ mpfsctl enable 172.24.107.242 s2_sh1 wflush
```

Request executed.

mpfsctl inq

The **mpfsctl inq** command lists the MPFS disks visible to the driver and groups multiple paths together.

Syntax mpfsctl inq

Example The following command lists the MPFS disks that the driver sees and groups multiple paths together:

```
$ mpfsctl inq
```

```

WWN                               VNX for file Signature
60:06:01:60:2e:9b:1e:00:54:81:38:a4:bc:62:dc:11 HK1908073200170000-0017
60:06:01:60:2e:9b:1e:00:1a:47:c2:e9:a7:62:dc:11 HK1908073200170000-0008
60:06:01:60:1d:84:1e:00:52:c4:d4:ef:a7:62:dc:11 HK1908073200170000-0018
60:06:01:60:2e:9b:1e:00:3e:06:ec:f5:a7:62:dc:11 HK1908073200170000-0009
60:06:01:60:1d:84:1e:00:0e:fa:f4:fb:a7:62:dc:11 HK1908073200170000-0019
60:06:01:60:2e:9b:1e:00:a0:d9:09:02:a8:62:dc:11 HK1908073200170000-0010
60:06:01:60:1d:84:1e:00:40:e0:25:08:a8:62:dc:11 HK1908073200170000-0020
60:06:01:60:2e:9b:1e:00:10:d4:2e:0e:a8:62:dc:11 HK1908073200170000-0011
60:06:01:60:1d:84:1e:00:b0:da:4a:14:a8:62:dc:11 HK1908073200170000-0021
60:06:01:60:2e:9b:1e:00:d4:f3:d6:22:a8:62:dc:11 HK1908073200170000-0012
60:06:01:60:1d:84:1e:00:b2:0e:e7:28:a8:62:dc:11 HK1908073200170000-0022
60:06:01:60:2e:9b:1e:00:62:73:62:2f:a8:62:dc:11 HK1908073200170000-0013
60:06:01:60:1d:84:1e:00:a2:9f:a4:35:a8:62:dc:11 HK1908073200170000-0023
60:06:01:60:2e:9b:1e:00:42:a6:c0:3b:a8:62:dc:11 HK1908073200170000-0014
60:06:01:60:1d:84:1e:00:c6:5e:ce:41:a8:62:dc:11 HK1908073200170000-0024
60:06:01:60:2e:9b:1e:00:0c:03:e8:47:a8:62:dc:11 HK1908073200170000-0015
60:06:01:60:1d:84:1e:00:36:59:f3:4d:a8:62:dc:11 HK1908073200170000-0025
60:06:01:60:2e:9b:1e:00:54:bf:4f:54:a8:62:dc:11 HK1908073200170000-0016
60:06:01:60:1d:84:1e:00:12:4b:d2:5a:a8:62:dc:11 HK1908073200170000-0026

```

mpfsctl list

The **mpfsctl list** command lists enabled MPFS shares and their properties.

The command shows the following information about each share:

- ◆ Server name
- ◆ Share name
- ◆ Minimal read I/O size in KB
- ◆ Read ahead section size in MB
- ◆ Write behind section size in MB
- ◆ Flag settings for flush after every write (wflush is on or -wflush is off)

Syntax `mpfsctl list`

Example The following example shows an MPFS share enabled on server g10dvt17s601:

```
$ mpfsctl list
...MPFS Share...
Server: g10dvt17s601, Share: svr6sh1
Minimal read IO size: 8KB
Read ahead section size: 1MB.
Write behind section size: 1 MB.
Flags: -wflush
Valid MPFS connection mapped to 2 disks
=== End of List ===
```

mpfsctl reset

The **mpfsctl reset** command clears previously accumulated counter statistics. Use the **mpfsctl reset** command to reset the counters to 0 before executing **mpfsctl stats**.

Syntax `mpfsctl reset`

Example The following command resets all statistic counters to 0:

```
$ mpfsctl reset
Request executed.
```

Output explanation

The output indicates the command was executed.

mpfsctl stats

The **mpfsctl stats** command shows MPFS driver statistics.

<disk>, an optional argument, causes the display to include statistics on disk response time. To display disk response-time statistics, execute the **mpfsctl stats disk** command twice:

- ◆ The first execution of **mpfsctl stats disk** turns on a collection of disk response-time statistics.
- ◆ The second execution shows all activity, including disk response time, since the first execution of the **mpfsctl stats disk** command.

All subsequent executions of **mpfsctl stats disk** display all activity, including disk response time, since the previous execution of the command.

To turn off collection and display of disk response-time statistics, simply execute the **mpfsctl stats** command without the optional **disk** argument.

Syntax `mpfsctl stats <disk>`

where *<disk>* specifies collection and display of statistics on disk response time.

Example The following command shows all statistics including disk response-time statistics:

```
$ mpfsctl stats
MPFS Statistics since last query:
  Bytes Written = 15466496
  Bytes Read = 0
  Last FMP error = 00000000
  Number of FMP errors = 0
  Number of FMP OKs = 252
  Number of FMP Fail Overs to CIFS = 0
  Total requests optimized to CIFS = 0
  Total requests passed to CIFS = 0

FMP error statistics:
  Number of FMP_SESSION_LOST = 1

External error statistic:
  Number of HR_GENERAL_TIMEOUT = 12
  Number of HR_NO_SERVICE = 6
```

Output explanation

The output shows the following activity since the last execution of the **mpfsctl stats** command:

- ◆ Number of bytes written directly to the storage-system volume from the server by using the FMP protocol: 15466496
- ◆ Number of bytes read directly from the storage-system volume to the server by using the FMP protocol: 0
- ◆ Hexadecimal error code of the most recent FMP error: 00000000
- ◆ Total number of FMP errors: 0
- ◆ Total number of successfully completed I/O operations going directly to and from the disk by using the FMP protocol: 252
- ◆ Total number of attempted FMP operations that failed over to CIFS: 0
- ◆ Total number of I/O requests optimized to CIFS: 0
- ◆ Total number of I/O requests passed to CIFS: 0
- ◆ Number of FMP sessions that were not recovered: 1
- ◆ Number of times the session timed out: 12
- ◆ Number of times service was not available: 6

mpfsctl version

The **mpfsctl version** command shows the MPFS software version number.

Syntax

```
mpfsctl version
```

Example

The following command shows which version of the MPFS software is installed on the server:

```
$ mpfsctl version
6.0.1.3
```

Output explanation

The output shows the MPFS software version number.

mpfsinfo

Use the **mpfsinfo** command to validate a Windows server installation by querying an FMP server (VNX Data Mover) and validating that the Windows server can access all the disks required to use MPFS for each exported file system.

The user must have CIFS authorization to the server by using the NETBIOS command **net use \\server\share** (**net use /?** provides help on the command) for the **mpfsinfo.exe** command to access the list of shares used by the server.

The user must supply the name or IP address of at least one FMP server. Multiple FMP servers are allowed, in which case the validation is done for the exported file systems on all the listed servers.

Syntax `mpfsinfo [-v] [-h] <fmpserver>`

where:

`[-v]` runs in verbose mode printing out additional information.

`[-h]` prints out the help information.

`<fmpserver>` is the name or IP address of the VNX Data Mover.

Example 1 The following command shows how to examine a VNX Data Mover with an IP address of 172.24.107.243 showing the IP address, server name, and disks:

```
$ mpfsinfo -v 172.24.107.243
```

Output:

```
172.24.107.243 : svr2sh1 - OK
    FNM000844000930000-0012 : OK
    FNM000844000930000-0007 : OK
172.24.107.243 : svr2sh2 - OK
    FNM000844000930000-0012 : OK
    FNM000844000930000-0007 : OK
172.24.107.243 : svr2sh3 - OK
    FNM000844000930000-0012 : OK
    FNM000844000930000-0007 : OK
```

Example 2 The following command shows how to examine a VNX Data Mover with an IP address of 172.24.107.243 showing the IP address and server name:

```
$ mpfsinfo -v 172.24.107.243
```

Output:

```
172.24.107.243 : svr2sh1 - OK
```

```
172.24.107.243 : svr2sh2 - OK
```

```
172.24.107.243 : svr2sh3 - OK
```

Example 3 The following command prints out the help information:

```
$ mpfsinfo -h
```

```
Usage: mpfsinfo [-v] [-h] <fmpserver>
```

```
-v          run in verbose mode printing out additional  
           information
```

```
-h          print out this help information
```

```
<fmpserver> name of the FMP server (VNX Data Mover)
```

mpfsinq

The **mpfsinq** command allows the display of MPFS storage devices and device IDs, along with the device path.

Syntax `mpfsinq`

Example To display MPFS storage devices, type:

```
$ mpfsinq
```

Output:

Port	Path	Tgt	LUN	Vendor	ID	Product	ID	Capacity	Signature
2	0	0	0	DGC	RAID	5		440.8	
2	0	0	1	DGC	RAID	5		440.8	
2	0	0	2	DGC	RAID	5		1073.5	HK1908073200170000-0008
2	0	0	3	DGC	RAID	5		1073.5	
2	0	0	4	DGC	RAID	5		536.7	HK1908073200170000-0009
2	0	0	5	DGC	RAID	5		536.7	
2	0	0	6	DGC	RAID	5		536.7	HK1908073200170000-0010
2	0	0	7	DGC	RAID	5		536.7	HK1908073200170000-0020
2	0	0	8	DGC	RAID	5		536.7	HK1908073200170000-0011
2	0	0	9	DGC	RAID	5		536.7	
2	0	0	10	DGC	RAID	5		536.7	HK1908073200170000-0012
2	0	0	11	DGC	RAID	5		536.7	
2	0	0	12	DGC	RAID	5		1073.5	HK1908073200170000-0013
2	0	0	13	DGC	RAID	5		1073.5	
2	0	0	14	DGC	RAID	5		536.7	HK1908073200170000-0014
2	0	0	15	DGC	RAID	5		536.7	
2	0	0	16	DGC	RAID	5		536.7	HK1908073200170000-0015
2	0	0	17	DGC	RAID	5		536.7	
2	0	0	18	DGC	RAID	5		536.7	HK1908073200170000-0016
2	0	0	19	DGC	RAID	5		536.7	
2	0	1	0	DGC	RAID	5		440.8	
2	0	1	1	DGC	RAID	5		440.8	
2	0	1	2	DGC	RAID	5		1073.5	HK1908073200170000-0008
2	0	1	3	DGC	RAID	5		1073.5	
2	0	1	4	DGC	RAID	5		536.7	HK1908073200170000-0009
2	0	1	5	DGC	RAID	5		536.7	
2	0	1	6	DGC	RAID	5		536.7	HK1908073200170000-0010
2	0	1	7	DGC	RAID	5		536.7	HK1908073200170000-0020
2	0	1	8	DGC	RAID	5		536.7	HK1908073200170000-0011
2	0	1	9	DGC	RAID	5		536.7	

Error Messages and Troubleshooting

This appendix describes Windows server error messages, informational messages, and references to troubleshooting documentation. This appendix includes the following topics:

- ◆ [EMC VNX MPFS messages sent to the event log.....](#) 114
- ◆ [Messages displayed in error message boxes.....](#) 115
- ◆ [Using Windows Event Viewer with MPFS.....](#) 116
- ◆ [Troubleshooting MPFS software.....](#) 121

EMC VNX MPFS messages sent to the event log

[Table 6 on page 114](#) shows messages that MPFS sends to the Windows Event Viewer System Log for an iSCSI configuration.

Table 6 MPFS iSCSI error messages

Error message	Explanation
<server_name> - lost session.	The Data Mover temporarily lost FMP communications.
<server_name> - Session recovered.	The Data Mover recovered FMP communications.
<server_name\share_name>\??\PhysicalDriveX - Device Error.- Session recovered.	A block device has become disabled.
<server_name\share_name> - I\O block device is back to normal.	An I/O block device has returned to normal after a device error.
<server_name\share_name> - Unmatched disk signature.	A disk you are attempting to use has an unmatched disk signature. VNX disk signatures are required.

[Table 7 on page 114](#) shows messages that MPFS sends to the Windows Event Viewer System Log for an FC configuration.

Table 7 MPFS FC error messages (page 1 of 2)

Error message	Explanation
BfxCommonDeviceControl: Wrong API_MAGIC value -> BFX_VERSION_MISMATCH!	Component versions are mismatched. Uninstall and then reinstall the MPFS software.
Could not Add Connection for <server>	The Windows server could not connect to the VNX for file.
Wrong notification code. Server = <server>	The VNX for file has malfunctioned, possibly because of a version mismatch between the VNX for file and the Windows server.
FMP_FLUSH failed for the <server>. ret = 0x99999	The server has probably crashed, or a cable is disconnected.
LowAlignedIo: cannot allocate IRP	An I/O Request Packet is not allocated, probably because of a lack of system resources.
FMP_MOUNT failed for the server: <server>. ret = 0x99999	The file system is not mounted.

Table 7 MPFS FC error messages (page 2 of 2)

Error message	Explanation
SessionCleanUpApc for the server: <server>. CreateSession failed.	An attempt to reconnect to the server has failed.
SessionCleanUpApc for the server: <server>. Session created.	An attempt to reconnect to the server has succeeded.
Unable to match disk signatures.	The Windows server was unable to locate or see one or more VNX disks. The mpfsctl enable command fails.

“Using Windows Event Viewer with MPFS” on page 116 explains how to view these error messages within the Windows environment.

Messages displayed in error message boxes

Table 8 on page 115 shows MPFS file system messages in standard error message boxes.

Table 8 Windows messages displayed in message boxes

Error message	Explanation
An incorrect version of WINSOCKET.DLL has been found. Network operations may not operate correctly.	MPFS has found and loaded an out-of-date version of WINSOCKET.DLL, a file used for initializing communication with the MPFS driver.
Windows sockets initialization failed.	Windows socket initialization failed. MPFS could not load or initialize WINSOCKET.DLL.
An error occurred while attempting to initialize the MPFS driver.	MPFS could not initialize the MPFS driver. The driver is not running.

Using Windows Event Viewer with MPFS

The Windows Event Viewer System Log contains three System Log types: Warning, Error, and Information.

Log type: Warning

A Warning log type is caused by multiple issues. A few examples are:

- ◆ A network connection was temporarily lost.
- ◆ A VNX for file experienced an error attempting to validate a Windows server's computer account.
- ◆ The Windows Time Service was not able to find a domain controller.

Log type: Error

Examples of Error log type messages are described in [Table 6 on page 114](#) and [Table 7 on page 114](#).

Note: Error messages described are MPFS specific. Other Windows specific error messages are displayed, some of which are shown in [Table 8 on page 115](#).

Log type: Information

An Information log type can describe several kinds of events. [Figure 31 on page 120](#) shows an example of an informational event.

Use the following steps to access the Windows Event Viewer:

1. Select **Start > Run** and type **eventvwr** in the **Run** window as shown in [Figure 27 on page 117](#).




Figure 27 Opening the Windows Event Viewer

2. Click **OK**, and the **Event Viewer** window opens as shown in [Figure 28 on page 118](#).




Figure 28 Event Viewer window

- To view an error in the System Log, double-click one of the **Error** icons listed under the Type column as shown in [Figure 29 on page 118](#).




Figure 29 Selecting an Error message

4. The **Event Properties** window opens and the selected error appears as shown in [Figure 30 on page 119](#).




Figure 30 Event Properties window

Note: The sequence k2seqa31s8 in [Figure 30 on page 119](#) is an example of a VNX for file name. Server names are different at each site.

5. Use the up and down arrows in the **Event Properties** window to scroll to the next event as shown in [Figure 31 on page 120](#).




Figure 31 Up and down arrows on Event Properties window

6. The **Information** log type in [Figure 31 on page 120](#) shows a description showing that the session recovered from the previous error shown in [Figure 30 on page 119](#).

Each MPFS error message is seen by using this method.

Troubleshooting MPFS software


This section lists procedures, problems, causes, and solutions in troubleshooting the MPFS software.

The *EMC VNX MPFS for Windows Clients Release Notes* provide additional information on troubleshooting, known problems, and limitations.

Confirming MPFS software installation

After installing the MPFS software, verify that the software was loaded correctly:

1. Type **mpfsctl version** in the **Command Prompt** window to confirm that the MPFS software is properly installed as shown in [Figure 32 on page 121](#).



```
C:\>mpfsctl version
6.0.1.0
C:\>
```

Figure 32 Using mpfsctl version to verify MPFS software is installed

If the MPFS version number shown in [Figure 32 on page 121](#) matches the version installed, the MPFS software was installed correctly.

2. If a message stating that the MPFS software is not installed appears, select **Start > Settings > Control Panel** and double-click **Add or Remove Programs**.
3. Uninstall the MPFS software as shown in [“Uninstalling the MPFS software” on page 94](#).
4. Restart the Windows server.
5. Install the MPFS software as shown in [“Installing the MPFS software” on page 86](#).
6. Restart the Windows server.
7. Verify the MPFS software version as shown in [step 1](#).

Installing difficulties

You can encounter the following problem during installation of the MPFS software.

Problem During installation, a message is displayed indicating you do not have proper access to install the service.

Cause

You lack administrative rights on the server on which you are installing the MPFS software.

Solution

Log in to the server with an administrative account and start the installation again.

Uninstalling difficulties

You can encounter the following problem during removal of the MPFS software.

Problem Locked files remain on the server after uninstalling the MPFS software.

Cause

You failed to restart the server after you uninstalled the MPFS software.

Solution

Restart the server and the locked files are removed.

Known problems and limitations

The *EMC VNX MPFS for Windows Clients Release Notes* provide known problems and limitations for MPFS clients.

This glossary defines terms useful for MPFS administrators.

C

**Challenge
Handshake
Authentication
Protocol (CHAP)**

Access control protocol for secure authentication using shared passwords called secrets.

client

Front-end device that requests services from a server, often across a network.

**command line
interface (CLI)**

Interface for typing commands through the Control Station to perform tasks that include the management and configuration of the database and Data Movers and the monitoring of statistics for the VNX for file cabinet components.

**Common Internet File
System (CIFS)**

File-sharing protocol based on the Microsoft Server Message Block (SMB). It allows users to share file systems over the Internet and intranets.

Control Station

Hardware and software component of the VNX for file that manages the system and provides the user interface to all VNX for file components.

D

daemon UNIX process that runs continuously in the background, but does nothing until it is activated by another process or triggered by a particular event.

Data Mover In a VNX for file, a cabinet component running its own operating system that retrieves data from a storage device and makes it available to a network client. This is also referred to as a blade.

disk volume On VNX for file, a physical storage unit as exported from the system. All other volume types are created from disk volumes. See also [metavolume](#), [slice volume](#), [stripe volume](#), and [volume](#).

E

extent Set of adjacent physical blocks.

F

fallthrough Fallthrough occurs when MPFS temporarily employs the NFS or CIFS protocol to provide continuous data availability, reliability, and protection while block I/O path congestion or unavailability is resolved. This fallthrough technology is seamless and transparent to the application being used.

Fast Ethernet Any Ethernet specification with a speed of 100 Mb/s. Based on the IEEE 802.3u specification.

Fibre Channel Nominally 1 Gb/s data transfer interface technology, although the specification allows data transfer rates from 133 Mb/s up to 4.25 Gb/s. Data can be transmitted and received simultaneously. Common transport protocols, such as Internet Protocol (IP) and Small Computer Systems Interface (SCSI), run over Fibre Channel. Consequently, a single connectivity technology can support high-speed I/O and networking.

File Mapping Protocol (FMP) File system protocol used to exchange file layout information between an MPFS client and the VNX for file. See also [Multi-Path File Systems \(MPFS\)](#).

file system Method of cataloging and managing the files and directories on a system.

G

gateway VNX for file that is capable of connecting to multiple systems, either directly (direct-connected) or through a Fibre Channel switch (fabric-connected).

Gigabit Ethernet Any Ethernet specification with a speed of 1000 Mb/s. IEEE 802.3z defines Gigabit Ethernet over fiber and cable, which has a physical media standard of 1000Base-X (1000Base-SX short wave, 1000Base-LX long wave) and 1000Base-CX shielded copper cable. IEEE 802.3ab defines Gigabit Ethernet over an unshielded twisted pair (1000Base-T).

H

host Addressable end node capable of transmitting and receiving data.

I

Internet Protocol (IP) Network layer protocol that is part of the Open Systems Interconnection (OSI) reference model. IP provides logical addressing and service for end-to-end delivery.

Internet Protocol address (IP Address) Address uniquely identifying a device on any TCP/IP network. Each address consists of four octets (32 bits), represented as decimal numbers separated by periods. An address is made up of a network number, an optional subnetwork number, and a host number.

Internet SCSI (iSCSI) Protocol for sending SCSI packets over TCP/IP networks.

iSCSI initiator iSCSI endpoint, identified by a unique iSCSI name, which begins an iSCSI session by issuing a command to the other endpoint (the target).

iSCSI target iSCSI endpoint, identified by a unique iSCSI name, which executes commands issued by the iSCSI initiator.

K

kernel Software responsible for interacting most directly with the computer's hardware. The kernel manages memory, controls user access, maintains file systems, handles interrupts and errors, performs input and output services, and allocates computer resources.

L

logical device	One or more physical devices or partitions managed by the storage controller as a single logical entity.
logical unit (LU)	For iSCSI on VNX for file, a logical unit is an iSCSI software feature that processes SCSI commands, such as reading from and writing to storage media. From a iSCSI host perspective, a logical unit appears as a disk drive.
logical unit number (LUN)	Identifying number of a SCSI or iSCSI object that processes SCSI commands. The LUN is the last part of the SCSI address for a SCSI object. The LUN is an ID for the logical unit, but the term is often used to refer to the logical unit itself.
logical volume	Logical devices aggregated and managed at a higher level by a volume manager. See also <i>logical device</i> .

M

metadata	Data that contains structural information, such as access methods, about itself.
metavolume	On a VNX for file, a concatenation of volumes, which can consist of disk, slice, or stripe volumes. Also called a hypervolume or hyper. Every file system must be created on top of a unique metavolume. See also <i>disk volume, slice volume, stripe volume, and volume</i> .
mirrored pair	Logical volume with all data recorded twice, once on each of two different physical devices.
mirroring	Method by which the storage system maintains two identical copies of a designated volume on separate disks.
mount	Process of attaching a subdirectory of a remote file system to a mount point on the local machine.
mount point	Local subdirectory to which a mount operation attaches a subdirectory of a remote file system.

MPFS over iSCSI	Multi-Path File System over iSCSI-based clients. MPFS client running an iSCSI initiator works in conjunction with an IP-SAN switch containing an iSCSI to SAN blade. The IP-SAN blade provides one or more iSCSI targets that transfer data to the storage area network (SAN) systems. See also <i>Multi-Path File Systems (MPFS)</i> .
MPFS session	Connection between an MPFS client and a VNX for file.
MPFS share	Shared resource designated for multiplexed communications by using the MPFS file system.
Multi-Path File Systems (MPFS)	VNX for file feature that allows heterogeneous servers with MPFS software to concurrently access, directly over Fibre Channel or iSCSI channels, shared data stored on a EMC Symmetrix or VNX for block. MPFS adds a lightweight protocol called File Mapping Protocol (FMP) that controls metadata operations.
N	
nested mount file system (NMFS)	File system that contains the nested mount root file system and component file systems.
nested mount file system root	File system on which the component file systems are mounted read-only, except for mount points of the component file systems.
network-attached storage (NAS)	Specialized file server that connects to the network. A NAS device, such as VNX for file, contains a specialized operating system and a file system, and processes only I/O requests by supporting popular file sharing protocols such as NFS and CIFS.
network file system (NFS)	Network file system (NFS) is a network file system protocol allowing a user on a client computer to access files over a network as easily as if the network devices were attached to its local disks.
P	
PowerPath	EMC host-resident software that integrates multiple path I/O capabilities, automatic load balancing, and path failover functions into one comprehensive package for use on open server platforms connected to Symmetrix or VNX for block.

R

Redundant Array of Independent Disks (RAID) Method for storing information where the data is stored on multiple disk drives to increase performance and storage capacities and to provide redundancy and fault tolerance.

S

server Device that handles requests made by clients connected through a network.

slice volume On a VNX for file, a logical piece or specified area of a volume used to create smaller, more manageable units of storage. See also *disk volume*, *metavolume*, *stripe volume*, and *volume*.

small computer system interface (SCSI) Standard set of protocols for host computers communicating with attached peripherals.

storage area network (SAN) Network of data storage disks. In large enterprises, a SAN connects multiple servers to a centralized pool of disk storage. See also *network-attached storage (NAS)*.

storage processor (SP) Storage processor on a VNX for block. On a VNX for block, a circuit board with memory modules and control logic that manages the VNX for block I/O between the host's Fibre Channel adapter and the disk modules.

Storage processor A (SP A) Generic term for the first storage processor in a VNX for block.

Storage processor B (SP B) Generic term for the second storage processor in a VNX for block.

stripe size Number of blocks in one stripe of a stripe volume.

stripe volume Arrangement of volumes that appear as a single volume. Allows for stripe units that cut across the volume and are addressed in an interlaced manner. Stripe volumes make load balancing possible. See also *disk volume*, *metavolume*, *slice volume*, and *volume*.

Symmetrix Remote Data Facility (SRDF)	EMC technology that allows two or more Symmetrix systems to maintain a remote mirror of data in more than one location. The systems can be located within the same facility, in a campus, or hundreds of miles apart using fiber or dedicated high-speed circuits. The SRDF family of replication software offers various levels of high-availability configurations, such as SRDF/Synchronous (SRDF/S) and SRDF/Asynchronous (SRDF/A).
T	
Transmission Control Protocol (TCP)	Connection-oriented transport protocol that provides reliable data delivery.
U	
unified storage	VNX for file that is connected to a captive system that is not shared with any other VNX for files and is not capable of connecting to multiple systems.
V	
Virtual Storage Area Network (VSAN)	SAN that can be broken up into sections allowing traffic to be isolated within the section.
VNX	EMC network-attached storage (NAS) product line.
VNX for block	EMC midrange block system.
VNX OE	Embedded operating system in VNX for block disk arrays.
volume	On a VNX for file, a virtual disk into which a file system, database management system, or other application places data. A volume can be a single disk partition or multiple partitions on one or more physical drives. See also <i>disk volume</i> , <i>metavolume</i> , <i>slice volume</i> , and <i>stripe volume</i> .

A

- Access Logix configuration 60
- accessing storage 64
- administering MPFS 99
- architecture
 - VNX
 - Fibre Channel 19, 35
 - iSCSI 21, 35
 - iSCSI/FC 36
 - VNX MPFS
 - over FC 19
 - over iSCSI 21
 - over iSCSI/FC 22
 - VNX VG2/VG8
 - Fibre Channel 20, 35
 - iSCSI 22, 23, 36
 - iSCSI/FC 22, 36
- arraycompath 60, 61, 62, 64, 65, 74, 77
- asynchronous flushing 81, 105
- authentication, CHAP 69

B

- best practices
 - file system 38, 42
 - LUNs 43
 - MPFS 28, 42, 49
 - MPFS threads 28
 - storage configuration 29
 - stripe size 42, 49
 - VNX for block 29, 54
 - VNX for file volumes 49
 - VNX for file with MPFS 28
 - VNX VG2/VG8 54

C

- Challenge Handshake Authentication Protocol.
See CHAP
- CHAP
 - one-way authentication 30
 - reverse authentication 68
 - secret 68
 - session authentication 69
- command line interface. See mpfs commands
- command line interface. See mpfsctl commands
- commands
 - mpfs copy 101
 - mpfs info 110
 - mpfs inq 112
 - mpfsctl config 104
 - mpfsctl disable 104
 - mpfsctl diskspeed 105
 - mpfsctl enable 105
 - mpfsctl inq 106
 - mpfsctl list 107
 - mpfsctl reset 107
 - mpfsctl stats 108
 - mpfsctl version 109
- comments 15
- configuration
 - overview 26
 - planning checklist 33
- configuring
 - Gigabit Ethernet ports 36
 - iSCSI target 67
 - storage 64
 - storage access 54
 - VNX for file 43

- VNX with MPFS configuration 28
 - zones 54
- connecting to a server 105
- creating
 - file system 43
 - metavolume 50
 - mountpoint 51
 - MPFS file system 50
 - security file 56
 - storage groups 62
 - stripe 48

D

- disabling
 - active MPFS session 97
 - arraycommpath 65, 74, 77
 - Asynch Close Policy 83
 - failovermode 65, 74, 77
 - MPFS using mpfsctl disable 104
- displaying
 - accessible LUNs 105
 - disks 45
 - MPFS devices 112
 - MPFS properties 80
 - MPFS shares 107
 - MPFS software version 109
 - MPFS statistics 108
 - MPFS storage devices 112
 - standard error message boxes 115

E

- enabling
 - arraycommpath 65, 74, 77
 - failovermode 65, 74, 77
 - MPFS using mpfsctl enable 105
- error messages 113

F

- failovermode 60, 61, 62, 64, 65, 74, 77
- Fibre Channel
 - adding hosts to storage groups 65, 74
 - driver installation 64
 - switch installation 55
 - switch requirements 39
- Fibre Channel over Ethernet (FCoE) 19

- File Mapping Protocol (FMP) 24
- file system
 - creating 43, 50
 - exporting 52
 - mounting 51
 - names of mounted 46
 - names of unmounted 47
 - setup 42
- firewall FMP ports 90, 91
- flush policy, setting
 - through the Properties dialog box 81
 - using mpfsctl enable 105

G

- Gigabit Ethernet port configuration 36

I

- I/O sizes 28
- installing
 - MPFS software 88
 - MPFS software, troubleshooting 122
 - storage configuration 38
- Installing
 - start the installation 88
- iSCSI CHAP authentication 30
- iSCSI discovery address 37
- iSCSI initiator
 - configure iSCSI discovery 69
 - configure with CHAP 68
 - configuring 64, 67
 - configuring ports 31
 - connecting to a SAN switch 29
 - installing 67
 - IQN to define a host 76
 - IQN to define a server 64
 - log into iSCSI target 71
 - names 55
 - register 68
 - show IQN name 68
- iSCSI port configuration 57
- iSCSI target configuration 57

L

- LUNs
 - access speed 103

- access target LUNs 71
- accessible by Data Mover 43
- adding 62
- best practices 43
- displaying all 105, 112
- failover 64
- maximum number supported 43
- mixed not supported 38
- total usable capacity 43

M

- managing using mpfs commands 99
- metavolume 50
- mounting
 - MPFS file system 51
 - MPFS file system troubleshooting 114
- mountpoint, creating 51
- mpfs commands
 - mpfs copy 101
 - mpfs info 110
 - mpfs inq 112
- MPFS configurations 19, 21, 22
- MPFS devices, displaying 112
- MPFS file system
 - creating 43, 50
 - exporting 52
 - mounting 51
 - setup 42
 - storage requirements 38
 - unmounting 29
- MPFS overview 18
- MPFS session
 - establishing using mpfsctl enable command 105
 - retaining across restarts 105
 - retaining properties 105
 - starting using mpfsctl enable 105
 - stopping through the Properties dialog box 97
 - stopping using mpfsctl disable 104
- MPFS share
 - connecting 105
 - disconnecting 104
 - enabling 81
 - map to a network drive 79
 - properties

- listing 107
 - retaining across restarts 81, 105
 - setting 105
- MPFS software
 - before installing 86
 - disabling active MPFS session 97
 - flush properties 81
 - install over existing 92
 - installation 121, 122
 - installing a later version 93
 - installing from a CD 87
 - installing from EMC Online Support 87
 - kernel driver 103, 105
 - managing using mpfsctl commands 99
 - mpfs commands 100
 - mpfsctl commands 103
 - Properties window 81
 - restarting 81, 93
 - start installation 88
 - starting 53
 - troubleshooting 121
 - uninstalling 94
 - uninstalling earlier version 95
 - upgrading 92
 - upgrading from an existing version 92
 - version number 109
 - Windows Event Viewer 116
- MPFS threads 28
- mpfsctl commands
 - mpfs inq 106
 - mpfsconfig 104
 - mpfsctl disable 104
 - mpfsctl enable 105
 - mpfsctl list 107
 - mpfsctl reset 107
 - mpfsctl stats 108
 - mpfsctl version 109
 - mpfsdiskspeed 105

O

- overview of configuring MPFS 26

P

- performance 28
 - asynchronous flushing 81, 105
 - file system 49

- Gigabit Ethernet ports 36
- improve read operation 105
- iSCSI ports 40
- iSCSI target ports 54
- MPFS 28, 29, 49
- MPFS threads 28
- MPFS with PowerPath 28
- read ahead 82, 105
- read operation 82
- storage configuration 29
- stripe size 42, 49
- VNX for block 40, 54
- VNX for block volumes 42
- VNX OE with MPFS 29
- Windows servers 40, 54
- persistent connection 81
- PowerPath support 28
- PowerPath with MPFS 60
- Prefetch requirements 29
- properties
 - listing 107
 - retaining across restarts 81, 105
 - setting 105

R

- Read ahead performance 82
- Read cache requirements 29
- removing MPFS software, troubleshooting 122
- restarting MPFS software 81, 93
- reverse authentication 67, 69

S

- SAN switch zoning 55
- secret (CHAP) 69
- security file creation 56
- SendTargets discovery 69
- setting
 - arraycommpath 65, 74, 77
 - failovermode 65, 74, 77
- setting up
 - MPFS file system 42
 - VNX 41
- software components
 - Java Runtime Environment 37
 - MPFS software 37
 - NAS software 37

- Windows iSCSI Initiator 37
- starting an MPFS session using mpfsctl enable 105
- starting MPFS 53
- statistics
 - displaying 108
 - resetting counters 107
- stopping an MPFS session
 - through the Properties dialog box 97
 - using mpfsctl disable 104
- storage configuration
 - configuring 55
 - installation 38
 - recommendations 29
 - requirements 38, 40
 - stripe size 42
- storage group
 - adding FC hosts 65, 74
 - adding initiators 74
 - configuring 64
- storage guidelines 29
- synchronous flushing 81, 105
- system component verification 35

T

- troubleshooting 113
 - administrative rights 122
 - confirming MPFS software installation 121
 - fibre channel error messages 114
 - installing MPFS software 122
 - iSCSI error messages 114
 - locked files on server 122
 - mounting an MPFS file system 114
 - MPFS file system error messages 115
 - MPFS software 121
 - removing MPFS software 122
 - uninstalling earlier MPFS software version 95
 - uninstalling MPFS software 122
 - Windows Event Viewer 116

U

- uninstalling
 - MPFS software 94, 122
 - MPFS software, earlier version 95
 - MPFS software, troubleshooting 122

- unmounting
 - MPFS file system 29
 - MPFS software 29
- upgrading MPFS software 92

V

- verifying an MPFS software installation 121
- version number, displaying 109
- VMware ESX server
 - limitations 31
 - requirements with Windows 31
- VNX
 - Fibre Channel 19, 35
 - iSCSI 21, 35
 - iSCSI/FC 36
 - setup 41
- VNX for block
 - best practices 54
 - configuring using CLI commands 54
 - iSCSI port configuration 57
 - storage requirements 40
 - system requirements 40
- VNX for file
 - configuring 43

- enabling MPFS 53
- VNX MPFS
 - configurations 35
 - over FC 19
 - over iSCSI 21
 - over iSCSI/FC 22
- VNX VG2/VG8
 - Fibre Channel 20, 35
 - iSCSI 22, 23, 36
 - iSCSI/FC 36
- volume stripe size 28
- volumes mounted 48
- volumes, names of 48

W

- Windows Event Viewer 116
- Windows server
 - configuration 29, 32
 - installation 86
 - Windows OS versions 87

Z

- zone activation 55

