



# EMC<sup>®</sup> PowerPath<sup>®</sup> Migration Enabler

Version 6.x

## User Guide

P/N 302-001-317

REV 03



EMC<sup>2</sup>

Copyright © 2006 - 2017 EMC Corporation. All rights reserved. Published in the USA.

Published July 2017

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.

EMC<sup>2</sup>, 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.....	12
Supported migration scenarios.....	12
Migrations using pseudo and native devices.....	13
Installing and licensing Migration Enabler.....	13
Understanding the Migration Enabler technology types.....	14
HostCopy.....	14
Open Replicator.....	15
Understanding the Migration Enabler migration workflow .....	15
Migration state descriptions .....	17

## Chapter 2 Migration Procedures

Before you start a migration.....	24
Perform a migration .....	25
After you finish a migration.....	26
Reclaiming unused space on the target.....	27

## Chapter 3 Migration Enabler Features

Migration Enabler commands .....	32
Migrating to devices in a Geo Cluster setup.....	33
Managing simultaneous migration sessions using -all and -file .....	33
Changing the synchronization speed for migrations .....	34
Specifying a HostCopy upper limit copy rate.....	34
powermig options command.....	35
Migrating to and from virtually provisioned devices .....	35
Migrating to and from devices with 4k sector size .....	35
Using the Solutions Enabler Thin Client for Migration Enabler .....	36
Changing the remote Solutions Enabler server.....	37
Viewing and updating disk labels with powerformat .....	38

## Chapter 4 Volume Managers and Cluster Environments

Migrations in a volume manager environment.....	40
---	----

Migrations in a Veritas Volume Manager (VxVM) environment..	40
Migrations in a Logical Volume Manager (LVM) environment....	40
Migrating a native-named device under VxVM control .....	40
Migrating devices under LVM control on UNIX and Linux hosts .	41
Dynamic disk support with Migration Enabler .....	44
Migrations in a cluster environment.....	44
Migrate in MSCS cluster environment .....	45
Migrating Cluster Disks in MSCS environment with multiple active nodes .....	45
Migrating CSV in MSCS environment with multiple active nodes ...	46
Migrating Cluster Disks and CSV in MSCS cluster environment with single active node.....	47
Migrating in HACMP environment (AIX hosts) .....	48
Migrating between pseudo-named devices in VCS environment	50
Migrating between pseudo-named devices in Sun Cluster with VxVM .....	53
Renaming a pseudo device .....	54

**Appendix A Migration Procedure Examples**

HostCopy migrations .....	58
HostCopy procedure .....	58
HostCopy throttle settings .....	59
HostCopy ceiling settings.....	60
Open Replicator migrations .....	60
Set up Symmetrix-to-Symmetrix migrations with Open Replicator .	60
Migrate EMC VNX or CLARiiON data to EMC Symmetrix arrays...	61
Open Replicator procedure .....	62
Understanding Open Replicator throttle settings.....	63

**Appendix B Troubleshooting a Migration**

Audit and error log messages .....	66
Audit log message example .....	66
Error log message example .....	66
Reporting a problem .....	66
Migration impact from a graceful host reboot.....	67
Faults after a host crash.....	67
Other reported faults .....	69
Resolving a process hang .....	69
Troubleshooting cluster migrations .....	70
Cluster migration best practices.....	71

	Common terminologies.....	72
<b>Appendix C</b>	<b>Support Summary</b>	
	Migration Enabler technology type support.....	74
<b>Appendix D</b>	<b>Disk Labeling powerformat Command for Solaris</b>	
	Solaris and PowerPath Migration Enabler disk labeling .....	76
	Multiple array type entries in the format utility output.....	76
	Viewing and updating disk labels with powerformat .....	78
<b>Index</b>		

## 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 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
03	July 20, 2017	Removed the powermigcl command and SYMCLONE technology type references.
02	March 2016	Added migrating to devices in a Geo Cluster setup content and migrating to devices with 4k sector size
01	December 2014	First release of the EMC PowerPath Migration Enabler Version 6.0

## 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:

<b>Normal</b>	Used in running (nonprocedural) text for: <ul style="list-style-type: none"> <li>Names of interface elements, such as names of windows, dialog boxes, buttons, fields, and menus</li> <li>Names of resources, attributes, pools, Boolean expressions, buttons, DQL statements, keywords, clauses, environment variables, functions, and utilities</li> <li>URLs, pathnames, filenames, directory names, computer names, links, groups, service keys, file systems, and notifications</li> </ul>
<b>Bold</b>	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"> <li>Names of interface elements, such as names of windows, dialog boxes, buttons, fields, and menus</li> <li>What the user specifically selects, clicks, presses, or types</li> </ul>
<i>Italic</i>	Used in all text (including procedures) for: <ul style="list-style-type: none"> <li>Full titles of publications referenced in text</li> <li>Emphasis, for example, a new term</li> <li>Variables</li> </ul>
<b>Courier</b>	Used for: <ul style="list-style-type: none"> <li>System output, such as an error message or script</li> <li>URLs, complete paths, filenames, prompts, and syntax when shown outside of running text</li> </ul>
<b>Courier bold</b>	Used for specific user input, such as commands
<i>Courier italic</i>	Used in procedures for: <ul style="list-style-type: none"> <li>Variables on the command line</li> <li>User input variables</li> </ul>
< >	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](mailto: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 ..... 12
- ◆ Installing and licensing Migration Enabler ..... 13
- ◆ Understanding the Migration Enabler technology types ..... 14
- ◆ Understanding the Migration Enabler migration workflow ..... 15

## 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 following table describes supported migration scenarios for PowerPath Migration Enabler.

### NOTICE

PowerPath Migration Enabler does not support migrating to smaller target devices. All migrations on HP-UX are disruptive.

**Table 2** Supported migration scenarios (1 of 2)

Source logical unit	Target logical unit	PowerPath platform supported	Technology	Notes
Native-device name <sup>1, 2</sup>	Native-device name	HP-UX, Solaris	HostCopy <sup>3</sup> Open Replicator	Open Replicator requires Solutions Enabler.
Native-device name	Pseudo-device name	Linux, Solaris		
Pseudo-device name <sup>4</sup>	Pseudo-device name	AIX, HP-UX, Linux, Solaris, Windows		

**Table 2** Supported migration scenarios (2 of 2)

Source logical unit	Target logical unit	PowerPath platform supported	Technology	Notes
Thin device	Thin device	AIX, HP-UX, Linux, Solaris, Windows	HostCopy Open Replicator	Open Replicator requires Solution Enabler. Migration Enabler detects thin devices by default. When using HostCopy, Migration Enabler preserves the thin device allocation savings. When using Open Replicator, refer to the Solutions Enabler documentation concerning support for thin devices.
Thin device	Thick device	AIX, HP-UX, Linux, Solaris, Windows		
Thick device	Thick device	Linux, Windows		
Thick device <sup>5</sup>	Thin device	AIX, HP-UX, Linux, Solaris, Windows		

1. Also called Legacy migrations on PowerPath for HP-UX. Legacy migrations are supported only on HP-UX 11i v2 (11.23).
2. On PowerPath for Linux, supported on SLES 10 SP4 only.
3. HostCopy migrations are independent of the Symmetrix, CLARiiON, or other operating environments.
4. Also called Agile migrations on PowerPath for HP-UX. Agile migrations are supported only on HP-UX 11i v3 (11.31).
5. The space efficiency of migrating from a thick to a thin device depends strictly on the amount of contiguous zero blocks on the thick LUN. To maximize the zero blocks, use external tools such as 'SDelete' for Windows before beginning a migration.

## 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.

The PowerPath multipathing license includes all Migration Enabler functionalities including all Migration Enabler technologies (HostCopy, Open Replicator). PowerPath Migration Enabler license keys are available at no cost to new and existing PowerPath Multipathing customers.

Migration Enabler requires a permanent license. You cannot use a temporary license with Migration Enabler.

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.

---

## HostCopy

PowerPath Migration Enabler HostCopy is a host-based migration tool that uses host resources to move data non-disruptively from a source to a target. HostCopy 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)

## 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.

---

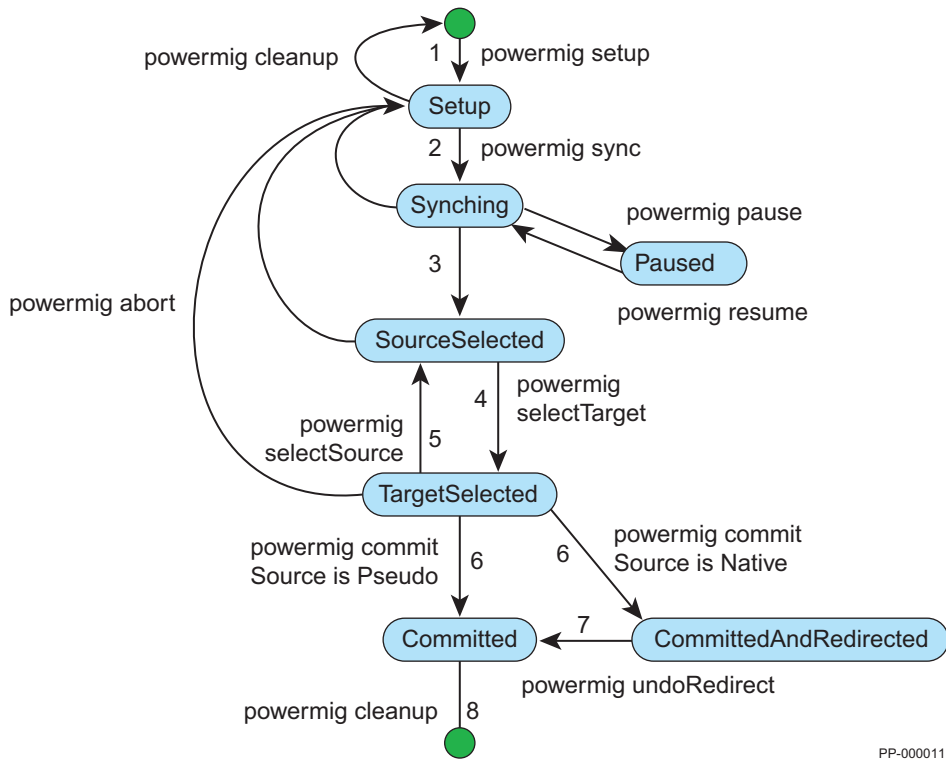
**Note:** Source from EMC Symmetrix, VNX, or CLARiiON cannot be in the same array.

---

## 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 16 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 3 describes the migration states.

**Table 3** Migration state descriptions

State	Description
Setup	A migration is in the <b>Setup</b> state after a <b>powermig setup</b> 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 <b>Syncing</b> state after running the <b>powermig sync</b> 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 <b>syncing [PAUSED]</b> state after running the <b>powermig pause</b> command. Bulk data copy and synchronization of the source and target logical units are paused. Run <b>powermig resume</b> to restart a migration that was paused.
SourceSelected	The migration enters the <b>SourceSelected</b> state once the bulk data copy of the source and target logical units completes. While the migration is in the <b>SourceSelected</b> state, the source and target logical units continue to synchronize, and I/O read requests are processed from the source logical unit.
TargetSelected <sup>1</sup>	The migration transitions to the <b>TargetSelected</b> state after the <b>powermig selectTarget</b> command completes. While the migration is in the <b>TargetSelected</b> 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 <b>powermig commit</b> command. While in this state, the application using the native source device must be brought down and the user must run the <b>powermig undoRedirect</b> command. The application can then be reconfigured and brought back up to use the target device. The migration will now be in the <b>Committed</b> state.
Committed	If using a pseudo source device, this state is entered upon issuing the <b>powermig commit</b> 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 <b>powermig undoRedirect</b> command (as described above in the <b>CommittedAndRedirected</b> state). In the <b>Committed</b> 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 for MSCS cluster migrations.

## 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.

When migrating Cluster Disk 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. For migrating Cluster Shared Volumes, the source should be the Cluster Shared Volume 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).

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 use 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 HostCopy migration synchronizations by setting the HostCopy 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:** When 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** 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 ..... 24
- ◆ Perform a migration..... 25
- ◆ After you finish a migration ..... 26
- ◆ Reclaiming unused space on the target ..... 27

## 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.
- ◆ If using the Solutions Enabler remote-only client with Open Replicator, 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.
- ◆ If you plan to use the `exclude_nodes` option, then do not connect the source and target devices on excluded nodes.
- ◆ 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):
  - For cluster disk migrations, ensure that PowerPath Migration Enabler cluster resource is configured for source cluster disk.



- 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 Enginuity 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 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.

## 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 mscd|mscsv] [-techHost <hostname>
[-techPort <port#>]] [-exclude_nodes <nodename1> <nodename2>]
[-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. Skip this step while migrating physical disks in an MSCS cluster environment.

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

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> [-format] [-force] [-noPrompt|-no]
```

## 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 [“Reclaiming unused space on the target” on page 27](#) to make unused space accessible.
- ◆ If you have EMC ControlCenter<sup>®</sup>, 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.

## Reclaiming unused space on the target

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 25](#).
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/emcpower*x*** 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.



# CHAPTER 3

## Migration Enabler Features

This chapter provides information about the following topics:

- ◆ Migration Enabler commands ..... 32
- ◆ Migrating to devices in a Geo Cluster setup ..... 33
- ◆ Managing simultaneous migration sessions using -all and -file..... 33
- ◆ Changing the synchronization speed for migrations..... 34
- ◆ Specifying a HostCopy upper limit copy rate ..... 34
- ◆ Migrating to and from virtually provisioned devices ..... 35
- ◆ Migrating to and from devices with 4k sector size ..... 35
- ◆ Using the Solutions Enabler Thin Client for Migration Enabler ..... 36
- ◆ Viewing and updating disk labels with powerformat ..... 38

## Migration Enabler commands

PowerPath Migration Enabler uses a common set of commands to support HostCopy and Open Replicator migration technology types. [Table 4 on page 32](#) 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 4** PowerPath Migration Enabler commands (1 of 2)

powermig command	Description
abort	Aborts a migration and returns the migration session to <b>Setup</b> 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 <b>powermig</b> command.
info	Displays information about a specific migration session or about all active migrations.
options	Used to set values that are global to all migrations, or all migrations that use a specific technology. For example, the option, <b>hostcopy_ceiling</b> , can be used to specify an upper limit on the aggregate rate of copying for all HostCopy migrations. When executed with no option specified, powermig options displays all available options and their values. Currently, <b>hostcopy_ceiling</b> is the only available option.
pause	Suspends a migration in the <b>Syncing</b> 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 <b>powermig pause</b> 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 <sup>1</sup>	Designates the target logical unit as the recipient of all I/O Read requests. Continues to send writes to both source and target.



**Table 4** PowerPath Migration Enabler commands (2 of 2)

powermig command	Description
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.

## Migrating to devices in a Geo Cluster setup

Beginning with PowerPath for Windows 6.0 SP1, you can perform migrations in Geo Cluster (multi-site cluster) setup where in not all nodes have access to all the devices. In earlier PowerPath releases, all nodes needed to have access to all devices for the migration to succeed.

Using the **exclude\_nodes** option, you can exclude nodes that do not have access to the source and target device when in a Geo Cluster setup. Migration Enabler will not look for any status or updates from excluded nodes and the nodes that are excluded cannot have migration devices connected during a migration. You cannot exclude the migration node.

You will receive an error if the node cannot be excluded or if the excluded node is not present in the cluster or is not reachable.

Migration Enabler supports only SQL AlwaysOn replication technology in a Geo Cluster setup.

## 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.

## 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 HostCopy* white paper available on EMC Online Support provides additional details about throttle impact. This command does not apply to Encapsulation migrations.

[“Specifying a HostCopy upper limit copy rate” on page 34](#) 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**.

---

EMC VMAX3 does not support throttle.

## Specifying a HostCopy upper limit copy rate

The HostCopy ceiling option lets you specify an upper limit on the aggregate rate of copying for all HostCopy migrations in the **Syncing** state on a host. In turn, this limits the I/O for all HostCopy migrations. With HostCopy 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 HostCopy technology type. Running **powermig options** with no option merely displays all available options and their values.

---

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

---

## 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 HostCopy 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.

HostCopy 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. HostCopy cannot determine whether or not a LUN is thin on a third-party array.

## Migrating to and from devices with 4k sector size

Beginning with PowerPath for Windows 6.0 SP1, migrations of devices 4k sector size is supported.

When migrating LUNs the source and target disk must have the same block size (512 bytes or 4096 bytes).

## 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 37](#). 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:

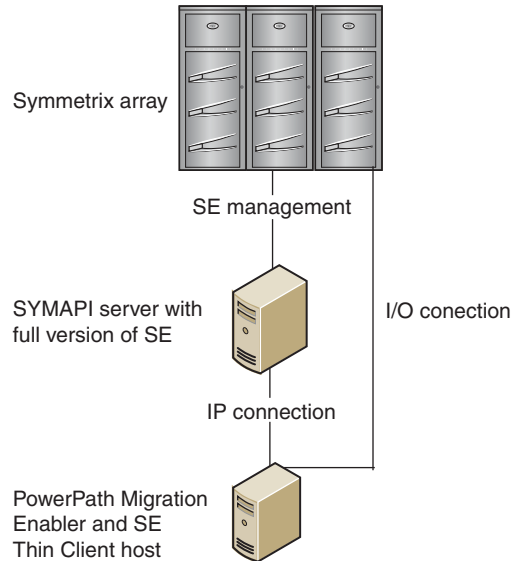
---

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

---

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.

# 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 .....](#) 40
- ◆ [Migrations in a cluster environment .....](#) 44

## 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 [“Migrating a native-named device under VxVM control” on page 40](#).

## 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 [“Migrating devices under LVM control on UNIX and Linux hosts” on page 41](#).

## Migrating 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 **vxvg 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 **vxvg import** *<disk\_group\_name>*
  - b. Type **vxvol -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 HostCopy 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 **vxdisk resize** to resize the logical unit that is under VxVM control.

## Migrating 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/rds/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 HostCopy migration with I/O running on iSCSI devices, `powermig cleanup -format` takes substantially longer since it performs a full format of the disk.

---

## Migrating 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 cleanup of 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 q1a2xxx            sdag      SP A0   active alive      0      0
1 q1a2xxx            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.

## Dynamic disk support with Migration Enabler

When using dynamic disks with Migration Enabler, you can:

- ◆ Migrate all of the members of a dynamic disk volume simultaneously
- ◆ Migrate some members of a dynamic disk volume, commit and cleanup the migration, and then migrate the rest of the members at a later time

While preparing to migrate, allocate the same number of target disks (preferably of the same size to avoid resizing afterwards) on the target array. Pair the target disks and then create a migration session per disk pair.

## Migrations in a cluster environment

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 as there is an approach for migration with multiple active nodes where you do not have to shutdown any nodes in the cluster to perform a cluster migration.

The *EMC E-Lab Interoperability Navigator™* provides additional information on supported cluster versions.

### 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, there are two approaches:

- ◆ Migration with a single active node (Windows Server 2003 and later)  
Where an *active node* is that the node's status is Up.
- ◆ Migration with multiple active nodes (Windows Server 2008 and later that are supported at GA. Windows Server 2003/2003 R2 requires Request for Product Qualification (RPQ)).

---

**Note:** Starting with Windows Server 2016, cluster nodes can be put into a special state called Quarantined. When nodes are in a Quarantined state, PowerPath Migration Enabler clustered migrations that are running in the cluster could return a target fault. Depending on the state of migration, a user intervention may be needed. This behavior occurs because the cluster resource communication needed for PowerPath Migration Enabler to function is hampered temporarily or permanently by the duration and result of the quarantine.

---

## Migrating Cluster Disks in MSCS environment with multiple active nodes

- ◆ [“Common terminologies” on page 72](#) describes some common terminologies associated with MSCS cluster migrations.
- ◆ [“Before you start a migration” on page 24](#) provides prerequisites for migrating physical disks in an MSCS cluster environment.
- ◆ [“Cluster migration best practices” on page 71](#) 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 70](#) provides more information on some of the common error messages and suggested actions while migrating physical disks in an MSCS cluster environment.

## Procedure

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 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 the **powermig setup -cluster** or **powermig setup -cluster mscd** command. 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 failback the resources to the migration node before moving to step 9.
9. Run the **powermig cleanup** command to clean up the migration.

Cluster groups can be safely unconfigured as needed after all migrations are cleaned up.

## Migrating CSV in MSCS environment with multiple active nodes

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.

This procedure is supported only with all migration technologies except InVista Encapsulation.

- ◆ [“Cluster migration best practices” on page 71](#) 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 70](#) provides more information on some of the common error messages and suggested actions while migrating physical disks in an MSCS cluster environment.

## Procedure

1. Ensure the target disk is initialized and accessible from all nodes.
2. Ensure that the source is a Cluster Shared Volume with a corresponding physical disk resource in the cluster and the target device is a non-cluster disk.
3. Ensure the underlying technology is configured, and then run the **powermig setup -cluster mscsv** command. The node where you run **powermig setup** is not required to own CSV resource.

The **powermig setup** command automatically creates the Migration Enabler cluster resource and assigns it to the cluster group **PPME Control Group**.

4. Run the **powermig sync** command to begin syncing the migrations.
5. When the synchronization is complete, type **powermig info -all** command to know the status of the migration.
6. In case of failover of CSV the resource during migration Migration Enabler could report target device fault. In this case abort the migration and restart migration. You do not need to failback CSV resource before restarting the migration after abort.
7. Run the **powermig commit** command to commit the migration.
8. Run the **powermig cleanup** command to clean up the migration.

On **powermig cleanup**, the Migration Enabler cluster resource is automatically removed.

## Migrating Cluster Disks and CSV in MSCS cluster environment with single active node

When migrating Cluster Disk or CSV you can use the approach for Migration with single active node and select one node to remain active during the migration. All failover groups must be moved to the active node.

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.

## Migrating 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.

## Migrating 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 HostCopy 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. The Veritas Cluster Server documentation 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 HostCopy 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

```

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

```

# hagrps -freeze service_group -persistent

```

7. 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.

8. Verify that emcpower names have swapped between source and target devices. Repeat [Step 6 on page 51](#) and [Step 7 on page 52](#) to verify the changes. If I/O was running to the disks during the migration, verify that there was no interruption or data corruption.
9. 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
```

10. Import disks labeled as **udid\_mismatch** as cloned disks by assigning a new disk group name:

```
# vxvg -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
```

11. Type **hagrp -unfreeze service\_group -persistent** to unfreeze the service groups that were previously set to freeze state in [Step 6 on page 51](#).
12. Use the **Cluster Management Console** or type **hagrp -switch service\_group -to linux116** 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
```

## Migrating 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 HostCopy 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 [“Renaming a pseudo device” on page 54](#).
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).

## Renaming 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:

- ◆ [HostCopy migrations](#) ..... 58
- ◆ [Open Replicator migrations](#) ..... 60

## HostCopy migrations

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

### HostCopy 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 mscd|mcsrv] [-techHost <hostname> [-techPort
<port#>]] [-exclude_nodes <nodename1> <nodename2>] [-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.

---

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 34](#)
- ◆ [“After you finish a migration” on page 26](#)

## HostCopy throttle settings

The HostCopy 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 HostCopy is 2.

**Table 5** HostCopy 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%

HostCopy throttle settings are ignored when you enable the HostCopy ceiling function.

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

---

**Note:** You can put as many HostCopy migrations into the **Syncing** state as desired; however only eight will be actively syncing at any given time. In HostCopy, 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, HostCopy 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.

---

## HostCopy ceiling settings

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

## 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 mscd|mscsv] [-techHost <hostname> [-techPort
<port#>]] [-exclude_nodes <nodename1> <nodename2>] [-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.

---

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 34](#)
- ◆ [“After you finish a migration” on page 26](#)

## 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 **symcopy** command with the **set pace** argument. A lower throttle value can significantly increase resource consumption on the target logical unit’s array.

## Related topics

- ◆ [“Changing the synchronization speed for migrations” on page 34](#)
- ◆ [“Using the Solutions Enabler Thin Client for Migration Enabler” on page 36](#)

- ◆ [“After you finish a migration” on page 26](#)



# APPENDIX B

## Troubleshooting a Migration

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

- ◆ [Audit and error log messages](#) ..... 66
- ◆ [Reporting a problem](#)..... 66
- ◆ [Migration impact from a graceful host reboot](#) ..... 67
- ◆ [Faults after a host crash](#) ..... 67
- ◆ [Other reported faults](#)..... 69
- ◆ [Resolving a process hang](#) ..... 69
- ◆ [Troubleshooting cluster migrations](#) ..... 70

## 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
```

The following audit log entry indicates that a **powermig setup** command was executed.

```
1/28/2016:9:40:09 PM:PPME: Info: handle=c2, event=stateChanged,
    state=setup, cmd=setup, tech=HostCopy, src=harddisk11
    (000195702076:57D4), tgt=harddisk25 (000197200145:01A1),
    clus=PPME BY2_E$_SQLDump1, exclude_nodes_mask=0xc
```

### 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 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 HostCopy, 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, 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, HostCopy 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, HostCopy 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 6 on page 68 shows the three states when Logical Unit faults can be reported.

**Table 6** 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	<b>sync or cleanup</b>	No
SourceSelected	Yes	TargetLUFault	Returns to Setup state	<b>sync or cleanup</b>	No
TargetSelected	Yes	SourceLUFault	pseudo devices: moves to Committed state	<b>cleanup</b>	No
			native devices: moves to Committed and Redirected state	<b>undoRedirect</b>	

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 7 on page 69](#) shows the three states when Logical Unit faults are reported.

**Table 7** Other reported faults

Current state	Source device failure	Target device failure	Fault state reported	Cause	powermig command	Data loss
Syncing	Yes	N/A	SourceLUfault <sup>1</sup>	I/O write to source fails	abort <sup>2</sup>	No
		Yes	TargetLUfault <sup>1</sup>	I/O write to target fails	abort	No
SourceSelected	Yes	N/A	SourceLUfault <sup>1</sup>	I/O write to source fails	selectTarget <sup>2</sup>	No
		Yes	TargetLUfault <sup>1</sup>	I/O write to target fails	abort	No
TargetSelected	Yes	N/A	SourceLUfault <sup>1</sup>	I/O write to source fails	commit	No
		Yes	TargetLUfault <sup>1</sup>	I/O write to target fails	abort <sup>2</sup>	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 8 on page 70](#) lists common error messages and suggested actions for physical disk migrations in an MSCS cluster environment.

**Table 8** 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 <b>powermig recover</b> 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.
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 <b>powermig abort</b> command to abort the migration.
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 <b>-cluster</b> to setup a cluster migration).	Cluster disk can be used as a source device for migration only if <b>-cluster</b> 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.
Error 98 Source must be a cluster shared volume	<b>-cluster mscsv</b> could be used only for migrating cluster shared volumes.

**Table 8** Cluster migration error messages

Message	Suggested action
Error 99 Unable to remove PPME cluster resource	Migration Enabler was unable to remove the Migration Enabler cluster resource. Run <b>powermig recover</b> .
Error 100 Unable to create PPME cluster resource for the source device.	Migration Enabler failed to create the Migration Enabler cluster resource for migration.
Error 101 Unable to register PPME cluster resource type.	Failed to register the Migration Enabler cluster resource type.
Error 103 CSV being migrated already has PPME cluster resource	The existing Migration Enabler cluster resource was found for the CSV being migrated.
Error 105 Some of cluster nodes are not in 'up' state	Some nodes do not respond in timely fashion. Check that all nodes are up and running. Check that cluster service is running on all the nodes.
Error 106 Some nodes have old version of PowerPath installed	Some node has an older PowerPath version installed than the current node. Install the same version of PowerPath on all nodes.
Error 107 Unable to communicate with some nodes	Some nodes do not respond in timely fashion. Check that all nodes are up and running. Check that cluster service is running on all the nodes.
Error 109 Source disk type was changed since migration was created	The type of source cluster resource was changed or it is impossible to check the type of source cluster resource. Move source cluster resource (Cluster Disk or CSV) to online state.
Error 110 Source disk's dependent PPME resource was changed	It was detected that the Migration Enabler cluster resource for the migration was changed or renamed. Restore the original Migration Enabler cluster resource.
Error 111 Target device is not attached to some nodes	It was detected that target device is not attached to some nodes. Attach target device to all nodes and retry.
Error 112 Source is a CSV (use '-cluster mscsv' to setup a cluster migration)	The -cluster option is used to setup a migration for Cluster Disk. To setup a migration for CSV use <b>-cluster mscsv</b> .

## 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.
- ◆ Do not add new cluster nodes during CSV migration as it could cause data corruption.

## Common terminologies

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

- ◆ **Active node** - Node's status in the cluster is Up.
- ◆ **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 9 on page 74](#) lists for all versions of PowerPath Migration Enabler the commands supported with each platform and technology type.

**Table 9** Migration Enabler technology type support

Commands	Supported technology types
	Version 6.0, 6.1, 6.2
powermig abort	HOSTCOPY, OR
powermig cleanup	
powermig commit	
powermig getHandle	
powermig help	
powermig info	
powermig options	HOSTCOPY
powermig pause	HOSTCOPY, OR
powermig query	
powermig recover	
powermig resume	
powermig selectSource	
powermig selectTarget	
powermig setup	
powermig sync	
powermig throttle	
powermig version	

# 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-RAID5-0206 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:  pcy1 17726, nhead 15, nsect 64, capacity 17016960
GEOM:  pcy1 10000, ncy1 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 10** Output explanation: preview mode (1 of 2)

Output line	Description
Inquiry: "EMC" "SYMMETRIX" "5568"	<b>Inquiry</b> displays the vendor, product, and revision information for the specified logical unit. This information does not derive from the disk label. When you run <b>powerformat</b> on the target logical unit after a migration, the inquiry data for the target displays.
Mode Sense: cyl 17726 hd 15 sec 64 cap 17016960	<b>Mode Sense</b> displays the physical logical unit's mode sense data, including the number of physical cylinders, heads, and sectors; and the disk capacity. This information does not derive from the disk label. When you run <b>powerformat</b> on the target logical unit after a migration, the mode sense data for the target displays.
GEOM (before): cyl 9998 alt 2 hd 10 sec 32 cap 3200000	<b>GEOM</b> displays the number of cylinders, alternate cylinders, heads, and sectors; and the disk capacity. This information derives from the disk label. When you run <b>powerformat</b> on the target logical unit after a migration, the source logical unit's disk geometry displays.

**Table 10** Output explanation: preview mode (2 of 2)

Output line	Description
VTOC (before): "DGC-RAID5-0217 cyl 9998 alt 2 hd 10 sec 32"	<p><b>VTOC</b> 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 <b>powerformat</b> 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><b>v_part</b> — volume partition number</p> <p><b>p_start</b> — sector number where the partition starts</p> <p><b>p_size</b> — number of sectors in the partition</p> <p>This information derives from the disk label. When you run <b>powerformat</b> on the target logical unit after a migration, the source logical unit's partition table displays.</p> <p><b>Note:</b> 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>
GEOM (after): cyl 53176 alt 2 hd 10 sec 32 cap 17016960	<p>The new geometry for the specified logical unit. This new geometry is installed if you execute the changes.</p> <p>The display includes the number of cylinders, alternate cylinders, heads, and sectors; and the disk capacity.</p>
VTOC (after): "EMC-SYMMETRIX-5568 cyl 53176 alt 2 hd 10 sec 32"	<p><b>VTOC (after)</b> displays the new VTOC information. The display includes the type of array and the number of cylinders, alternate cylinders, heads, and sectors.</p> <p>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.</p>
<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 <b>p_size</b> column for partitions <b>2</b> and <b>6</b> 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"
```

Disk Labeling powerformat Command for Solaris

Mode Sense: pcyl 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: pcyl 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

Disk Labeling powerformat Command for Solaris

# INDEX

## A

- all option 33
- audit log message 66

## C

- ceiling copy rate 32
- changing the remote SE server environment variables 37
- CLARiiON and VNX to Symmetrix migration 61
- cleanup migration state 21
- command support summary table 74
- commands, Migration Enabler 32
- committed migration state 17, 20
- committedAndRedirected migration state 17

## D

- disk label
  - installing 83
  - overview 76
  - previewing 79
  - querying 78, 79

## E

- EMC ControlCenter post-migration steps 26
- EMC MirrorView prerequisites 24
- EMC Online Support website 7
- emcgrab utility 66
- Encapsulation technology type
  - licensing 14
- environment variables, changing the remote SE server 37
- error log message 66
- errors or faults 69
- execute operation, powerformat utility 78

## F

- file option 33

## G

- gkselect selection file 25
- graceful host reboot (HostCopy) 67

## H

- HACMP 48
- host crash troubleshooting 67
- HostCopy technology type 14
  - graceful host reboot 67
  - licensing 14
  - migrations 58
  - throttle settings 59
- hostcopy\_ceiling 32
- hostcopy\_ceiling settings 60

## L

- labeling, disk 76
- licensing Migration Enabler 14
- Logical Disk Manager (LDM) 40
- Logical Volume Manager (LVM) 40

## M

- Microsoft Cluster Server 45
- Migration Enabler
  - commands 32
  - licensing 14
  - migration prerequisites 24
  - overview 12
  - post-migration steps 26
  - support summary table 74
- migration speed, tuning 34
- migration states
  - cleanup 21
  - committed 17, 20
  - committedAndRedirected 17
  - overview 15
  - paused 17
  - setup 17, 18
  - sourceSelected 17, 19
  - syncing 17, 19, 27

- targetSelected 17, 20
- migration synchronization speed 34
- migration workflow 16

## N

- native devices, defined 13

## O

- Open Replicator technology type
  - licensing 14
  - migrations 60
  - overview 15
  - prerequisites 24, 25
  - Symmetrix-to-Symmetrix migrations 60
  - throttle settings 63
- OR. *See also* Open Replicator 60

## P

- paused migration state 17
- post-migration steps
  - EMC ControlCenter 26
  - Migration Enabler 26
- powerformat utility
  - execute operation 78
  - overview 38
  - preview operation 78, 79
  - query operation 78, 79
  - updating disk labels 78
  - viewing disk labels 78
- powermig abort command 32
- powermig cleanup command 21, 32
- powermig commands and migration states 16
- powermig commit command 17, 32
- powermig getHandle command 32
- powermig help command 32
- powermig info command 32
- powermig options command 32, 35
- powermig pause command 17, 32
- powermig query command 32
- powermig recover command 32
- powermig resume command 32
- powermig selectSource command 17, 32
- powermig selectTarget command 17, 20, 32
- powermig setup command 17, 18, 33
- powermig sync command 17, 19, 33
- powermig throttle command 33, 34

- powermig undoRedirect command 17, 33
- powermig version command 33
- PPME\_REMOTE\_SE\_PORT environment variable 37
- PPME\_REMOTE\_SE\_SERVER environment variable 37
- prerequisites, Migration Enabler migration 24
- preview operation, powerformat utility 78, 79
- problems, reporting 66
- process hang, resolving 69
- pseudo devices
  - defined 13
  - renaming 54

## Q

- query operation, powerformat utility 78
- querying disk label 78, 79

## R

- reclaim unused space 27
- renaming, pseudo devices 54
- reported faults 68
- reporting problems 66
- resolving a process hang 69

## S

- set pace argument 63
- setup migration state 17, 18
- simultaneous migration sessions, managing 33
- Solaris disk labeling 76
- Solaris host prerequisites 25
- Solutions Enabler (SE) Thin Client 24, 36
- Solutions Enabler symrcopy command 63
- sourceSelected migration state 17, 19
- Sun Cluster with VxVM 53
- Symmetrix Remote Data Facility (SRDF) prerequisites 24
- Symmetrix-to-Symmetrix migration with Open Replicator 60
- synchronization speed, changing 34
- syncing migration state 17, 19, 27

## T

- target device prerequisites 24
- targetSelected migration state 17, 20
- technology type
  - HostCopy 14, 58



- Open Replicator 15, 60
- technology type licensing 14
- thin devices (virtually provisioned devices) 35
- throttle settings
  - HostCopy technology type 59
- TimeFinder/Clone technology type
  - licensing 14
- troubleshooting
  - host crashes 67
  - migrations 65
- tuning migration speed 34

## U

- unused space
  - reclaim 27
- updating disk labels with powerformat 78

## V

- Veritas Cluster Server (VCS) 50
- Veritas Volume Manager (VxVM) 40
- viewing disk labels with powerformat 78
- virtually provisioned devices (thin devices) 35
- VNX/CLARiiON to Symmetrix migration 61

## W

- Windows host prerequisites 24

