

EMC[®] PowerPath[®] Migration Enabler

Version 5.7

User Guide

P/N 300-014-263

REV 05



EMC²

Copyright © 2006 - 2014 EMC Corporation. All rights reserved. Published in the USA.

Published January 2014

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.

EMC², EMC, and the EMC logo are registered trademarks or trademarks of EMC Corporation in the United States and other countries. All other trademarks used herein are the property of their respective owners.

For the most up-to-date regulatory document for your product line, go to EMC Online Support (<https://support.emc.com>).

CONTENTS

Preface

Chapter 1

PowerPath Migration Enabler Overview

Migration Enabler overview.....	11
Supported migration scenarios.....	11
Migrating data onto an encrypted LUN	11
Migrations using pseudo and native devices.....	12
Installing and licensing Migration Enabler.....	12
Understanding the Migration Enabler technology types.....	13
Host Copy	13
Encapsulation.....	13
Open Replicator.....	14
TimeFinder/Clone	14
Understanding the Migration Enabler migration workflow	14
Migration state descriptions	16

Chapter 2

Migration Procedures

Before you start a migration.....	21
Configuring PowerPath Migration Enabler cluster resources	23
Perform a migration	24
After you finish a migration.....	25
Reclaim unused space on the target for plaintext data	25
Reclaim unused space on the target for encrypted text.....	28

Chapter 3

Migration Enabler Features

Migration Enabler commands	29
Managing simultaneous migration sessions using -all and -file	30
Changing the synchronization speed for migrations	31
Specifying a Host Copy upper limit copy rate.....	31
powermig options command.....	32
Migrating to and from virtually provisioned devices	32
Using the Solutions Enabler Thin Client for Migration Enabler	32
Changing the remote Solutions Enabler server.....	33
Viewing and updating disk labels with powerformat	34

Chapter 4	Volume Managers and Cluster Environments	
	Migrations in a volume manager environment.....	35
	Migrations in a Veritas Volume Manager (VxVM) environment..	35
	Migrations in a Logical Volume Manager (LVM) environment....	36
	Migrate a native-named device under VxVM control.....	36
	Migrate devices under LVM control on UNIX and Linux hosts....	37
	Migrations in a cluster environment.....	39
	Migrate in MSCS cluster environment	40
	Migrate physical disks.....	40
	Migrate physical disks on Windows Server 2003.....	42
	Migrate cluster shared volume.....	42
	Migrate in HACMP environment (AIX hosts).....	43
	Migrate between pseudo-named devices in VCS environment..	45
	Migrate between pseudo-named devices in Sun Cluster with VxVM	48
	Rename a pseudo device	49
Appendix A	Migration Procedure Examples	
	Host Copy migrations.....	51
	Host Copy procedure	51
	Host Copy throttle settings.....	53
	Host Copy ceiling settings.....	53
	Migrating to encrypted devices	54
	Encapsulation migrations	54
	Set up Storage Element for Migration Enabler	54
	Identify target-device name	55
	Encapsulation procedure	56
	Open Replicator migrations	58
	Set up Symmetrix-to-Symmetrix migrations with Open Replicator .	58
	Migrate EMC VNX or CLARiiON data to EMC Symmetrix arrays...	59
	Open Replicator procedure	60
	Understanding Open Replicator throttle settings.....	61
	TimeFinder/Clone migrations.....	61
	TimeFinder/Clone procedure.....	61
	Understanding TimeFinder/Clone throttle settings	63
Appendix B	Troubleshooting a Migration	
	Audit and error log messages	65
	Audit log message example	65
	Error log message example	65

- Reporting a problem 66
- Migration impact from a graceful host reboot..... 66
- Faults after a host crash..... 66
- Other reported faults 68
- Resolving a process hang 69
- Troubleshooting cluster migrations 69
 - Cluster migration best practices..... 70
 - Common terminologies..... 70

- Appendix C Support Summary**
 - Migration Enabler technology type support..... 73

- Appendix D Disk Labeling powerformat Command for Solaris**
 - Solaris and PowerPath Migration Enabler disk labeling 75
 - Multiple array type entries in the format utility output..... 75
 - Viewing and updating disk labels with powerformat 77

- Index**

Contents

PREFACE

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

Contact your EMC representative if a product does not function properly or does not function as described in this document.

Note: This document was accurate at publication time. New versions of this document might be released on the EMC Online Support website. Check the EMC Online Support website to ensure that you are using the latest version of this document.

Purpose

This document describes how to configure and use EMC PowerPath Migration Enabler.

Audience

This manual is part of the PowerPath documentation set. It is intended for use by storage administrators and other information system professionals responsible for using PowerPath Migration Enabler.

Readers of this guide are expected to be familiar with the operating system on which you are running PowerPath Migration Enabler and with EMC storage systems.

Related documentation

The following EMC publications provide additional information:

- ◆ *EMC PowerPath Installation and Administration Guide*
- ◆ *EMC PowerPath Encryption with RSA User Guide*
- ◆ *EMC PowerPath and PowerPath /VE for Windows Installation and Administration Guide*
- ◆ *EMC PowerPath Family CLI and System Messages Reference*
- ◆ *EMC PowerPath Family Product Guide*

- ◆ *EMC PowerPath Management Pack for Windows User Guide*

Revision history

The following table presents the revision history of this document.

Table 1 Revision history

Revision	Date	Description
05	January 2014	Release of PowerPath for Windows 5.7 SP2 <ul style="list-style-type: none"> • Added “Migrate physical disks on Windows Server 2003” on page 42 beginning with PowerPath 5.7 SP2 for Windows.
04	June 2013	Addition of a new MSCS cluster functionality with PowerPath Migration Enabler beginning with PowerPath 5.7 SP1 for Windows.
03	February 2013	Release of EMC PowerPath Migration Enabler Version 5.7 for Windows.
02	December 2012	Update to the PowerPath Migration Enabler license beginning with PowerPath 5.7 SP1 for Linux.
01	September 2012	First release of the EMC PowerPath Migration Enabler Version 5.7 User Guide for Linux.

Conventions used in this document

EMC uses the following conventions for special notices:

NOTICE

NOTICE is used to address practices not related to personal injury.

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

IMPORTANT

An important notice contains information essential to software or hardware operation.

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, and utilities URLs, pathnames, filenames, directory names, computer names, links, groups, service keys, file systems, and notifications
Bold	Used in running (nonprocedural) text for names of commands, daemons, options, programs, processes, services, applications, utilities, kernels, notifications, system calls, and 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 the 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
<code>Courier bold</code>	Used for specific user input, such as commands
<i><code>Courier italic</code></i>	Used in procedures for: <ul style="list-style-type: none"> Variables on the 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 enclose content that the user must specify, such as x or y or z
...	Ellipses indicate nonessential information omitted from the example

Where to get help

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

Product information — For documentation, release notes, software updates, or information about EMC products, licensing, and service, go to the EMC Online Support website (registration required) at:

<https://support.emc.com>

Technical support — For technical support, go to EMC Online Support and select Service Center. On the Service Center page, you will see several options, including one to create a service request. Note that to open a service request, you must have a valid support agreement. Contact your EMC sales representative for details about obtaining a valid support agreement or with questions about your account.

Your comments

Your suggestions will help us continue to improve the accuracy, organization, and overall quality of the user publications. Send your opinions of this document to:

techpubcomments@emc.com

CHAPTER 1

PowerPath Migration Enabler Overview

This chapter provides an overview of the EMC PowerPath Migration Enabler product, discusses the various technology types and migration states, and provides licensing information for migration enabler:

- ◆ [Migration Enabler overview](#) 11
- ◆ [Installing and licensing Migration Enabler](#) 12
- ◆ [Understanding the Migration Enabler technology types](#) 13
- ◆ [Understanding the Migration Enabler migration workflow](#) 14

Migration Enabler overview

EMC® PowerPath® Migration Enabler is a migration tool that enables non-disruptive or minimally disruptive data migration between storage systems or between logical units within a single storage system. Migration Enabler resides on the host and allows applications continued data throughout the migration process. Migration Enabler works independently of PowerPath multipathing. However, PowerPath must be installed.

PowerPath Migration Enabler integrates with other technologies to minimize or eliminate application downtime while migrating data.

Supported migration scenarios

The *EMC PowerPath Family Release Notes* includes supported migration scenarios for each operating system. The release notes include information about supported pseudo and native device names, technology types, data types (encrypted and plaintext), and operating systems.

Migrating data onto an encrypted LUN

The Migration Enabler Host Copy technology migrates data onto an encrypted LUN. The *PowerPath Encryption with RSA User Guide* provides additional information.

Migrations using pseudo and native devices

PowerPath Migration Enabler supports migrations using PowerPath devices (pseudo devices) or native operating system devices.

Migrations are non-disruptive with pseudo devices, which means an application never has to be brought down through the entire migration process. With native devices, you must manually configure the application after the migration to allow the host to see the new target device. All migrations on HP-UX are disruptive migrations.

Installing and licensing Migration Enabler

On all PowerPath supported operating systems, PowerPath Migration Enabler automatically installs when PowerPath multipathing is installed. Run **powermt check_registration** command to confirm your license registration.

With the exception of the PowerPath 5.7 for Linux GA release version, the PowerPath multipathing license includes all Migration Enabler functionalities including all Migration Enabler technologies (Host Copy, Open Replicator, TimeFinder/Clone, and Encapsulation). PowerPath Migration Enabler license keys are available at no cost to new and existing PowerPath Multipathing customers.

NOTICE

The PowerPath Encryption with RSA license supports Host Copy data migration to or from encrypted devices without the need to install the PowerPath multipathing license. The PowerPath multipathing license is required when migrating with only plaintext devices (no encrypted devices as source or target).

You do not need to obtain a separate license for every server. A single license supports multiple PowerPath migrations.

Note: Multiple PowerPath migrations are supported using different technologies on the same host.

License key ordering and activation

Effective February 15, 2011, instead of the physical Right To Use (RTU), the default delivery method for PowerPath licenses is electronic. An electronic License Authorization Code (LAC) is sent by email in order to redeem the license key on the Powerlink Licensing Portal. This does not affect upgrades because PowerPath retains existing license information.

Physical RTU cards are still available as an option. EMC Global Support, at 1-800-svc4emc or 1-800-782-4362, can provide more information. The *EMC PowerPath Family Electronic License Ordering Process Technical Notes* provides more information about the PowerPath license electronic ordering process.

Understanding the Migration Enabler technology types

NOTICE

Migration Enabler does not support migrations with boot devices for any Migration Enabler technology.

Host Copy

PowerPath Migration Enabler Host Copy is a host-based migration tool that uses host resources to move data nondisruptively from a source to a target. Host Copy supports the following:

- ◆ Migrations to and from a device managed by PowerPath (boot and swap devices are not supported)
- ◆ Migrations to and from virtually provisioned devices (thin devices)
- ◆ Migrations to and from encrypted LUNs

Encapsulation

PowerPath Migration Enabler Encapsulation supports migration of EMC Symmetrix[®], EMC VNX[™], and EMC CLARiiON[®] devices to EMC Invista[®] devices. The target logical unit in the migration is a Virtual Volume that encapsulates a Symmetrix, VNX, or CLARiiON Storage Element. Because both logical units represent the same physical storage, no data moves during the migration.

Open Replicator

EMC Open Replicator for EMC Symmetrix works along with PowerPath Migration Enabler and with Solutions Enabler and supports the online pull method where data is copied through the SAN from a remote logical unit to a Symmetrix storage system.

Note: Source from EMC Symmetrix, VNX, or CLARiiON cannot be in the same array.

TimeFinder/Clone

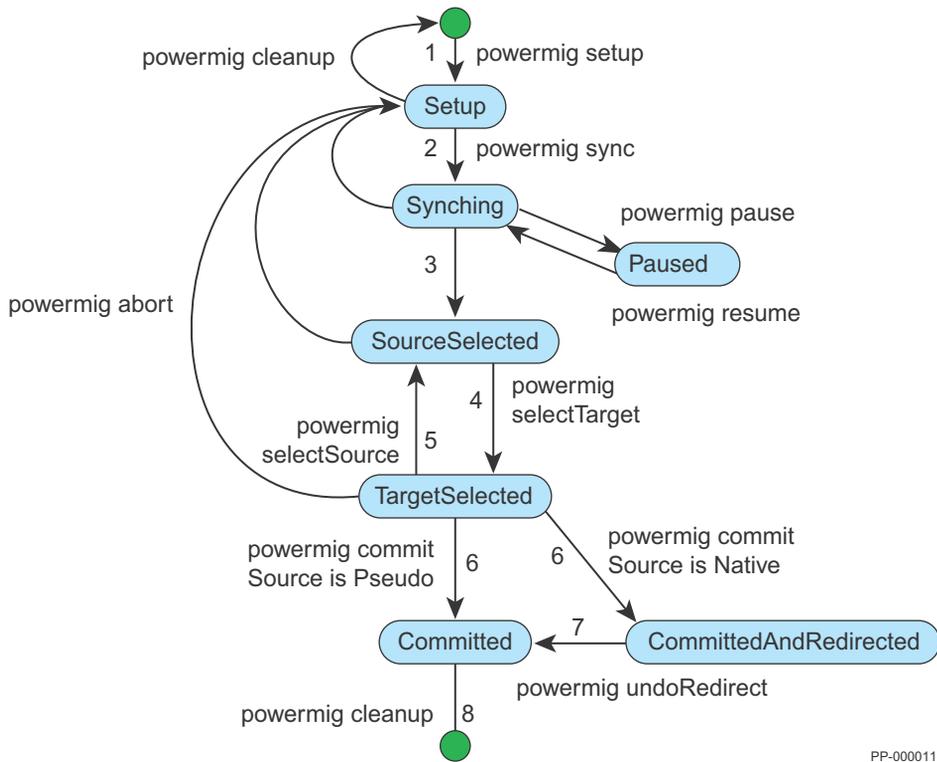
EMC TimeFinder®/Clone migrations work along with PowerPath Migration Enabler and with Solutions Enabler to migrate data between devices within a single EMC Symmetrix system, including devices with different RAID type.

Note: The source and target devices must be from the same EMC Symmetrix array.

Understanding the Migration Enabler migration workflow

A migration session transitions through several states as you complete each step in the migration process.

Figure 1 on page 15 shows the overall migration workflow, including the sequence of **powermig** commands and migration states. Numbers 1–8 in the graphic display the commands involved in a normal workflow.



PP-000011

Figure 1 Migration commands and states

Migration state descriptions

Table 2 describes the migration states.

Table 2 Migration state descriptions

State	Description
Setup	A migration is in the Setup state after a powermig setup command completes successfully. Before the migration can enter this state, synchronization prerequisites must be met, including technology-specific criteria, where relevant.
Syncing	A migration enters the Syncing state after running the powermig sync command. Bulk data copy initiates in the Syncing state. Data is read from the source and is written to the target. Application reads are directed to the source. Application writes are written to the source and cloned to the target.
syncing [PAUSED]	A migration enters the syncing [PAUSED] state after running the powermig pause command. Bulk data copy and synchronization of the source and target logical units are paused. Run powermig resume to restart a migration that was paused.
SourceSelected	The migration enters the SourceSelected state once the bulk data copy of the source and target logical units completes. While the migration is in the SourceSelected state, the source and target logical units continue to synchronize, and I/O read requests are processed from the source logical unit.
TargetSelected ¹	The migration transitions to the TargetSelected state after the powermig selectTarget command completes. While the migration is in the TargetSelected state, the source and target logical units continue to synchronize, and I/O read requests are processed from the target logical unit.
CommittedAndRedirected	If you are migrating data from a native-named device, this state occurs after you issue the powermig commit command. While in this state, the application using the native source device must be brought down and the user must run the powermig undoRedirect command. The application can then be reconfigured and brought back up to use the target device. The migration will now be in the Committed state.
Committed	If using a pseudo source device, this state is entered upon issuing the powermig commit command. During this transition, Migration Enabler swaps all underlying paths on both source and target logical units causing all I/O to be sent to the target logical unit without an application having to be brought down or reconfigured. If using a native source device, this state is entered after issuing the powermig undoRedirect command (as described above in the CommittedAndRedirected state). In the Committed state, source and target devices are no longer kept synchronized and I/O only goes to the target device.

1. TargetSelected state is not supported while migrating physical disks in an MSCS cluster environment. You need to skip the targetSelected state from the migration workflow and commit directly from the sourceSelected state. On device fault, migrations are always aborted and never committed.

Setup migration state

After the prerequisites are met, run the **powermig setup** command to enter the **Setup** state. You can perform the following actions:

- ◆ Specify the source and target device names.

NOTICE

Do not rename a logical unit that is involved in a migration session.

Note: While migrating physical disks in Microsoft Cluster Server (MSCS) environment, the source should be a cluster disk with a corresponding physical disk resource in the cluster and the target device should be a non-cluster disk.

- ◆ Specify the technology type.
- ◆ Optionally set the throttle value to configure the migration speed of the synchronization.
- ◆ Optionally specify the remote host or the port for Solutions Enabler (Open Replicator and TimeFinder/Clone migrations only).

To successfully enter the **Setup** state, PowerPath Migration Enabler performs the following actions:

- ◆ Checks devices for availability
- ◆ Checks that the devices are not part of another migration
- ◆ Checks that the target is large enough for the migration

Note: During **Setup**, the target is made unavailable to all applications on the host for the duration of the migration.

Synchronization migration state

After the Setup state, run the **powermig sync** command to enter the **Syncing** state. You must specify the handle (returned from Setup) for the migration pair. In the **Syncing** state, host applications continue to the source logical unit, as data is copied to the

target. Migration Enabler also ensures all new host writes are written to both source and target devices. The time for completing synchronization depends in part on the size of the source logical unit.

While in the **Syncing** state, you can perform the following options:

- ◆ Query one or more migrations to determine percentage completion and throttle value.
- ◆ Increase or decrease the migration synchronization rate by changing the throttle value.
- ◆ Control the speed of all Host Copy migration synchronizations by setting the Host Copy ceiling (in MB/second).
- ◆ Pause the migration and then resume at a later time.
- ◆ Abort the migration, returning it to the **Setup** state.

Once in the **Syncing** state, PowerPath Migration Enabler performs the following actions:

- ◆ Initiates write cloning of application data. Any new I/O writes to the source device are automatically cloned and written to the target device.
- ◆ Initiates bulk data copy from source device to target device.

SourceSelected migration state

After all data has been copied from source to target, the migration enters the **SourceSelected** state where the following occurs:

- ◆ I/O reads are directed to the source.
- ◆ I/O writes are written to the source and target devices to keep them synchronized.

While in the **SourceSelected** state, you can perform the following options:

- ◆ Transition to the **TargetSelected** state.

Note: While migrating physical disks in an MSCS environment, you need to skip the **TargetSelected** state from the migration workflow since migrations are committed directly from the **SourceSelected** state.

- ◆ Abort the migration and return the migration to the **Setup** state.

TargetSelected migration state

When you run the **powermig selectTarget** command, the migration transitions to the **TargetSelected** state from the **SourceSelected** state. The **TargetSelected** state must be entered before committing the migration. In this state, all I/O reads are sent to the target while writes continue to be written to both source and target. This gives you an opportunity to determine that all data is intact on the target and service levels are in line with application requirements.

Note: **TargetSelected** state is not supported while migrating physical disks in an MSCS environment. You need to skip the **targetSelected** state from the migration workflow since migrations are committed directly from the **sourceSelected** state. On device fault, migrations are always aborted and never committed.

While in the **TargetSelected** state, you can perform the following operations:

- ◆ Move to the **Committed** state, (or **CommittedAndRedirected** state if migrating a native source device).
- ◆ Abort the migration and return the migration to the **Setup** state.
- ◆ Return to the **SourceSelected** state.

Once in the **TargetSelected** state, PowerPath Migration Enabler behaves as follows:

- ◆ I/O reads are directed to the target.
- ◆ I/O writes are written to the source and target.

Committed migration state

The Committed state permanently designates the target as the recipient of all I/O requests. After this command is run, PowerPath Migration Enabler no longer keeps the source and target synchronized, and I/O is no longer sent to the source.

After entering the **Committed** state, perform the following action:

Issue the **Cleanup** command.

To successfully enter the **Committed** state, PowerPath Migration Enabler performs the following actions:

- ◆ When committing pseudo devices, source and target LUNs are swapped, allowing an application to continue using the original source pseudo device without any disruption.

- ◆ When committing native devices, the migration enters the **CommittedAndRedirected** state. In this state you must take down any application using the source device and execute the **undoRedirect** command, which stops the redirection of all I/O from source to target. At this point the application can be reconfigured to use the new target device and brought back online.



Once in the Committed state, there is no returning to the previous state.

- ◆ All reads and writes are directed to the target.
- ◆ The source and target are no longer synchronized.

Cleanup migration state

Cleanup is the final step in the migration process, where the following occurs:

- ◆ When entered from the **Committed** state, the **powermig cleanup** command removes enough data from the former **source** device to prevent the existence of two identical logical units, ensuring that the operating system or applications do not become confused.

Note:

—For encapsulation migrations, remove source to the host before cleanup.
—On some operating systems, **powermig cleanup -format** option can take substantially longer since it performs a full format of the disk.

- ◆ When entered from the **Setup** state (after setup or abort) this command removes enough data from the **target** device to prevent the possibility of two identical logical units.
- ◆ The migration handle is removed and the migration no longer exists.

CHAPTER 2

Migration Procedures

This chapter provides the steps for performing the various PowerPath technology type migrations, and includes the migration prerequisites and post-migration steps:

- ◆ [Before you start a migration](#) 21
- ◆ [Configuring PowerPath Migration Enabler cluster resources](#) 23
- ◆ [Perform a migration](#)..... 24
- ◆ [After you finish a migration](#) 25
- ◆ [Reclaim unused space on the target for plaintext data](#) 25
- ◆ [Reclaim unused space on the target for encrypted text](#) 28

Before you start a migration

The following prerequisites apply to all the Migration Enabler technology types:

Note: Some prerequisites may only apply to specific technology types, as noted.

- ◆ Ensure that the PowerPath multipathing license is installed.
- ◆ Make target devices visible to the host, such as by provisioning Symmetrix logical units or adding volumes to the storage group.
- ◆ The target device must be the same size or larger than the source device and cannot be in use by any application.
- ◆ For PowerPath Encryption with RSA, the target logical unit must be at least 65 KB (Solaris or Windows) or 68 KB (Linux) larger than the source when migrating from an unencrypted to an encrypted virtual logical unit to allow space for PowerPath Encryption metadata.
- ◆ If using the Solutions Enabler remote-only client with Open Replicator or Timefinder/Clone technologies, download and install the Solutions Enabler Thin Client from EMC Online Support.
- ◆ With replication, such as EMC MirrorView™ or Symmetrix Remote Data Facility® (SRDF), if you replicate your source to a remote site, you must also replicate your target to a remote site.

- ◆ Use your HBA management utility to change the maximum transfer length on the target device to be equal to or larger than the source device. If the target maximum transfer length is smaller than the source device an error may display. Refer to your HBA manufacturers' documentation for more information on changing the maximum transfer length.
- ◆ The following applies to Windows:
 - Ideally, the target logical unit is a freshly initialized MBR partition on a basic disk (or a GPT basic disk if the platform supports it). Regardless of the source type, the target logical unit should not contain a volume or file system.
 - Ensure that the Startup type is set to enabled to run PowerPath Migration Enabler. Turn on Services if you turned them off.
- ◆ The following applies when you migrate physical disks in an MSCS cluster environment (Windows platform):
 - Ensure that PowerPath Migration Enabler cluster resource is configured for source cluster disk. [“Configuring PowerPath Migration Enabler cluster resources” on page 23](#) provides the required steps.
 - Ensure that the same version of PowerPath is installed on all nodes in the cluster.
 - Ensure that PowerPath Migration Enabler feature is installed on all nodes in the cluster.
 - Ensure that all nodes in the cluster are up and running.
- ◆ The following applies to Solaris:
 - Before using the **powermig setup** command, both the source and target must be labeled, and the labels must be of the same type (VTOC/SMI or EFI).
- ◆ The following applies to Open Replicator:
 - The target logical unit must reside on a Symmetrix system with Engenuity Operating Environment 5x71 or later.
 - The two storage systems where the source and target reside must be connected directly or through a switch and they must be able to communicate.
 - Every port on the target array that allows access to the target logical unit must be accessible to the source logical unit through at least one port on the source array.

- ◆ The following applies to TimeFinder/Clone and Open Replicator
 - Logical units involved in a migration should not be used as Symmetrix gatekeeper devices. Add the names of devices you want to use as gatekeepers to the gatekeeper selection file, **gkselect**. Even if the device names are already configured as gatekeepers on the Symmetrix system, add them to the **gkselect** file to avoid a potential conflict between SYMAPI and Migration Enabler. The *EMC Solutions Enabler Symmetrix Base Management CLI Product Guide* contains more information about gatekeeper devices.

Related topics

[“Volume Managers and Cluster Environments” on page 35](#)

[“Migrate in MSCS cluster environment” on page 40](#)

Configuring PowerPath Migration Enabler cluster resources

Use the **powermigcl config** command to configure PowerPath Migration Enabler cluster resources for physical disks. During the configuration, a new PowerPath Migration Enabler cluster resource is created for each cluster physical disk resource and is added to its dependencies.

1. Open a command prompt window.
2. Run the **powermigcl config** command and perform any of the following operations depending on your requirement:
 - *powermigcl config -all -force*: Configures all available cluster groups on all nodes.
 - *powermigcl config -all*: Configures all available cluster groups on the current node.
 - *powermigcl config -group <group1>*: Configures the specified cluster group on the current node.

The *PowerPath Family CLI and System Messages Reference Guide* available on EMC Online Support provides more information on the **powermigcl** command, its syntax, and the associated arguments.

Perform a migration

1. Run the following command to set up the migration from the source to the target, and to specify the technology:

```
powermig setup -techType <techType> -src <sourceName> -tgt
<targetName>|-file <fileName> [-throttleValue
<throttleValue>] [-cluster] [-techHost <hostname> [-techPort <port#>]]
[-nothin] [-force] [-noprompt|-no]
```

Note: The **-cluster** argument is used only when migrating devices in an MSCS environment. This argument is not supported on Windows Server 2003.

2. Run the following command to start the bulk copy and cloning of application write I/Os:

```
powermig sync -handle <migrationHandle>|-all [-noPrompt|-no]
```

3. Run the following command to check the status of the migration:

```
powermig query -handle <migrationHandle>|-all [-noPrompt|-no]
```

4. Run the following command to designate the target logical unit as the recipient of all I/O read requests:

```
powermig selectTarget -handle <migrationHandle>|-all [-noPrompt|-no]
```

- Skip this step while migrating physical disks in an MSCS cluster environment.

5. Run the following command to commit the migration to designate the target as the recipient of all I/O requests:

```
powermig commit -handle <migrationHandle>|-all [-noPrompt|-no]
```

6. If you are migrating native devices, run the following command to stop the redirection of I/O from the source to the target:

```
powermig undoRedirect -handle <migrationHandle>|-all [-force]
[-noPrompt|-no]
```

7. Run the following command to clean up the migration. The migration ends after the **cleanup** step.

```
powermig cleanup -handle <migrationHandle>|-all
[-format] [-force] [-noPrompt|-no]
```

Related topics

[“After you finish a migration” on page 25](#)

[Appendix A, “Migration Procedure Examples”](#)

After you finish a migration

Complete the post migration steps that apply to your environment:

- ◆ If the target is larger than the source, follow the steps in [“Reclaim unused space on the target for plaintext data” on page 25](#) or [“Reclaim unused space on the target for encrypted text” on page 28](#) to make unused space accessible.
- ◆ If you have EMC ControlCenter[®], update it to recognize the new device name mappings after a migration. The EMC ControlCenter documentation provides more information.
- ◆ Update any replication software to ensure that the new target device is replicated. This could also be done prior to the syncing state.

Reclaim unused space on the target for plaintext data

After a migration completes, the usable space on the target is equivalent to the size of the source. When the target logical unit is larger than the source, operating system commands must be run to make the unused space accessible.

Note: This section provides sample procedures, but it does not include steps for all supported operating systems and versions.

AIX hosts

1. Run **chvg -g** *<volume_group>* to increase the size of the volume. Where *<volume_group>* is the volume group name.
2. If a logical volume exists, run **chlv -x** *<MaxLPs>* *<logical_volume>* to increase the size of the logical volume. Where:
 - *<MaxLPs>* must be between **1** and **32512**. This value cannot be smaller than the existing number of logical partitions (LPs).
 - *<logical_volume>* is the logical volume name.

3. Run `chfs -a size=+<new_size> |<mount-point>` to increase the size of the file system.

Linux hosts

1. Complete the migration procedure as described under [“Perform a migration” on page 24](#).
2. Unmount the file system.
3. Run `vgchange -an <volume_group>` to make the volume group unavailable.
where `<volume_group>` is the volume group name.
4. Run `pvresize /dev/emcpowerx` to resize the physical volumes.
where `x` is the name of the target logical unit after the migration.
5. Run `vgchange --resizeable` to make the volume group resizeable if it is not already.
6. Run `vgdisplay <volume_group>` to verify how many physical extents are free.
7. Run `vgchange -ay <volume_group>` to make the volume group available.
8. Run `lvextend -l <phys_ext> /dev/<volume_group>/<log_vol#>` to extend the logical volume.

where:

`<phys_ext>` is the number of physical extents you want to add

`<log_vol#>` is the logical volume number

For example, `lvextend -l +6343 /dev/vg01/lv01`

9. Resize the file system. For example, for ext3 file systems, type:

```
e2fsck -f /dev/<volume_group>/<log_vol#>  
resize2fs -p /dev/<volume_group>/<log_vol#>
```

10. Mount the file system.

Solaris hosts

SMI labeled devices

For Solaris hosts, when both source and target devices are SMI labeled, to reclaim unused space on the original target device after migration and cleanup, perform the following steps:

1. Type `umount <mount-point>` to unmount the file system if already mounted.

2. Type **powerformat -w -x** *<target_raw_device>* to use PowerPath powerformat utility to reclaim unused space on the target device.
3. Use the Solaris **format** utility to verify that the new label is applied (use **verify** option).
4. Quit format.
5. Grow the file system. Type **growfs** [-M mount-point] [newfs-options] **raw-special-device**

EFI labeled devices

For Solaris hosts, where both the source and target are EFI labeled, to reclaim unused space on the original target device after migration and cleanup, perform the following steps:

1. Type **umount** *<mount-point>* to unmount the file system if already mounted.
2. Type **format -e** *<target_device>* to relabel the target device.
3. Select **type**.
4. Choose **0: Auto Configuration**.
5. Select **label** and choose **1** for EFI label.
6. Quit format.
7. Grow the file system. Type **growfs** [-M mount-point] [newfs-options] **raw-special-device**

Windows hosts

DiskPart enables administrators to manage objects (disks, partitions, or volumes) by using scripts or direct input at a command prompt. To claim space on the target:

1. At a Windows command prompt, type **diskpart** to start DiskPart.
2. Type **list volume** to list all volumes seen by this host.
3. Type **select volume** *<volume>* to select the specific volume to shift the focus to it.
4. Type **extend** to extend the volume with focus into the next contiguous unallocated space.
5. Exit DiskPart.

Reclaim unused space on the target for encrypted text

AIX hosts

1. Run **powervt update -dev** *<device_name>*.
2. Run **powervt resize -dev** *<device_name>*.
3. Run **chvg -g** *<volume_group>* to increase the size of the volume.
4. If a logical volume exists, run **chlv -x** *<MaxLPs>* *<logical_volume>* to increase the size of the logical volume. Where:
 - *<MaxLPs>* must be between **1** and **32512**. This value cannot be smaller than the existing number of logical partitions (LPs).
 - *<logical_volume>* is the logical volume name.
5. Run **chfs -a size=+<new_size> /<mount-point>** to increase the size of the file system.

Linux hosts

Follow the procedure in [“Linux hosts” on page 26](#) to reclaim unused space on the target for Linux hosts.

Solaris and Windows hosts

1. Run **powervt update** to move and re-read the PowerPath encryption metadata on the disk to another location when the LUN has grown.
2. Run **powervt resize** to adjust the usable space on an encrypted virtual logical unit to reflect its actual capacity.
3. Perform one of the following:
 - On Solaris hosts, grow the file system:
 - a. Type **growfs [-M mount-point] [newfs-options] raw-special-device**
 - b. Use the Solaris format **utility** to verify that the new label is applied (use **verify** option).
 - On Windows hosts, if there is volume build on top of this partition and file system, then follow the steps under [“Windows hosts” on page 27](#) to utilize new capacity available on the virtual encrypted disk.

CHAPTER 3

Migration Enabler Features

This chapter provides information about the following topics:

- ◆ Migration Enabler commands 29
- ◆ Managing simultaneous migration sessions using -all and -file..... 30
- ◆ Changing the synchronization speed for migrations..... 31
- ◆ Specifying a Host Copy upper limit copy rate 31
- ◆ Migrating to and from virtually provisioned devices 32
- ◆ Using the Solutions Enabler Thin Client for Migration Enabler 32
- ◆ Viewing and updating disk labels with powerformat 34

Migration Enabler commands

PowerPath Migration Enabler uses a common set of commands to support Host Copy, Encapsulation, Open Replicator, and TimeFinder/Clone migration technology types.

[Table 3 on page 29](#) lists the migration **powermig** commands.

Note: Some commands may vary in usage and meaning depending on technology type. The *PowerPath Family CLI and System Messages Reference Guide* describes how these commands are used with the various technology types.

Table 3 PowerPath Migration Enabler commands (1 of 2)

powermig command	Description
abort	Aborts a migration and returns the migration session to Setup state.
cleanup	Cleans up data on the source or target logical unit, depending on the state of the migration when you run the command.
commit	Commits the migration by permanently designating the target as the recipient of all I/O requests.
getHandle	Displays the handle for a migration session in progress.
help	Displays a brief description and usage summary for the specified powermig command.
info	Displays information about a specific migration session or about all active migrations.

Table 3 PowerPath Migration Enabler commands (2 of 2)

powermig command	Description
options	Used to set values that are global to all migrations, or all migrations that use a specific technology. For example, the option, hostcopy_ceiling , can be used to specify an upper limit on the aggregate rate of copying for all host copy migrations. When executed with no option specified, powermig options displays all available options and their values. Currently, hostcopy_ceiling is the only available option.
pause	Suspends a migration in the Syncing state and pauses the synchronization of the source and target logical units.
query	Displays the state of a migration.
recover	Recovers the migration command that was in progress when an interruption occurred due to a migration error or process crash.
resume	Restarts a migration that was paused by the powermig pause command.
selectSource	Designates the source logical unit as the recipient of all I/O Read requests. Continues to send writes to both source and target.
selectTarget ¹	Designates the target logical unit as the recipient of all I/O Read requests. Continues to send writes to both source and target.
setup	Specifies the source and target devices and technology type used in the migration.
sync	Starts the bulk data copy and write cloning process while application continues uninterrupted.
throttle	Sets the speed at which data is synchronized between a source-and-target pair.
undoRedirect	Stops the redirection of I/O from the source to the target.
version	Displays the version of PowerPath Migration Enabler running on the host.

1. selectTarget is not supported while migrating physical disks in an MSCS cluster environment.

Managing simultaneous migration sessions using -all and -file

The **-all** and **-file** options allow you to manage simultaneous migration sessions. You can use a file of migration pairs with the **powermig setup** command (**-file** option) or manage all migration sessions that are in the state required for that **powermig** command (**-all** option).

The **-all** option performs the action for all migration sessions in the valid state for the command. For example, if you want to pause a migration and have migrations in various states (such as, Syncing and Setup), then **powermig pause -all** runs on migrations only in the valid state (Syncing). Migrations not in the valid state will be ignored (Setup). If there are no migrations in the valid state, then an error displays. After the command has run for all relevant handles, **powermig info -all -query** runs to display resulting state changes.

The *EMC PowerPath Family CLI and System Messages Reference Guide* provides more information.

Changing the synchronization speed for migrations

The **powermig throttle** command changes the synchronization speed for a migration, and can be set to increase or decrease the rate at which bulk data is copying. As throttle value increases, the bulk data copy rate decreases. The command works differently for the supported technology types, as described in this section. The *EMC PowerPath Migration Enabler Host Copy* white paper available on EMC Online Support provides additional details about throttle impact. This command does not apply to Encapsulation migrations.

[“Specifying a Host Copy upper limit copy rate” on page 31](#) discusses an option for setting the ceiling rate, which ignores **powermig throttle**.

Note: Throttle can also be set during **powermig setup**, and can be changed with **powermig throttle** while the migration is in the **setup** or **syncing** state, except in the case of OR, where throttle can be set when you run **powermig setup**, but cannot be changed until after running **powermig sync**.

The *EMC PowerPath Family CLI and System Messages Reference Guide* provides more information.

Specifying a Host Copy upper limit copy rate

The Host Copy ceiling option lets you specify an upper limit on the aggregate rate of copying for all host copy migrations in the **Syncing** state on a host. In turn, this limits the I/O for all host copy migrations. With Host Copy ceiling enabled, per-migration throttle values are ignored.

powermig options command

The **powermig options** command lets you set PowerPath Migration Enabler options which apply to all migrations using the Host Copy technology type. Running **powermig options** with no option merely displays all available options and their values.

Note: Host Copy ceiling value is not set on a per migration basis. Once set, the value is applicable for all the host copy migrations on the server. The behavior is observed during **syncing** state. Therefore it is preferable to set the ceiling value before starting the host copy migrations.

The *EMC PowerPath Family CLI and System Messages Reference Guide* provides more information.

Migrating to and from virtually provisioned devices

Migrations to and from virtually provisioned devices (thin devices) offer improved storage capacity utilization and simplified storage management. Virtual provisioning presents the application with sufficient capacity for an extended period of time, reducing the need to provision new storage frequently, and avoiding costly allocated but unused storage.

In Host Copy migrations, the **Setup** state identifies the target devices as either thick or thin. When copy operations begin with thin target devices, PowerPath Migration Enabler copies only those blocks with data.

Host Copy operates in thin copying mode only if it recognizes the target device as a virtually provisioned device. It checks an array-specific SCSI-mode page value to determine EMC Symmetrix or VNX/CLARiiON virtual LUNs. Host Copy cannot determine whether or not a LUN is thin on a third-party array.

Using the Solutions Enabler Thin Client for Migration Enabler

Migration Enabler supports Solutions Enabler (SE) Thin Client that is installed on the same host as Migration Enabler. The SE Thin Client communicates with a remote server running a full version of SE. The remote server can be a dedicated SYMAPI server, as shown in [Figure 2 on page 33](#). The SE Thin Client allows you to limit the number of servers running a full version of SE.

The SE Thin Client uses a reduced API and communicates with a SYMAPI daemon on the remote server. The SE Thin Client host is unable to directly access a base daemon, Symmetrix devices, or **symmwin**. Additionally, the Thin Client is restricted to the Migration Enabler API, as SYMCLI is unavailable:

Note: The Solutions Enabler Thin Client for PowerPath Migration Enabler applies only to technology types that require Solutions Enabler (that is, Open Replicator and TimeFinder/Clone).

For remote-only installations using the SE Thin Client, download the latest SE Thin Client version from EMC Online Support.

Note: Both the host where the SE Thin Client is installed and the server running the full SE version must be able to see the source and target devices.

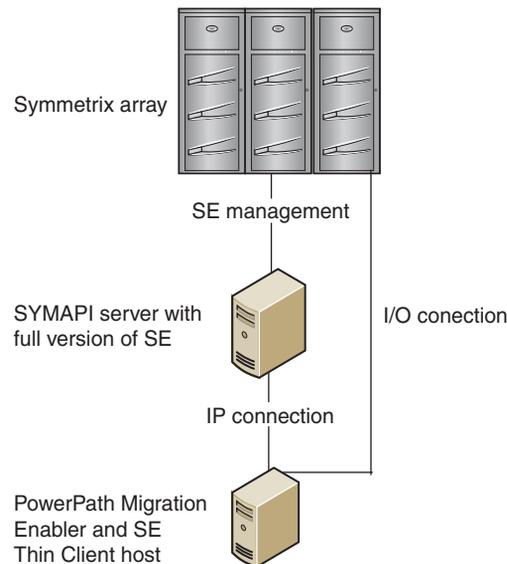


Figure 2 Solutions Enabler Thin Client for Migration Enabler

Changing the remote Solutions Enabler server

In rare cases if the remote server (or port) changes after a migration has started, two environment variables are provided to allow updating values that were set during **powermig setup**. The variables are:

```
PPME_REMOTE_SE_SERVER  
PPME_REMOTE_SE_PORT
```

- ◆ These variables can only be used with post-setup **powermig** commands (that is, they cannot be used with setup).
- ◆ An error is returned if these variables are present during **powermig setup**.
- ◆ Once used, values are saved and used for the duration of that migration (variables can and should be removed).

Viewing and updating disk labels with powerformat

For Solaris, PowerPath Migration Enabler includes a utility called **powerformat** that allows you to safely update disk-label information, preserve partition definitions and data, and make newly available disk capacity available for use.

When you migrate data to a target logical unit that is larger than the source logical unit, **powerformat** can adjust the disk label and partition table to make available the entire disk capacity on the target logical unit. The **powerformat** command takes partitions that extended to the end of the currently labeled logical unit (the label migrated from the source logical unit), and extends them to the end of the target logical unit. This allows you to use the additional space on the target.

Related topic

[Appendix D, “Disk Labeling powerformat Command for Solaris”](#)

CHAPTER 4

Volume Managers and Cluster Environments

This chapter discusses support with volume managers and migrations in cluster environments:

- ◆ [Migrations in a volume manager environment](#) 35
- ◆ [Migrations in a cluster environment](#) 39

Migrations in a volume manager environment

Currently, Veritas Volume Manager (VxVM) and Logical Volume Manager (LVM) are supported with PowerPath Migration Enabler.

Note: The target logical unit in a migration cannot be controlled by a volume manager. For example, when migrating Logical Disk Manager (LDM) or VxVM logical units on a Windows host, the target logical unit should not be a member of an LDM dynamic disk group or a VxVM disk group.

Migrations in a Veritas Volume Manager (VxVM) environment

When setting up a migration, with the target device online, if VxVM holds open the device, then PowerPath Migration Enabler reports a target device in use error message. You can resolve this issue in one of the following ways:

- ◆ Remove the device from VxVM control
- ◆ Run the **-force** option
- ◆ Run **fuser** to discover which process is holding open the device (UNIX only)

Linux

- ◆ The target logical unit should be removed from VxVM configuration. Run **vxdisk rm** *<target_device>* to remove from VxVM configuration.
- ◆ For native-named device migrations, refer to [“Migrate a native-named device under VxVM control” on page 36](#).

Migrations in a Logical Volume Manager (LVM) environment

The Logical Volume Manager included with AIX, HP-UX, and Linux is supported with PowerPath Migration Enabler. Refer to [“Migrate devices under LVM control on UNIX and Linux hosts” on page 37](#).

Migrate a native-named device under VxVM control

Note: PowerPath for AIX does not support native devices.

Migrating a native-named device under Veritas Volume Manager (VxVM) control requires additional steps, as described in the procedure below. Follow the standard migration procedure when migrating a pseudo-named device:

1. Follow the migration steps in [Chapter 2, “Migration Procedures,”](#) until committing the migration. Choose a convenient time to take applications or file systems offline and reconfigure them.
2. After committing the migration, close the application and unmount the volumes in the disk group.
3. Type **vx dg deport** *<disk_group_name>*.
4. Run the **powermig undoRedirect** command to turn off I/O redirection.
5. Import the disk group into VxVM, and restart the volumes.
 - a. Type **vx dg import** *<disk_group_name>*
 - b. Type **vx vol -g** *<disk_group_name>* **startall**
6. Mount the volumes.
7. Run the **powermig cleanup** command to clean up the migration.

Note: During VNX/CLARiiON to VNX/CLARiiON Host Copy migration with I/O running on iSCSI devices, **powermig cleanup -format** takes substantially longer since it performs a full format of the disk.

Optionally, use **vx disk resize** to resize the logical unit that is under VxVM control. If dealing with an encrypted LUN, use **powervt resize** on applicable platforms and then **vx disk resize** followed by **growfs** to reclaim unused space.

Migrate devices under LVM control on UNIX and Linux hosts

To migrate native-named devices under LVM control on UNIX and Linux Hosts:

1. Follow the migration steps in [Chapter 2, “Migration Procedures,”](#) until committing the migration. Choose a convenient time to take applications or file systems offline and reconfigure them.

Note: Native to pseudo migrations are not supported on Linux SLES 11 kernels.

2. After committing the migration, close the application and unmount file systems that use the volume group containing the source device name.
3. Close the logical volumes and then type **vgchange -a n <volume_group>** to deactivate the specified volume group.

Note: You must close the logical volumes before executing this command. For example, if the logical volume contains a file system, unmount the file system.

4. Type **vgexport /tmp/ <volume_group_export_file> <volume_group>** to capture all devices in the volume group so you can later import them back into the system.
5. Edit the device mapping from the source to the target-device names. Type:

```
sed -e s/<source>/<target>/ /tmp/ <VG_export_file> /tmp/<VG_import_file>
```

where:

- *<source>* is the source-device name in the migration
- *<target>* is the target-device name in the migration
- *<VG>* is the name of the volume group

Note: Do not use the fully qualified path names (*/dev/rdsk/c#t#d#*) when specifying the device names.

6. Type **powermig undoRedirect -hd <handle>** to turn off the redirection of I/O from the source to the target.
7. Make a directory, create a special file, and recreate the volume-group device node entry to import. Type:

```
mkdir /dev/<VG>  
mknod /dev/<VG>/group c 64 <minor_number>
```

Note: The minor number must be unique among existing volume groups.

8. Type **vgimport -f /tmp/** *<VG_import_file>* to import the volume group back into the system.
9. Type **vgchange -a y** *<volume_group>* to set the LVM volume group availability.
10. Type **vgcfgbackup** *<volume_group>* to create or update the LVM volume group configuration backup file.
11. Restart applications or mount file systems that use the volume group.
12. Optionally, clean up temporary files. Type **rm /tmp/***<VG_export_file>*
/tmp/*<VG_import_file>*

Note: If the LVM-controlled device does not use LVM volumes for storage, you must reconfigure the application to use the new target-device name before restarting it.

13. Run **powermig cleanup** to clean up the migration.

Note: During VNX/CLARiiON to VNX/CLARiiON Host Copy migration with I/O running on iSCSI devices, **powermig cleanup -format** takes substantially longer since it performs a full format of the disk.

Migrate pseudo-named devices under LVM control on UNIX and Linux hosts

1. Follow the migration steps in [Chapter 2, “Migration Procedures,”](#) until committing the migration.
2. After committing the migration, Migration Enabler swaps all underlying paths and designates the target device as the recipient of all I/O requests. Run **pvdisplay** to confirm that the former source device is still the active device for your physical volume.
3. Run **powermt display dev=***<source_device>* to display the source device information and to confirm the corresponding native devices that are to be removed.

For example, where emcpowerb is the source device (whose native devices are to be removed):

```

# powermt display dev=emcpowerb
Pseudo name=emcpowerb
CLARiiON ID=FNM00022300718
Logical device ID=6006016038E02400BCCBAF9428AADE1
state=alive; policy=CLAROpt; priority=0; queued-IOs=0
Owner: default=SP A, current=SP A      Array failover mode: 4
=====
----- Host ----- - Stor - -- I/O Path -- -- Stats ---
### HW Path          I/O Paths Interf.  Mode  State    Q-IOs  Errors
=====
1 qla2xxx             sdag      SP A0  active alive    0      0
1 qla2xxx             sdq       SP B1  active alive    0      0

```

4. Run **powermt remove dev=<source_device>**.
5. Run **powermt display dev=<source_device>** to confirm that the source device no longer exists. This should fail.
6. On Linux hosts only,
 - a. Run **powermt release**.
 - b. Run **echo "1" > /sys/block/<native_device>/device/delete** on all corresponding native devices. For example,

```
# echo "1" > /sys/block/sdag/device/delete
# echo "1" > /sys/block/sdq/device/delete
```
7. Verify that the pseudo and corresponding native devices have been removed.
8. Run **powermig cleanup** to clean up the migration.

Migrations in a cluster environment

Note: When migrating devices in a cluster configuration, you must select one node to remain active during the migration, the only exception being the MSCS environment, where you do not have to shutdown any nodes in the cluster to perform a cluster migration.

[Table 4 on page 40](#) lists the supported clusters. The *EMC E-Lab Interoperability Navigator™* provides additional information on supported cluster versions.

Table 4 Supported clusters

	AIX HACMP	Microsoft Cluster Server (MSCS)	Solaris Cluster	Veritas Cluster Server (VCS)
AIX	Yes	N/A	N/A	Yes
HP-UX	N/A	N/A	N/A	Yes
Linux	N/A	N/A	N/A	No
Solaris	N/A	N/A	Yes	Yes
Windows	N/A	Yes	N/A	Yes

NOTICE

Quorum device migration is not supported. You can add a new device (quorum) from the target array and then remove the older one.

Migrate in MSCS cluster environment

Microsoft Cluster Server (MSCS) provides a clustering technology that keeps server-based applications highly available, regardless of individual component failures. The Microsoft MSCS documentation provides complete information on planning, configuring, managing, and troubleshooting MSCS.

When migrating devices in an MSCS cluster environment with PowerPath Migration Enabler, failover groups need not be disabled and you do not have to shutdown any nodes in the cluster to perform a cluster migration. PowerPath Migration Enabler supports multiple cluster migrations from multiple cluster nodes. You can perform non-cluster migrations between non-cluster disks in any cluster node.

Migrate physical disks

- ◆ [“Common terminologies” on page 70](#) describes some common terminologies associated with MSCS cluster migrations.
- ◆ [“Before you start a migration” on page 21](#) provides prerequisites for migrating physical disks in an MSCS cluster environment.

- ◆ “Cluster migration best practices” on page 70 describes some of the best practices that you need to follow while migrating physical disks in an MSCS cluster environment.
 - ◆ “Troubleshooting cluster migrations” on page 69 provides more information on some of the common error messages and suggested actions while migrating physical disks in an MSCS cluster environment.
1. Ensure that the target disk is initialized and accessible from all nodes.
 2. Ensure that the PowerPath Migration Enabler cluster resource is configured for the source cluster disk.

The *EMC PowerPath Family CLI and System Messages Reference Guide* provides more information on configuring cluster resources.

Note: The source should be a cluster disk with a corresponding physical disk resource in the cluster and the target device should be a non-cluster disk.

3. Ensure that the underlying technology is configured, and then run **powermig setup -cluster** command.

Note: The cluster group of the source disk should be owned by the node where this command is being executed.

4. Run **powermig sync** command to begin syncing the migrations.
5. If a failover occurs during migration, the migration is automatically aborted. You need to restart the synchronization once again. Perform the following steps:
 - a. Failback the resources to the migration node.
 - b. Run **powermig sync** command to begin syncing the migrations once again.
6. Run the **powermig info -query** command when the synchronization is complete.
7. Run the **powermig commit** command to commit the migration.
8. If failover occurs after commit, you need to perform the following before moving to step 9.
 - Failback the resources to the migration node.
9. Run the **powermig cleanup** command to clean up the migration.

Migrate physical disks on Windows Server 2003

Note: When migrating physical disks on Windows Server 2003, you must select one node to remain active during the migration. All failover groups must be moved to the active node.

To migrate physical disks on Windows Server 2003

1. Ensure the target disk is initialized and accessible from all nodes.
2. Select the node that will remain active during the migration (only one node with access to the source and target devices can remain active). Move all failover groups to the active node.
3. Shut down the remaining nodes in the cluster.
4. Ensure the underlying technology is configured, and then run the **powermig setup -force** command.
5. Run the **powermig sync** command to begin syncing the migrations.
6. When the synchronization is complete, type **powermig info -all** command to know the status of the migration.
7. Run the **powermig selectTarget** command to select the target logical unit.
8. Run the **powermig commit** command to commit the migration.
9. Run the **powermig cleanup** command to clean up the migration.

Migrate cluster shared volume

A cluster shared volume is a shared disk containing an NTFS volume that is made accessible for read and write operations by all nodes within a Windows Server Failover Cluster. Cluster shared volume does not have a physical disk resource.

Note: When migrating devices in a cluster shared volume, you must select one node to remain active during the migration. All failover groups must be moved to the active node.

Note: This procedure is supported only with Host Copy.

To migrate cluster shared volume:

1. Ensure the target disk is initialized and accessible from all nodes.
2. Select the node that will remain active during the migration. (Only one node with access to the source and target devices can remain active.) Move all failover groups to the active node.
3. Shut down the remaining nodes in the cluster.
4. Ensure the underlying technology is configured, and then run **powermig setup** command.
5. Run **powermig sync** command to begin syncing the migrations.
6. When the synchronization is complete, type **powermig info -all** command to know the status of the migration.
7. Run the **powermig selectTarget** command to select the target logical unit.
8. Run the **powermig commit** command to commit the migration.
9. Run the **powermig cleanup** command to clean up the migration.

Migrate in HACMP environment (AIX hosts)

Note: This section applies only to AIX hosts.

PowerPath Migration Enabler supports migrations in a standby configuration, where one cluster node is active and the other is passive.

In the following procedure:

- ◆ hdiskpower7 is the source device
- ◆ hdiskpower9 is the target device

Note: Steps 3 through 11 in the following procedure must be done for each source and target migration pair.

1. Map the intended target device(s) to the active and passive nodes, and then run the **cfgmgr** command.
2. Stop the HACMP service on the passive node.

3. On the active node, set up the migration and synchronize the devices, as described in [Chapter 2, “Migration Procedures.”](#)
4. On the active node, run **powermt display dev=all** and note the unique device ID for the **target device (hdiskpower9)** and the **source device (hdiskpower7)** in the migration. The unique device ID consists of the following information:
 - Symmetrix device ID = Symmetrix ID + logical device ID
 - VNX/CLARiiON device ID = VNX/CLARiiON ID + logical device ID (LUN)
5. On the passive node:
 - a. Run **powermt display dev=all**, and locate the device whose ID matches the device ID of the **target device**, noted in step 4. (That is, the device ID associated with **hdiskpower9** on the active node.)
 - b. On the passive node, run **chdev -a pv=yes -l hdiskpowerX** on the **hdiskpower** device identified in step 5 a.

where *X* is the **hdiskpower** number identified in step 5.

For example, **chdev -a pv=yes -l hdiskpower9**
6. Type **lspv | grep hdiskpowerX** to verify the device you modified in step 5 b has the same PVID as the **source device** on the active node. (Identified in step 4, in this example, the source-device name is **hdiskpower7**).

where *X* is the source device on the active node.
7. Complete the remaining migration steps, as outlined in [Chapter 2, “Migration Procedures.”](#)

Note: After completing the migration steps, the device names on the active node only are swapped.

8. On the active node, determine the device ID of the original source device (in this example, **hdiskpower7**, which is now the target device on the active node).
9. On the passive node, run **powermt display dev=all**, and locate the device whose device ID is the same as the ID of the original source device (**hdiskpower7**, which is now the target device on the active node).
10. On the passive node, run **chdev -a pv=clear -l hdiskpowerX** on the **hdiskpower** device identified in step 9:

where *X* is the device identified in step 9.

For example, `chdev -a pv=clear -l hdiskpower7`

11. Use the `lspv` command to verify each node has only one device name associated with the PVID of the target device on the active node.
12. Synchronize the cluster on both nodes and make sure synchronization completes.
13. Start HACMP on the passive node, and ensure the cluster service starts.
14. Fail the resource group from the active to the passive node and verify the Volume Group is activated.

Migrate between pseudo-named devices in VCS environment

Follow the steps in the example given below to migrate data between pseudo-named array devices in a Veritas Cluster Server (VCS) environment for Host Copy and Open Replicator technologies.

Note: PowerPath must be installed and properly configured on all nodes in the cluster. The VCS cluster must be configured with at least one service group running on the active node. The service group must be associated with a VxVM disk group with EMC devices. "Configuring applications and resources in VCS" in the *Veritas Cluster Server User Guide* provides additional information.

1. Zoning recommendations:

Currently, VCS is only working properly in a single FA zoning for Symmetrix devices environment (each hosts HBA is zoned per only one of the Symmetrix FAs). VNX/CLARiiON array ports should be visible on Symmetrix arrays, and must be added to zoning configuration to perform PowerPath Migration Enabler Open Replicator migrations from VNX/CLARiiON to Symmetrix.

HBA1 - FA1, SPA1, SPB0 (hba1, Target Symm FA1, Source VNX/CLARiiON SP ports).

HBA2 - FA2, SPA0, SPB1 (hba2, Target Symm FA2, Source VNX/CLARiiONs SP ports):

```
# powermt display paths | more
```

```
Symmetrix logical device count=14
```

```
=====
----- Host Bus Adapters ----- Storage System ----- I/O Paths -
### HW Path ID Interface Total Dead
=====
```

4	lpfc	000187461234	FA	7dA	14	0
5	lpfc	000187461234	FA	7cA	14	0

CLARiiON logical device count=16

```

=====
----- Host Bus Adapters ----- Storage System ----- - I/O Paths -
### HW Path ID Interface Total Dead
=====
4 lpfc APM00090500212 SP B0 16 0
4 lpfc APM00090500212 SP A1 16 0
5 lpfc APM00090500212 SP A0 16 0
5 lpfc APM00090500212 SP B1 16 0
=====

```

2. In the active node, configure PowerPath Migration Enabler Open Replicator or Host Copy as the underlying technology:

- Register Symmetrix ports on a VNX/CLARiiON side and add Symmetrix to VNX/CLARiiON storage group with VCS devices connected to nodes.
- Mask VNX/CLARiiON devices to the Symmetrix.

3. Display information for devices associated with the service group.

```

# vxdisk list -o alldgs | grep mydg
emcpowerd auto:cdsdisk mydg01 mydg online
emcpowere auto:cdsdisk mydg02 mydg online
emcpowerf auto:cdsdisk mydg03 mydg online

```

4. Run the **powermt display** command to display source device information.

```

# powermt display dev=emcpowerd | grep "device ID\|LUN "
Logical device ID=60060160574023003472AFDDA10ADE11 [LUN 7]

# powermt display dev=emcpowere | grep "device ID\|LUN "
Logical device ID=60060160574023003672AFDDA10ADE11 [LUN 9]

# powermt display dev=emcpowerf | grep "device ID\|LUN "
Logical device ID=60060160574023004172AFDDA10ADE11 [LUN 20]

```

5. Select target devices and display information using the **powermt display** command:

```

# powermt display dev=emcpowerw | grep "device ID\|LUN "
Logical device ID=03A3

# powermt display dev=emcpowerx | grep "device ID\|LUN "
Logical device ID=03A5

# powermt display dev=emcpowerz | grep "device ID\|LUN "
Logical device ID=04C3

```

- Before starting PowerPath Migration Enabler migrations, freeze any service group containing disk groups associated with the migrations:

```
# hagr -freeze service_group -persistent
```

- In the active node (service group is currently running) start PowerPath Migration Enabler migrations between source and target devices, and continue to completion.

```
# powermig setup -techType OR -src <source_dev> -tgt <target_dev> -no
# powermig sync -hd <handle_number> -no
# powermig selecttarget -hd <handle_number> -no
# powermig commit -hd <handle_number> -no
# powermig cleanup -hd <handle_number> -no
```

Run **powermig info -query -all** to monitor the current state of the migration, especially during the **syncing** state.

- Verify that emcpower names have swapped between source and target devices. Repeat [Step 6 on page 47](#) and [Step 7 on page 47](#) to verify the changes. If I/O was running to the disks during the migration, verify that there was no interruption or data corruption.
- From the active node display disk information for the VxVM disk group being used:

```
# vxdisk list -o alldgs | grep mydg
emcpowerd    auto:cdsdisk    mydg01  mydg online
emcpowere    auto:cdsdisk    mydg02  mydg online
emcpowerf    auto:cdsdisk    mydg03  mydg online
emcpowerw    auto:cdsdisk    -        (mydg) online udid_mismatch
emcpowerx    auto:cdsdisk    -        (mydg) online udid_mismatch
emcpowerz    auto:cdsdisk    -        (mydg) online udid_mismatch
```

- Import disks labeled as **udid_mismatch** as cloned disks by assigning a new disk group name:

```
# vxdg -n newdg -o useclonedev=on -o updateid import mydg
# vxdisk list -o alldgs
emcpowerd    auto:cdsdisk    mydg01  mydg online
emcpowere    auto:cdsdisk    mydg02  mydg online
emcpowerf    auto:cdsdisk    mydg03  mydg online
emcpowerw    auto:cdsdisk    -        (newdg) online clone_disk
emcpowerx    auto:cdsdisk    -        (newdg) online clone_disk
emcpowerz    auto:cdsdisk    -        (newdg) online clone_disk
```

- Type **hagr -unfreeze service_group -persistent** to unfreeze the service groups that were previously set to freeze state in [Step 6 on page 47](#).

12. Use the **Cluster Management Console** or type **hagrp -switch service_group -to linx116** to fail resource group to a passive node by using the or the following command.
13. Run **hastatus -sum** to verify that the service group has successfully switched to the other node.
14. Verify that I/O is still running after switching the service group to the new active node.
15. In the new active node, verify that the emcpower devices currently used in the VxVM disk group correspond to the physical devices used as targets during the PowerPath Migration Enabler migration.

Note: Names of pseudo devices may be different from the original node.

```
# powermt display dev=emcpowerd | grep "device ID\|LUN "
Logical device ID=03A3

# powermt display dev=emcpowere | grep "device ID\|LUN "
Logical device ID=03A5

# powermt display dev=emcpowerf | grep "device ID\|LUN "
Logical device ID=04C3
```

Migrate between pseudo-named devices in Sun Cluster with VxVM

Note: Quorum device migration is not supported. You can add a new device (quorum) from the target array and then remove the older one.

1. Configure target array devices only on the primary node.
2. Setup a migration on the primary node.
3. Sync the migration.
4. Prior to running **powermig selectTarget**, perform the following cluster specific steps that prevent failover during migration:

Note: Full cluster functionality will not be available from [step 4](#) through [step 12](#) .

- a. Type **scconf -c -D name=<DG_NAME>,numsecondaries=1** to set the number of the secondary node for VxVM disk group to 1 if it is more than one.

- b. Type **scconf -r -D name=<DG_NAME>,nodelist=<nodename>** to remove all secondary nodes from the VxVM device group.
 - c. Type **scstat -D** to verify that no secondary node is listed.
5. Run **powermig selectTarget**.
6. Run **powermig commit**.
7. Configure target array devices on all secondary nodes.
8. Update the VxVM database on the primary node:
 - a. Type **rm /etc/vx/disk.info**
 - b. Type **vxconfigd -kc**
9. Run **powermig cleanup**.

Note: During VNX/CLARiiON to VNX/CLARiiON Host Copy migration with I/O running on iSCSI devices, **powermig cleanup -format** takes substantially longer since it performs a full format of the disk.

10. Optionally rename the pseudo devices on all secondary nodes, as described in [“Rename a pseudo device” on page 49](#).
11. Update the VxVM database on all secondary nodes:
 - a. Type **rm /etc/vx/disk.info**
 - b. Run **vxconfigd -kc**
12. Type **scconf -a -D type=vxvm,name=<DG_NAME>,nodelist=nodename** to add all secondary nodes to the VxVM device group.
13. Type **scconf -c -D name=<DG_NAME>,numsecondaries=<no_of_secondarynode>** to set the number of secondary node for VxVM DG as required (optional: in case if more than one secondary node is required).

Rename a pseudo device

Follow the steps detailed below to rename a pseudo device, and refer to EMC Knowledgebase Solution emc166911 for additional information.

Note: Renaming a pseudo device is not supported on HP-UX.

1. Type **vxdisk offline** *<devices_need_to_rename>*.
2. Type **emcpadm getfreepseudos -n** *<#_of_pseudo_instances>* to check free available pseudo instances.
3. Rename source pseudo to available free pseudo.
4. Rename target pseudo to source pseudo name.
5. Type **vxdisk online** *<source_pseudo_name>*.

APPENDIX A

Migration Procedure Examples

This appendix provides information about the PowerPath Migration Enabler technology types and includes migration procedure examples for each of the technology types:

- ◆ [Host Copy migrations](#) 51
- ◆ [Encapsulation migrations](#) 54
- ◆ [Open Replicator migrations](#) 58
- ◆ [TimeFinder/Clone migrations](#) 61

Host Copy migrations

Host Copy migrations work in conjunction with the host operating system to migrate data from the source logical unit to the target. A Host Copy migration does not use or require a direct connection between the arrays containing the source and target logical units.

Host Copy procedure

1. Run the following command to setup the migration from the source to the target, and to specify the technology:

```
powermig setup -techType <techType> -src <sourceName> -tgt  
<targetName>|-file <fileName> [-throttleValue  
<throttleValue>] [-cluster] [-techHost <hostname> [-techPort <port#>]]  
[-nothin] [-force] [-noprompt]
```

Note: The **-cluster** argument is used only when migrating devices in an MSCS environment. This argument is not supported on Windows Server 2003.

For example,

```
powermig setup -techType hostcopy -src emcpower72a -tgt emcpower52a
```

Optionally, you can set throttle during **Setup**.

2. From the **Setup** state, type **powermig sync -handle** <handle> to start the bulk copying and cloning of application write I/Os.

- Optionally, you can set a throttle value of 0 through 9 in the **Syncing** state.
 - Optionally, you can pause the migration during peak I/O times. When you resume the migration, it continues from where it left off.
 - Optionally, you can abort the migration.
3. While the migration is syncing, type **powermig query -handle <handle>** to check the status.
 4. After synchronization, the migration automatically goes into the **SourceSelected** state.
 5. Type **powermig selectTarget -handle <handle>** to designate the target logical unit as the recipient of all I/O read requests.

Note: Before you commit the migration, you have the option of returning to the **SourceSelected** state or aborting the migration.

Skip this step while migrating physical disks in an MSCS cluster environment.

6. Type **powermig commit -handle <handle>** to commit the migration to designate the target as the recipient of all I/O requests.

Note: There is no returning to the source after you perform a **commit**.

- With native operating system devices as the source, the **commit** command places the migration into the **CommittedAndRedirected** state.
 - Also with native devices, you must execute the **undoRedirect** command to point to the new device and restart the application. This places the migration into the **Committed** state.
7. Type **powermig cleanup -handle <handle>** to clean up the migration. The migration ends after the **cleanup** step.

Related topics

- ◆ [“Changing the synchronization speed for migrations” on page 31](#)
- ◆ [“After you finish a migration” on page 25](#)

Host Copy throttle settings

The Host Copy technology type copies data from source to target in cycles that last five seconds each. shows the percentage of the host's time spent copying data as defined by the throttle value. The default value for Host Copy is 2.

Table 5 Host Copy throttle values

Throttle value	0	1	2	3	4	5	6	7	8	9
Percentage of time the host spends copying data	100%	60%	36%	22%	13%	7.8%	4.7%	2.8%	1.7%	1.0%

Host Copy throttle settings are ignored when you enable the Host Copy ceiling function.

The *EMC PowerPath Migration Enabler Host Copy* white paper available on EMC Online Support provides additional details about throttle settings.

Note: You can put as many Host Copy migrations into the **Syncing** state as desired; however only eight will be actively syncing at any given time. In Host Copy, a daemon in UNIX and a Service in Windows manages the bulk data copy. The daemon or Service initiates periodic calls to the kernel to perform copy operations in 5 second cycles, and maintains a separate log showing the progress of each migration it is managing.

During synchronization, Host Copy reads from the source device and writes to the target device in chunks of 256 KB.

Note: Synchronizing multiple migrations impacts host performance. One or more migrations may be paused, or throttle values may be modified to reduce application impact.

Host Copy ceiling settings

When you enable Host Copy ceiling, this ignores per migration throttling. Host Copy ceiling is specified on the **powermig options** command. provides additional information.

Migrating to encrypted devices

PowerPath Encryption with RSA is a software-based encryption tool that leverages RSA technology to encrypt user-specified PowerPath-controlled devices. Writes are encrypted and reads are decrypted to and from the encrypted LUN. Data must be migrated to an encrypted target device to encrypt plaintext data. Migration Enabler Host Copy is a non-disruptive method for performing this migration.

When migrating to an encrypted device, the target must be at least 65 KB (AIX, Solaris, or Windows) or 68 KB (Linux) larger than the source device to accommodate metadata requirements.

Encapsulation migrations

Set up Storage Element for Migration Enabler

To set up a Storage Element for PowerPath Migration Enabler, identify the LU that will be the source for the migration and expose it to the Invista (for Symmetrix LUs, set up masking; for VNX/CLARiiON LUs, add it to the appropriate storage group). Then, using the Invista Element Manager (IEM):

1. Right-click **Invista Instance** and select **Rediscover Backend SEs and Paths**.
2. Right-click an unimported **Storage Element** that corresponds to the source device (compare the "Unique ID" to the VNX/CLARiiON WWN or to the Symmetrix ID and LUN), and select **Set Not Ready**.

Putting the element in the **Not Ready** state (also called the **Faulted** state in some IEM views) disables I/O to and from the Storage Element and prevents from using it in any Clones, Data Mobility, or volume resizing operations.

3. Import the Storage Element.
4. Create a Virtual Volume from the Storage Element. Make the volume the full size of the Storage Element.

Note: You cannot expand this Virtual Volume until after the migration has been completed.

5. Add the volume to a Virtual Frame and expose it to the host as you would any other Virtual Volume.

The *EMC Invista Element Manager Administration Guide* contains more information on Virtual Frames and Virtual Volumes.

6. On the host running PowerPath Migration Enabler, make the volume visible to the host.
 - On AIX, type **emc_cfgmgr** or **cfgmgr -v**
 - On HP-UX, type **ioscan -fnC disk** and then type **insf -eC disk**
 - On Linux, refer to Dynamic LUN addition and removal in the *EMC PowerPath Installation and Administration Guide for Linux*.
 - On Solaris, type **devfsadm -C**
 - On Windows, perform device rescan from the Device Manager utility.
7. On UNIX and Linux hosts, type **powermt config** to update the configuration.
8. On UNIX and Linux hosts, type **powermt save** to save the configuration.

Identify target-device name

Once the Virtual Volume is visible on the host, identify the target-device name to use when you set up the migration session. The examples in this section use pseudo-device names in a Solaris host environment. For example: **emcpower44a**

Type **powermt display nonvirtual** to identify the device names for a migration.

The nonvirtual information is displayed for a device only if that device is qualified to be the target of an INVE (or) migration:

- ◆ The Virtual Volume is associated with exactly one Storage Element.
- ◆ The Storage Element is associated with exactly one Virtual Volume.
- ◆ The Virtual Volume and the Storage Element have the same capacity.
- ◆ The Storage Element has been set to **Not Ready**.

In the nonvirtual output:

- Pseudo name is the target-device name used setting up a migration.
- Nonvirtual pseudo name is the source-device name used when setting up a migration.

The following output, from a Solaris host, is truncated:

Note: Other operating systems may display differently. For example, you may not see **Pseudo name=emcpower** as shown in the output above. Also, in **NonVirtual pseudo name=emcpower** the word pseudo may be replaced with the word **device**.

```

Pseudo name=emcpower44a
  ID=APM00051002822
Logical device ID=60001440B0601C6D00589C10390003B2
NonVirtual Symmetrix ID=000187461238
NonVirtual device ID=067C
NonVirtual pseudo name=emcpower15a
state=alive; policy=ADaptive; priority=0; queued-IOS=0
=====
----- Host ----- - Stor - -- I/O Path - -- Stats ---
### HW Path          I/O Paths  Interf.  Mode state  Q-IOS Errors
=====
2306 pci@1e,600000/lpfc@2      c5t6d2s0  10          active alive    0    0
2306 pci@1e,600000/lpfc@2      c5t7d2s0  11          active alive    0    0
2307 pci@1f,700000/lpfc@2      c6t6d2s0  0C          active alive    0    0
2307 pci@1f,700000/lpfc@2      c6t7d2s0  0D          active alive    0    0

Pseudo name=emcpower9a
  ID=APM00051002822
Logical device ID=60001440B0601C6D00589C80390003B6
NonVirtual device ID=60060160EEE215000A5F05D3A915DB11
NonVirtual pseudo name=emcpower35a
state=alive; policy=ADaptive; priority=0; queued-IOS=0
=====
----- Host ----- - Stor - -- I/O Path - -- Stats ---
### HW Path I/O Paths  Interf.  Mode state  Q-IOS Errors
=====
2306 pci@1e,600000/lpfc@2      c5t6d4s0  10          active alive    0    0
2306 pci@1e,600000/lpfc@2      c5t7d4s0  11          active alive    0    0
2307 pci@1f,700000/lpfc@2      c6t6d4s0  0C          active alive    0    0
2307 pci@1f,700000/lpfc@2      c6t7d4s0  0D          active alive    0    0
    
```

Encapsulation procedure

1. Run the following command to setup the migration from the source to the target, and to specify the technology type:

```

powermig setup -techType <techType> -src <sourceName> -tgt
  <targetName>|-file <fileName> [-throttleValue
  <throttleValue>][-cluster] [-techHost <hostname> [-techPort <port#>]]
  [-nothin] [-force] [-noprompt]
    
```

Note: The **-cluster** argument is used only when migrating devices in an MSCS environment. This argument is not supported on Windows Server 2003.

For example,

```
powermig setup -techType INVE -src emcpower72a -tgt emcpower52a
```

2. Type **powermig sync -handle** *<handle>* to move to the **Syncing** state.

Optionally, you can abort the migration.

Encapsulation migrations are instantaneously synchronized; there is no bulk copying or cloning of I/Os. Once synchronized, the migration automatically goes into the **SourceSelected** state.

Note: Throttle and Pause and Resume commands are not supported for the encapsulation migrations.

3. Type **powermig query -handle** *<handle>* to check the status.
4. If you are not migrating in an MSCS cluster environment, type **powermig selectTarget -handle** *<handle>* to designate the target logical unit as the recipient of all I/O read requests.

Note: Before you commit the migration, you have the option of returning to the **SourceSelected** state or aborting the migration.

5. Type **powermig commit -handle** *<handle>* to commit the migration to designate the target as the recipient of all I/O requests.

Note: There is no returning to the source after you perform a **commit**.

- With native operating system devices as the source, the **powermig commit** places the migration into the **CommittedAndRedirected** state.
 - In this case, the application must now be reconfigured to point to the new (target) device and restarted. Then execute **powermig undoRedirect** to place the migration into the **Committed** state.
6. Remove host to the source device. Unmask the LU on a Symmetrix array or remove the LU from the storage group on a VNX/CLARiON array.

7. Type **powermig cleanup -handle** *<handle>* to clean up the migration. The migration ends after the **cleanup** step.

Related topic

[“After you finish a migration” on page 25](#)

Open Replicator migrations

PowerPath Migration Enabler works in conjunction with the Open Replicator hot pull method, which copies data through the SAN infrastructure from a remote logical unit to a Symmetrix storage system. Migration Enabler mirrors I/O to keep the source and target logical units synchronized throughout the migration process. For more information on Open Replicator, refer to the *EMC Solutions Enabler Symmetrix Open Replicator CLI Product Guide*, available on EMC Online Support.

Set up Symmetrix-to-Symmetrix migrations with Open Replicator

1. Run **powermt display dev=all** to verify the host can see the source and target devices.
2. Verify and update the zoning configuration as necessary. If zoning is enabled on the switch, every port on the target array that allows to the target logical unit must have to the source logical unit through at least one port on the source array.
3. If the storage port with the source logical unit is running Volume Logix (VCM), grant permission to the source logical unit through the target storage port.
4. On the host running PowerPath Migration Enabler, make the volume visible to the host.
 - On AIX, type **emc_cfgmgr** or **cfgmgr -v**
 - On HP-UX, type **ioscan -fnC disk** and then type **insf -eC disk**
 - On Linux, refer to Dynamic LUN addition and removal in the *EMC PowerPath Installation and Administration Guide for Linux*.
 - On Solaris, type **devfsadm -C**
 - On Windows, perform device rescan from the Device Manager utility.
5. Install either Solutions Enabler or Solutions Enabler Thin Client on the host.
6. On UNIX and Linux hosts, type **powermt config** to update the configuration.

7. On UNIX and Linux hosts, type **powermt save** to save the configuration.

Migrate EMC VNX or CLARiiON data to EMC Symmetrix arrays

1. Run **powermt display dev=all** to verify the host can see the source and target devices.
2. Verify and update the zoning configuration as necessary. If zoning is enabled on the switch, every port on the target array that allows to the target logical unit must have to the source logical unit through at least one port on the source array.
3. Add the EMC Symmetrix target WWN to the Storage Group that contains the source logical unit(s).

Note: A Symmetrix storage port must be registered as an initiator to be added to a VNX/CLARiiON Storage Group. You must try to initiate a PowerPath Migration Enabler session for the Symmetrix to appear in connectivity status.

4. On the host running PowerPath Migration Enabler, make the volume visible to the host.
 - On AIX, type **emc_cfgmgr** or **cfgmgr -v**
 - On HP-UX, type **ioscan -fnC disk** and then type **insf -eC disk**
 - On Linux, refer to Dynamic LUN addition and removal in the *EMC PowerPath Installation and Administration Guide for Linux*.
 - On Solaris, type **devfsadm -C**
 - On Windows, perform device rescan from the Device Manager utility.
5. On UNIX and Linux hosts, type **powermt config** to update the configuration.
6. On UNIX and Linux hosts, type **powermt save** to save the configuration.

Migration Enabler supports Solutions Enabler Thin Client for Open Replicator migrations.

Open Replicator procedure

1. Setup the migration from the source to the target, and specify the technology type:

```
powermig setup -techType <techType> -src <sourceName> -tgt
<targetName>|-file <fileName> [-throttleValue
<throttleValue>] [-cluster] [-techHost <hostname> [-techPort <port#>]]
[-nothin] [-force] [-noprompt]
```

Note: The **-cluster** argument is used only when migrating devices in an MSCS environment. This argument is not supported on Windows Server 2003.

For example,

```
powermig setup -techType OR -src emcpower72a -tgt emcpower52a
```

If using the SE Thin Client, specify the server that has the full Solutions Enabler.

For example,

```
powermig setup -techType OR -src harddisk7 -tgt harddisk5 -techHost
192.168.1.1
```

Optionally, you can set throttle during **Setup**, but cannot change it until the migration enters the **Syncing** state.

2. In the **Syncing** state, type **powermig sync -handle <handle>** to start the bulk copying and cloning of application write I/Os. Once synchronized, the migration automatically goes into the **SourceSelected** state.
 - Optionally, you can set a throttle value of 0 through 9 in the **Syncing** state.
 - Optionally, you can pause the migration during peak I/O times. When you resume the migration, it continues from where it left off.
 - Optionally, you can abort the migration.
3. Type **powermig query -handle <handle>** to check the status.
4. Type **powermig selectTarget -handle <handle>** to designate the target logical unit as the recipient of all I/O read requests.

Note: Before you commit the migration, you have the option of returning to the **SourceSelected** state or aborting the migration.

- Skip this step while migrating physical disks in an MSCS cluster environment.

5. Type **powermig commit -handle** *<handle>* to commit the migration to designate the target as the recipient of all I/O requests.

NOTICE

There is no returning to the source after you perform a **commit**.

- a. With native operating system devices as the source, the **commit** command places the migration into the **CommittedAndRedirected** state.
 - b. You must run the **undoRedirect** command to point to the new device and restart the application. This places the migration into the **Committed** state.
6. Type **powermig cleanup -handle** *<handle>* to clean up the migration. The migration ends after the **cleanup** step.

Related topics

- ◆ [“Changing the synchronization speed for migrations” on page 31](#)
- ◆ [“After you finish a migration” on page 25](#)

Understanding Open Replicator throttle settings

For Open Replicator migrations, the throttle value adjusts the speed at which data is transferred between the source and target logical units. For Open Replicator, the throttle works similarly to the Solutions Enabler **symrcopy** command with the **set pace** argument. A lower throttle value can significantly increase resource consumption on the target logical unit’s array.

TimeFinder/Clone migrations

The TimeFinder/Clone technology is used to copy from LUNs within the same EMC Symmetrix array. Migration Enabler supports Solutions Enabler (SE) Thin Client for TimeFinder/Clone migrations.

TimeFinder/Clone procedure

1. Setup the migration from the source to the target, and specify the technology Type:

```
powermig setup -techType <techType> -src <sourceName> -tgt  
<targetName> -file <fileName> [-throttleValue  
<throttleValue>] [-cluster] [-techHost <hostname> [-techPort <port#>]]  
[-nothin] [-force] [-noprompt]
```

Note: The **-cluster** argument is used only when migrating devices in an MSCS environment. This argument is not supported on Windows Server 2003.

For example,

```
powermig setup -techType symclone -src emcpower72a -tgt emcpower52a
```

If using the SE Thin Client, specify the server that has the full Solutions Enabler.

For example,

```
powermig setup -techType symclone -src harddisk7 -tgt harddisk5 -techHost  
192.168.1.1
```

Note: Symclone is the name that represents the TimeFinder/Clone technology type.

- Optionally, you can set throttle during **Setup**.
 - Optionally, you can set the Thin Client.
2. In the **Syncing** state, type **powermig sync -handle** <handle> to start the bulk copying and cloning of application write I/Os. Once synchronized, the migration automatically goes into the **SourceSelected** state.
 - Optionally, you can set a throttle value of 0 through 9 in the **Syncing** state.
 - Optionally, you can pause the migration during peak I/O times. When you resume the migration, it continues from where it left off. Pause is supported only when the source and target are the same size.
 - Optionally, you can abort the migration.
 3. Type **powermig query -handle** <handle> to check the status.
 4. Type **powermig selectTarget -handle** <handle> to designate the target logical unit as the recipient of all I/O read requests.

Note: Before you commit the migration, you have the option of returning to the **SourceSelected** state or aborting the migration.

- Skip this step while migrating physical disks in an MSCS cluster environment.
5. Type **powermig commit -handle** *<handle>* to commit the migration to designate the target as the recipient of all I/O requests.

NOTICE

There is no returning to the source after you perform a commit.

- a. With native operating system devices as the source, the **commit** command places the migration into the **CommittedAndRedirected** state.
 - b. You must run the **undoRedirect** command to point to the new device and restart the application. This places the migration into the **Committed** state.
6. Type **powermig cleanup -handle** *<handle>* to clean up the migration. The migration ends after the **cleanup** step.

Related topics

- ◆ [“Changing the synchronization speed for migrations” on page 31](#)
- ◆ [“Using the Solutions Enabler Thin Client for Migration Enabler” on page 32](#)
- ◆ [“After you finish a migration” on page 25](#)

Understanding TimeFinder/Clone throttle settings

For TimeFinder/Clone migrations, the throttle value changes the Symmetrix Quality of Service (QoS) value for the source and target logical units. The QoS value affects all I/O performance for the devices, and not just migration performance. The **powermig setup** command sets the QoS value to 0, if not specified with the throttle option during setup. After a migration completes, the QoS value set during the migration remains on the source and target.

Symmetrix QoS values range from 0 to 16 and Migration Enabler throttle values range from 0 to 9. Migration Enabler maps the throttle values to QoS values as shown in .

Table 6 TimeFinder/Clone throttle and QoS values

Throttle value	0	1	2	3	4	5	6	7	8	9
Corresponding QoS value	0	2	4	5	7	9	10	12	14	16

APPENDIX B

Troubleshooting a Migration

This appendix provides guidance for troubleshooting problems you may encounter during a migration:

- ◆ [Audit and error log messages](#) 65
- ◆ [Reporting a problem](#)..... 66
- ◆ [Migration impact from a graceful host reboot](#) 66
- ◆ [Faults after a host crash](#) 66
- ◆ [Other reported faults](#)..... 68
- ◆ [Resolving a process hang](#)..... 69
- ◆ [Troubleshooting cluster migrations](#) 69

Audit and error log messages

PowerPath Migration Enabler uses audit logs for recording events that change the migration state, and error logs that capture unexpected events that occur during the course of a migration. The *PowerPath Family CLI and System Messages Reference Guide* describes audit logging, including instructions for setting up logging in your environment. It also lists Migration Enabler error messages, descriptions, and corrective action suggestions.

Audit log message example

The following audit log entry indicates that a **powermig sync** command was executed on Migration Handle number 3 on December 21st at approximately 3 p.m. In this example from an audit log, the state of the migration is now Syncing:

```
Dec 21 15:00:23 Ippa020 emcpAudit: PPME: Info:  
    handle=3,event=stateChanged, state=syncing, cmd=sync
```

Error log message example

The following error log entry indicates that the **powermig setup** command failed because either the source or the target device was already a part of another Migration Enabler migration session:

```
Dec 21 14:49:16 Ippa020 PPME: API: Error: Setup failed: PPME  
error(14): Devices already involved in a migration
```

Reporting a problem

If you encounter a problem with PowerPath Migration Enabler, run the **emcgrab** utility, which collects system and software configuration information, and then follow the instructions in this section to collect Migration Enabler information. You can download emcgrab from the EMC Online Support. Refer to the emcgrab ReadMe.txt file on EMC Online Support for prerequisite information and instructions for using the utility.

Note: EMC Grab 3.9 or later is required to collect information about PowerPath Encryption with RSA or Host Copy migrations. EMC Grab version 3.7 or later collects information about PowerPath Migration Enabler for Open Replicator or Encapsulation.

Migration impact from a graceful host reboot

When a host gracefully shuts down, there is no impact to the migration. On Windows, the maximum number of concurrent migrations that continue after a graceful reboot is 60.

With Host Copy, Migration Enabler records a checkpoint in the Migration Enabler database. The migration progress checkpoint records approximately every 30 seconds on a local or SAN disk, depending on where the boot disk is located. When the host starts up after a graceful shutdown, the migration resumes from its last checkpoint. This behavior is for migrations that are syncing or paused.

With Open Replicator and TimeFinder/Clone, the copy continues on the array.

Faults after a host crash

When a host crashes, PowerPath Migration Enabler cannot guarantee the success of application writes, and cannot guarantee that the data on the source and target devices are the same after such a crash. The fault state (either TargetLUfault or SourceLUfault) displays in the **powermig info** output. PowerPath Migration Enabler reacts depending on the state of the migration at the time of the crash. Faults are recorded in the **Syncing**, **SourceSelected**, and **TargetSelected** states.

Note: This does not apply to Encapsulation (INVE) migrations.

States before the host crash:

- ◆ Syncing
- ◆ SourceSelected
- ◆ TargetSelected

Each state causes a different result when the host reboots after the crash:

- ◆ In the **Syncing** state, the synchronization has been stopped. PowerPath Migration Enabler reports a TargetLUFault. Application reads and writes are still directed to the source device. The migration returns to the **Setup** state. Restart the migration with the **powermig sync** command.
- ◆ In the **SourceSelected** state, the migration has completed. Both source and target devices are fully synchronized. However, when the host crashes, PowerPath Migration Enabler reports a TargetLUFault. I/O writes may have been in flight during the time of the crash. Therefore, Host Copy can no longer guarantee that the data is the same on both the source and the target. The migration returns to the **Setup** state. Restart the migration with the **powermig sync** command.
- ◆ In the **TargetSelected** state, the migration has completed. Both source and target devices are fully synchronized. However, when the host crashes, PowerPath Migration Enabler reports a SourceLUFault. I/O writes may have been in flight during the time of the crash. Therefore, Host Copy can no longer guarantee that the data is the same on both the source and the target. The migration proceeds to the **Committed** state.

[Table 7 on page 67](#) shows the three states when Logical Unit faults can be reported.

Table 7 Reported faults after a host crash

State at time of Host Crash	Host crash	Fault state reported	State after crash	powermig command	Data loss
Syncing	Yes	TargetLUFault	Returns to Setup state	sync or cleanup	No
SourceSelected	Yes	TargetLUFault	Returns to Setup state	sync or cleanup	No
TargetSelected	Yes	SourceLUFault	pseudo devices: moves to Committed state	cleanup	No
			native devices: moves to Committed and Redirected state	undoRedirect	

When a host crashes, PowerPath Migration Enabler returns errors to both reads and writes for any device that is in a migration where the source and target is kept synchronized. This ends when PowerPath Migration Enabler **startup recover** automatically runs. This protects data integrity until device default states are processed. Once startup recovery runs, it is safe for reads and writes to flow to these devices.

Other reported faults

In addition to a host crash, errors or faults may occur due to changes in the storage, the host, or the SAN. If an I/O failure occurs in the course of a write I/O operation, and if Migration Enabler is currently cloning write I/Os (Invista Encapsulations do not clone writes or reads) and then Migration Enabler records a Logical Unit fault, the fault state (either TargetLUfault or SourceLUfault) appears in the powermig info output. Migration Enabler reports different types of faults depending on the state of the migration at the time the I/O failure occurs. Faults are recorded in the **Syncing**, **SourceSelected**, and **TargetSelected** states.

States at the time of the fault:

- ◆ Syncing
- ◆ SourceSelected
- ◆ TargetSelected

[Table 8 on page 68](#) shows the three states when Logical Unit faults are reported.

Table 8 Other reported faults

Current state	Source device failure	Target device failure	Fault state reported	Cause	powermig command	Data loss
Syncing	Yes	N/A	SourceLUfault ¹	I/O write to source fails	abort ²	No
		Yes	TargetLUfault ¹	I/O write to target fails	abort	No
SourceSelected	Yes	N/A	SourceLUfault ¹	I/O write to source fails	selectTarget ²	No
		Yes	TargetLUfault ¹	I/O write to target fails	abort	No
TargetSelected	Yes	N/A	SourceLUfault ¹	I/O write to source fails	commit	No
		Yes	TargetLUfault ¹	I/O write to target fails	abort ²	No

1. Migration Enabler Encapsulation (INVE) migrations do not report LU faults.

2. When this fault occurs under these conditions, errors return on all I/Os until you perform this transaction.

Resolving a process hang

Some **powermig** commands do not respond when you press **Ctrl** and **C**. If a **powermig** command hangs, you will be unable to type any subsequent **powermig** commands.

On UNIX hosts, to clear the hung process:

1. Run the **ps -ef | grep** command to retrieve the process ID.
2. Run **kill -9 <process_ID>** to forcefully kill the process.

Troubleshooting cluster migrations

[Table 9 on page 69](#) lists common error messages and suggested actions for physical disk migrations in an MSCS cluster environment.

Table 9 Cluster migration error messages

Message	Suggested action
Error 75 Cluster disk being migrated is not owned by this node.	Migration node should own the cluster group for the source device to perform necessary migrations. You need to failover the required cluster group to the migration node.
Error 79 Not all nodes have reached the correct state to complete this operation, please try this or the recover command again later.	Ensure that all cluster nodes are up and running. Ensure that the cluster service is running on all nodes. Make sure that the target device is connected and is available to all the nodes. Run the powermig recover command.
Error 82 The source device does not appear to have a dependent PPME resource instance.	The PowerPath Migration Enabler resource for the source cluster disk may not be configured. Use powermigcl command to configure the PowerPath Migration Enabler resource. “Configuring PowerPath Migration Enabler cluster resources” on page 23 provides the required steps.
Error 86 Cluster migration must be aborted	If migration is in the needsRecovery (commit)state, but roll forward point has not been reached, the migration should be aborted. You need to run powermig abort command to abort the migration.

Table 9 Cluster migration error messages

Message	Suggested action
Error 89 Source must be a cluster disk	Source device for cluster migration should be a cluster disk. You need to add a cluster disk resource for the device to setup a cluster migration.
Error 90 Source is a cluster disk (use -cluster to setup a cluster migration).	Cluster disk can be used as a source device for migration only if -cluster option is specified.
Error 91 Target must not be a cluster disk	The source should be a cluster disk with a corresponding physical disk resource in the cluster and the target device should be a non-cluster disk. You need to remove cluster resource for the device to use it as a target in migration.

Cluster migration best practices

This section describes some of the best practices that needs to be followed while migrating physical disks in an MSCS cluster environment.

- ◆ **Do not remove the PowerPath Migration Enabler cluster resource** - Do not remove the PowerPath Migration Enabler cluster resource dependency that is listed on the cluster disk if there are still pending cluster migrations.
- ◆ **Do not evict cluster nodes** - Do not evict the cluster nodes if there are pending cluster migrations.

Common terminologies

This section describes some common terminologies associated with MSCS cluster migrations.

- ◆ **Migration node** - Migration node (MN) is the node where the migration was initiated. This is the node that owns the application or the physical disk. Each migration activity can only take place on one node, i.e., the migration node. Migration node maintains the migration states and controls the migration technology.
- ◆ **Failover node** - Failover node (FO) is any other node in the cluster other than the migration node. It applies access control to source and target devices as appropriate and swaps source and target devices during commit.

- ◆ **Roll forward point** - This is the point during a commit operation, just before write cloning from source to target device is disabled. Once the migration reaches this point, migration cannot be aborted and must be committed.

APPENDIX C

Support Summary

This appendix provides a table of the commands supported for the various platforms and technology types.

Migration Enabler technology type support

[Table 10 on page 73](#) lists for all versions of PowerPath Migration Enabler the commands supported with each platform and technology type.

Table 10 Migration Enabler technology type support (1 of 2)

Commands	Supported technology types				
	Version 5.2	Version 5.3	Version 5.5	Version 5.6	Version 5.7
	HP-UX, Solaris, Windows	AIX, Linux, Solaris, Windows	AIX, Linux, Solaris, Windows	Linux	Linux, Windows
powermig abort	HOSTCOPY, INVE, OR, SYMCLONE ¹	HOSTCOPY, INVE ² , OR, SYMCLONE	HOSTCOPY, INVE, OR, SYMCLONE	HOSTCOPY, INVE, OR, SYMCLONE	HOSTCOPY, INVE, OR, SYMCLONE
powermig cleanup					
powermig commit					
powermig getHandle					
powermig help					
powermig info					
powermig options	HOSTCOPY (HP-UX only)	N/A ³	HOSTCOPY (AIX, Linux, Solaris only)	HOSTCOPY	HOSTCOPY
powermig pause	HOSTCOPY, OR, SYMCLONE ⁴	HOSTCOPY, OR, SYMCLONE	HOSTCOPY, OR, SYMCLONE	HOSTCOPY, OR, SYMCLONE	HOSTCOPY, OR, SYMCLONE
powermig query	HOSTCOPY, INVE, OR, SYMCLONE	HOSTCOPY, INVE, OR, SYMCLONE	HOSTCOPY, INVE, OR, SYMCLONE	HOSTCOPY, INVE, OR, SYMCLONE	HOSTCOPY, INVE, OR, SYMCLONE
powermig recover					

Table 10 Migration Enabler technology type support (2 of 2)

Commands	Supported technology types				
	Version 5.2	Version 5.3	Version 5.5	Version 5.6	Version 5.7
	HP-UX, Solaris, Windows	AIX, Linux, Solaris, Windows	AIX, Linux, Solaris, Windows	Linux	Linux, Windows
powermig resume	HOSTCOPY, OR, SYMCLONE				
powermig selectSource	HOSTCOPY, INVE, OR, SYMCLONE				
powermig selectTarget					
powermig setup					
powermig sync					
powermig throttle	HOSTCOPY, OR, SYMCLONE				
powermig version	HOSTCOPY, INVE, OR, SYMCLONE				

1. Migration Enabler 5.2 for Solaris and Windows does not support SYMCLONE. Supported only on Migration Enabler 5.2 HP-UX.
2. Encapsulation (INVE) is supported only with Invista storage systems. It is not supported for EMC VPLEX storage systems.
3. N/A means the command or utility is not supported on the platform.
4. TimeFinder/Clone supports pause and resume only when the source and target devices are of the same size.

APPENDIX D

Disk Labeling powerformat Command for Solaris

This appendix provides detailed information about the **powerformat** command for Solaris.

Solaris and PowerPath Migration Enabler disk labeling

When you perform a heterogeneous migration with PowerPath Migration Enabler on a Solaris host, the source logical unit's disk label is copied to the target along with the source's data. PowerPath Migration Enabler changes the ASCII name of the target logical unit's disk label to one that matches the logical unit type of the target array.

For example, when migrating data from a VNX/CLARiiON logical unit (DGC-RAID-0206) to a Symmetrix logical unit (EMC-SYMMETRIX-5671) using Open Replicator as the underlying technology, the VNX/CLARiiON logical unit's disk label (DGC-RAID-0206) is copied to the target Symmetrix logical unit. In VNX/CLARiiON-to-Symmetrix migrations only, PowerPath Migration Enabler changes portions of the ASCII disk label name to match the target logical-unit type. In this example, the ASCII name of the disk label is changed from *DGC-RAID-0206 cyl...* to *EMC-SYMMETRIX-5671 cyl...* This change to the ASCII disk-label name ensures the correct logical-unit type (for example, *EMC-SYMMETRIX-5671*) appears when you run the Solaris **format** utility.

Multiple array type entries in the format utility output

PowerPath Migration Enabler changes the ASCII disk-label name of the target logical unit as part of the migration process; however, the disk geometry on the target logical unit is unchanged. (It has the disk geometry of the source.) You can change the disk geometry and partition information after a migration using the **powerformat** command. Refer to for more information.

Usually the **format type** command displays one entry for each drive type. After a heterogeneous migration, the **format type** command sees a different geometry for the target (Symmetrix) logical unit. It sees the source (VNX/CLARiiON) logical unit's

geometry. The logical-unit label EMC-SYMMETRIX-5671 appears in the **format** output twice because the two Symmetrix 5671 logical units are recognized as different array types.

format utility example

The following example shows the **format** utility before and after a VNX/CLARiiON logical unit (DGC-RAID-0206) is migrated to a Symmetrix logical unit (EMC-SYMMETRIX-5671) using Open Replicator as the underlying technology.

Before the migration

The Solaris **format** utility shows the following output.

Note: A logical unit with the EMC-SYMMETRIX-5671 identification exists on the host before the migration is performed.

```
format>type
```

```
AVAILABLE DRIVE TYPES:
 0. Auto configure
 1. Quantum ProDrive 80S
 2. Quantum ProDrive 105S
 3. CDC Wren IV 94171-344
 4. SUN0104
 5. SUN0424
 6. SUN1.0G
 7. SUN1.05
 8. SUN1.3G
 9. SUN2.1G
10. SUN2.9G
11. Zip 100
12. EMC-SYMMETRIX-5568
13. EMC-SYMMETRIX-5671    <— original drive type
14. DGC-RAID5-0206
```

After the migration

After you migrate a VNX/CLARiiON source logical unit with disk label DGC-RAID-0206 to a Symmetrix target logical unit with disk label EMC-SYMMETRIX-5671, **format** displays the following output.

Note: Two entries now appear for EMC-SYMMETRIX-5671.

```
format>type
```

```
AVAILABLE DRIVE TYPES:
```

- 0. Auto configure
- 1. Quantum ProDrive 80S
- 2. Quantum ProDrive 105S
- 3. CDC Wren IV 94171-344
- 4. SUN0104
- 5. SUN0424
- 6. SUN1.0G
- 7. SUN1.05
- 8. SUN1.3G
- 9. SUN2.1G
- 10. SUN2.9G
- 11. Zip 100
- 12. EMC-SYMMETRIX-5568
- 13. EMC-SYMMETRIX-5671 ← **original drive type**
- 14. DGC-RAID5-0206
- 15. EMC-SYMMETRIX-5671 ← **new drive type with a different geometry**

The **format** utility recognizes that the original logical unit and the migrated logical unit have different disk geometries. Therefore, two entries for EMC-SYMMETRIX-5671 appear.

Note: More duplicate drive entries are created with PowerPath Migration Enabler and Open Replicator if you migrate from logical units with different geometries to the same type of EMC-SYMMETRIX-5671 logical units.

Exception to disk-label renaming

If PowerPath Migration Enabler cannot open the target logical unit (for example, because all slices are exclusively opened), the **powermig cleanup** command displays the following message:

```
Unable to edit array name in disk label
```

EMC does not recommend manually editing the disk label. However, you can use the **powerformat** utility to update the label, as described in the next section.

Viewing and updating disk labels with powerformat

On Solaris hosts, a logical unit's disk label contains information about the vendor, product, geometry, and slices. PowerPath Migration Enabler includes a utility called **powerformat** that allows you to safely update disk-label information, preserve partition definitions and data, and make newly available disk capacity available for use.

When you migrate data to a target logical unit that is larger than the source logical unit, **powerformat** can adjust the disk label and partition table to make available the entire disk capacity on the target logical unit. The **powerformat** command takes partitions that extended to the end of the currently labeled logical unit (the label migrated from the source logical unit), and extends them to the end of the target logical unit. This allows you to use the additional space on the target.

The **powerformat** three main operation modes are as follows:

- ◆ Query — Allows you to view the current disk information associated with the specified logical unit. Refer to .
- ◆ Preview — Allows you to view the current disk information associated with the specified logical unit, as well as new disk information calculated by **powerformat**. You can preview any changes before implementing them.
- ◆ Execute — Executes disk-label changes on the specified disk. Refer to .

Query for disk label information

The **powerformat** utility allows you to query the specified logical unit for information about the logical unit and its label. The output includes the logical unit's vendor, product, and revision; and mode sense, disk geometry, and VTOC information.

Run **powerformat <device>** to view disk information for the specified logical unit.

Where *<device>* is the device name whose disk information you want to display.

For example, **powerformat /dev/rdisk/c1t2d0s0**

Output similar to the following appears:

```
Inquiry:  "EMC          ", "SYMMETRIX          ", "5568"
Mode Sense:  pctl 17726, nhead 15, nsect 64, capacity 17016960
GEOM:  pctl 10000, nctl 9998, acyl 2, nhead 10, nsect 32
VTOC:  "DGC-RAID5-0217 cyl 9998 alt 2 hd 10 sec 32"
```

v_part	p_start	p_size
-----	-----	-----
0	0	263040
1	263040	263040
2	0	3199360
6	526080	2673280

The output in this example is a subset of the output in . Refer to the relevant output lines and their corresponding descriptions.

Preview disk information

Preview mode allows you to view the current disk information as well as the new disk information calculated by **powerformat** without implementing any changes. The procedure in this section does not write a new disk label.

Type **powerformat -<n | g | a | w> [-p <partition>] <device>** to preview the new disk information for the specified logical unit.

Arguments

<device> is the device name whose disk information you want to change.

-a installs a new ASCII text label. The ASCII text label includes the vendor, product, and revision; and the number of cylinders, alternate cylinders, heads, and sectors associated with the logical unit. This option does not change the disk geometry.

-g installs a new geometry and partitions based on the label cylinder geometry of the specified logical unit.

The **-g** option uses the label cylinder geometry of the specified logical unit, and increases the number of cylinders to accommodate a larger target logical unit than the label implies. If the maximum number of cylinders allowed on a logical unit is reached before the entire disk is covered, then the maximum number of cylinders is used, and additional space becomes unavailable. When the logical unit is not in use, use the **-n** option to avoid this scenario.

-n tries the native disk geometry before installing a new geometry, as follows:

NOTICE

Do not use the **-n** option while the logical unit is in use. Use the **-g** option instead.

- If the partitions can be preserved using the native geometry, and the native geometry can be represented by the VTOC label, then the native geometry is installed.
- Otherwise, if the label cylinder geometry does not result in wasted disk space, then the label cylinder geometry is installed. In this case, the number of cylinders is increased to cover the whole disk.
- Otherwise, a new geometry that can preserve partitions and cover as much of the disk as possible is invented and installed.
- By default, **powerformat** grows all partitions that extend to the end of the currently labeled logical unit.

-p allows you to specify partitions to preserve. Without this option, all partitions are preserved. When you specify partitions, powerformat selects a new geometry based on the partitions you want to preserve. It then preserves other partitions as possible given the new geometry, and deletes the remaining partitions. To specify multiple partitions, include multiple **-p** options on a single command line.

-w tells **powerformat** to extend only the partitions that were whole disk partitions before the migration. Partitions that start in the middle of the disk and continue to the end are not extended.

Output

```
Inquiry: "EMC      " "SYMMETRIX      " "5568"
Mode Sense: p_cyl 17726 hd 15 sec 64 cap 17016960
GEOM (before): cyl 9998 alt 2 hd 10 sec 32 cap 3200000
VTOC (before): "DGC-RAID5-0217 cyl 9998 alt 2 hd 10 sec 32"
```

v_part	p_start	p_size
-----	-----	-----
0	0	263040
1	263040	263040
2	0	3199360
6	526080	2673280

```
GEOM (after): cyl 53176 alt 2 hd 10 sec 32 cap 17016960
VTOC (after): "EMC-SYMMETRIX-5568 cyl 53176 alt 2 hd 10 sec 32"
```

v_part	p_start	p_size
-----	-----	-----
0	0	263040
1	263040	263040
2	0	17016320
6	526080	16490240

Output explanation

explains the output shown in the previous examples.

Table 11 Output explanation: preview mode (1 of 2)

Output line	Description
Inquiry: "EMC" "SYMMETRIX" "5568"	<p>Inquiry displays the vendor, product, and revision information for the specified logical unit.</p> <p>This information does not derive from the disk label. When you run powerformat on the target logical unit after a migration, the inquiry data for the target displays.</p>
Mode Sense: cyl 17726 hd 15 sec 64 cap 17016960	<p>Mode Sense displays the physical logical unit's mode sense data, including the number of physical cylinders, heads, and sectors; and the disk capacity.</p> <p>This information does not derive from the disk label. When you run powerformat on the target logical unit after a migration, the mode sense data for the target displays.</p>
GEOM (before): cyl 9998 alt 2 hd 10 sec 32 cap 3200000	<p>GEOM displays the number of cylinders, alternate cylinders, heads, and sectors; and the disk capacity.</p> <p>This information derives from the disk label. When you run powerformat on the target logical unit after a migration, the source logical unit's disk geometry displays.</p>
VTOC (before): "DGC-RAID5-0217 cyl 9998 alt 2 hd 10 sec 32"	<p>VTOC displays the VTOC information for the logical unit, including the type of array and the number of cylinders, alternate cylinders, heads, and sectors.</p> <p>This information derives from the disk label. When you run powerformat on the target logical unit after a migration, the source logical unit's VTOC information displays (except in cases where Migration Enabler automatically updates parts of the ASCII name—as in EMC VNX/CLARiiON-to-EMC Symmetrix migrations).</p>
<pre> v_part p_start p_size ----- - 0 0 263040 1 263040 263040 2 0 3199360 6 526080 2673280 </pre>	<p>The partition table for the specified logical unit.</p> <p>v_part — volume partition number</p> <p>p_start — sector number where the partition starts</p> <p>p_size — number of sectors in the partition</p> <p>This information derives from the disk label. When you run powerformat on the target logical unit after a migration, the source logical unit's partition table displays.</p> <p>Note: If the target to which you migrate data is larger than the source, you cannot take advantage of additional space on the target logical unit until you update the target's partition table.</p>

Table 11 Output explanation: preview mode (2 of 2)

Output line	Description
GEOM (after): cyl 53176 alt 2 hd 10 sec 32 cap 17016960	The new geometry for the specified logical unit. This new geometry is installed if you execute the changes. The display includes the number of cylinders, alternate cylinders, heads, and sectors; and the disk capacity.
VTOC (after): "EMC-SYMMETRIX-5568 cyl 53176 alt 2 hd 10 sec 32"	VTOC (after) displays the new VTOC information. The display includes the type of array and the number of cylinders, alternate cylinders, heads, and sectors. In this example, the VNX/CLARiiON logical unit information copied to the target during the migration would be replaced by the target logical unit's Symmetrix information shown here.
<pre> v_part p_start p_size ----- - 0 0 263040 1 263040 263040 2 0 17016320 6 526080 16490240 </pre>	<p>This output shows the new partition information.</p> <p>In this example, the values in the p_size column for partitions 2 and 6 increased compared with the <i>before</i> partition information earlier in this table.</p> <p>This output shows how partitions 2 and 6 would be extended if you installed the new label.</p>

Install a new disk label

Before you install a new disk label, preview the changes as described in [“Preview disk information” on page 79](#). Once you preview the information and decide you want to install a new label, you can execute the changes by including the **-x** option.

Type **powerformat -x <n/g/a/w> -p <partitions> <device>** to install the new disk label. For example,

```
powerformat -xg /dev/rdisk/c1t2d0s0
```

```
powerformat -xg -p 0 -p 6 /dev/rdisk/c1t5006048AD52CB8EDd342s0
```

where:

<device> is the device name whose disk information you want to change.

-xn executes the changes previewed with the **powerformat -n <device>** command. **-xn** tries the native disk geometry before installing a new geometry.

NOTICE

Do not execute this command while the logical unit is in use. If you want to run **powerformat** while the logical unit is in use, then use the **-g** option instead of **-n**.

-xg executes the changes previewed with the **powerformat -g <device>** command. **-xg** installs the new disk-label geometry and the current and new partition tables.

-xa executes the changes previewed with the **powerformat -a <device>** command. **-xa** installs the new ASCII disk label.

Note: Refer to [“Preview disk information” on page 79](#) for more information on these options.

-w tells powerformat to extend only the partitions that were whole-disk partitions before the migration. Partitions that start in the middle of the disk and continue to the end are not extended.

-p allows you to specify particular partitions to preserve. Without this option, all partitions are preserved.

Output

```
powerformat -xg /dev/rdisk/c1t2d0s0
Inquiry: "EMC      " "SYMMETRIX      " "5568"
Mode Sense: pcyll 17726 hd 15 sec 64 cap 17016960
GEOM (before): cyl 9998 alt 2 hd 10 sec 32 cap 3200000
VTOC (before): "DGC-RAID5-0217 cyl 9998 alt 2 hd 10 sec 32"
```

v_part	p_start	p_size
-----	-----	-----
0	0	263040
1	263040	263040
2	0	3199360
6	526080	b2673280

```
GEOM (after): cyl 53176 alt 2 hd 10 sec 32 cap 17016960
VTOC (after): "EMC-SYMMETRIX-5568 cyl 53176 alt 2 hd 10 sec 32"
```

v_part	p_start	p_size
-----	-----	-----
0	0	263040
1	263040	263040
2	0	17016320
6	526080	16490240

```
powerformat -xg -p 1 -p 2 /dev/rdisk/emcpower17c
Inquiry: "EMC      " "SYMMETRIX      " "5671"
Mode Sense: pcyll 25314 hd 15 sec 64 cap 24301440
GEOM (before): cyl 32766 alt 2 hd 8 sec 16 cap 4194304
VTOC (before): "EMC-SYMMETRIX-5671 cyl 32766 alt 2 hd 8 sec 16"
```

v_part	p_start	p_size
--------	---------	--------

```

-----
      0          0          131072
      1        131072        262144
      2          0          4194048
      6        393216        3800832
Using max dkg_ncyl. Wasting capacity.
GEOM (after): cyl 65535 alt 2 hd 8 sec 16 cap 8388736
VTOC (after): "EMC-SYMMETRIX-5671 cyl 65535 alt 2 hd 8 sec 16"
  v_part    p_start    p_size
-----
      0          0          131072
      1        131072        262144
      2          0          8388480
      6        393216        7995264
Output
Inquiry: "EMC          " "SYMMETRIX          " "5771"
Mode Sense: pcyl 1916 hd 15 sec 128 cap 3678720
GEOM (before): cyl 2078 alt 2 hd 15 sec 64 cap 1996800
VTOC (before): "EMC-SYMMETRIX-5267 cyl 2078 alt 2 hd 15 sec 64"
  v_part    p_start    p_size
-----
      0          0          205440
      2          0          1994880
      4        205440          205440
      5        410880          205440
      6        616320          1024320

GEOM (after): cyl 3830 alt 2 hd 15 sec 64 cap 3678720
VTOC (after): "EMC-SYMMETRIX-5771 cyl 3830 alt 2 hd 15 sec 64"

```

```

  v_part    p_start    p_size
-----
      0          0          205440
      2          0          3676800
      4        205440          205440
      5        410880          205440
      6        616320          1024320

```

INDEX

- A**
 - all option 30
 - audit log message 65

- C**
 - ceiling copy rate 30
 - changing the remote SE server environment variables 33
 - CLARiiON and VNX to Symmetrix migration 59
 - cleanup migration state 20
 - clusters
 - AIX HACMP 40
 - MSCS (Microsoft Cluster Server) 40
 - Solaris Cluster 40
 - VCS (Veritas Cluster Server) 40
 - command support summary table 73
 - commands, Migration Enabler 29
 - committed migration state 16, 19
 - committedAndRedirected migration state 16

- D**
 - disk label
 - installing 82
 - overview 75
 - previewing 79
 - querying 77, 78

- E**
 - EMC ControlCenter post-migration steps 25
 - EMC MirrorView prerequisites 21
 - EMC Online Support website 7
 - emcgrab utility 66
 - Encapsulation technology type
 - licensing 12
 - migrations 54
 - overview 13
 - encrypted devices 54
 - Encryption with RSA target logical unit prerequisites 21
 - environment variables, changing the remote SE server 33
 - error log message 65
 - errors or faults 68
 - execute operation, powerformat utility 78

- F**
 - file option 30

- G**
 - gkselect selection file 23
 - graceful host reboot (Host Copy) 66

- H**
 - HACMP 43
 - Host Copy technology type 13
 - graceful host reboot 66
 - licensing 12
 - migrations 51
 - throttle settings 53
 - host crash troubleshooting 66
 - hostcopy_ceiling 30
 - hostcopy_ceiling settings 53

- L**
 - labeling, disk 75
 - licensing Migration Enabler 12
 - Logical Disk Manager (LDM) 35
 - Logical Volume Manager (LVM) 36

- M**
 - Microsoft Cluster Server 40
 - Migration Enabler
 - commands 29
 - licensing 12
 - migration prerequisites 21
 - overview 11
 - post-migration steps 25
 - support summary table 73

migration speed, tuning 31
 migration states
 cleanup 20
 committed 16, 19
 committedAndRedirected 16
 overview 14
 paused 16
 setup 16, 17
 sourceSelected 16, 18
 syncing 16, 17, 25
 targetSelected 16, 19
 migration synchronization speed 31
 migration workflow 15

N

native devices, defined 12

O

Open Replicator technology type
 licensing 12
 migrations 58
 overview 14
 prerequisites 21, 22, 23
 Symmetrix-to-Symmetrix migrations 58
 throttle settings 61
 OR. *See also* Open Replicator 58

P

paused migration state 16
 post-migration steps
 EMC ControlCenter 25
 Migration Enabler 25
 powerformat utility
 execute operation 78
 overview 34
 preview operation 78, 79
 query operation 77, 78
 updating disk labels 77
 viewing disk labels 77
 powermig abort command 29
 powermig cleanup command 20, 29
 powermig commands and migration states 15
 powermig commit command 16, 29
 powermig getHandle command 29
 powermig help command 29
 powermig info command 29

powermig options command 30, 32
 powermig pause command 16, 30
 powermig query command 30
 powermig recover command 30
 powermig resume command 30
 powermig selectSource command 16, 30
 powermig selectTarget command 16, 19, 30
 powermig setup command 16, 17, 30
 powermig sync command 16, 17, 30
 powermig throttle command 30, 31
 powermig undoRedirect command 16, 30
 powermig version command 30
 PPME_REMOTE_SE_PORT environment variable 34
 PPME_REMOTE_SE_SERVER environment variable 34
 prerequisites, Migration Enabler migration 21
 preview operation, powerformat utility 78, 79
 problems, reporting 66
 process hang, resolving 69
 pseudo devices
 defined 12
 renaming 49

Q

query operation, powerformat utility 78
 querying disk label 77, 78

R

reclaim unused space for encrypted data 28
 reclaim unused space for plaintext data 25
 renaming, pseudo devices 49
 reported faults 67
 reporting problems 66
 resolving a process hang 69

S

set pace argument 61
 setup migration state 16, 17
 simultaneous migration sessions, managing 30
 Solaris disk labeling 75
 Solaris host prerequisites 22
 Solutions Enabler (SE) Thin Client 21, 32
 Solutions Enabler symrcopy command 61
 sourceSelected migration state 16, 18
 Sun Cluster with VxVM 48
 Symclone. *See also* TimeFinder/Clone 62

- Symmetrix Quality of Service (QoS) value,
 - TimeFinder/Clone 63
- Symmetrix Remote Data Facility (SRDF) prerequisites 21
- Symmetrix-to-Symmetrix migration with Open Replicator 58
- synchronization speed, changing 31
- syncing migration state 16, 17, 25

T

- target device prerequisites 21
- targetSelected migration state 16, 19
- technology type
 - Encapsulation 13, 54
 - Host Copy 13, 51
 - Open Replicator 14, 58
 - TimeFinder/Clone 14, 61
- technology type licensing 12
- thin devices (virtually provisioned devices) 32
- throttle settings
 - Host Copy technology type 53
 - TimeFinder/Clone technology type 63
- TimeFinder/Clone technology type 14
 - licensing 12
 - migrations 61
 - prerequisites 21, 23
 - QoS (Quality of Service) values 63
 - throttle settings 63
 - throttle value 63
- troubleshooting
 - host crashes 66
 - migrations 65
- tuning migration speed 31

U

- unused space
 - reclaim for encrypted data 28
 - reclaim for plaintext data 25
- updating disk labels with powerformat 77

V

- Veritas Cluster Server (VCS) 45
- Veritas Volume Manager (VxVM) 35
- viewing disk labels with powerformat 77
- virtually provisioned devices (thin devices) 32
- VNX/CLARiiON to Symmetrix migration 59

W

- Windows host prerequisites 22

