



# EMC<sup>®</sup> PowerPath<sup>®</sup> Family

Version 5.7

## CLI and System Messages Reference

P/N 300-014-168

REV 10

Copyright © 2008–2014 EMC Corporation. All rights reserved. Published in the USA.

Published August 2014

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

The information in this publication is provided as is. EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose. Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

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

## Part 1 PowerPath Family CLI Reference

### Chapter 1 PowerPath Multipathing CLI

Multipathing new and changed commands .....	14
powermt: Monitor and manage a PowerPath environment .....	17
Syntax summary.....	20
Running powermt.....	22
Command syntax: dev argument specification .....	23
Command syntax: class argument specification .....	24
powermt check.....	26
powermt check_registration .....	28
powermt config.....	30
powermt disable hba .....	32
powermt display .....	33
powermt display alua.....	45
powermt display hba_mode .....	47
powermt display latency .....	49
powermt display nonvirtual.....	53
powermt display options.....	58
powermt display bus.....	60
powermt display perf dev .....	62
powermt display perf bus.....	65
powermt display port_mode.....	68
powermt_display_dead.....	71
powermt display ports.....	72
powermt display unmanaged .....	75
powermt enable hba .....	76
powermt help.....	77
powermt load.....	78
powermt manage .....	80
powermt release .....	81
powermt remove .....	82
powermt restore.....	84
powermt save .....	86
powermt set autostandby.....	88
powermt set autostandby agingperiod .....	90
powermt set autostandby iopflimit.....	91
powermt set mode .....	92
powermt set path_latency_monitor .....	94
powermt set path_latency_threshold .....	95
powermt set perfmon .....	97
powermt set periodic_autorestore.....	98
powermt set policy .....	99
powermt set port_disable .....	103
powermt set reactive_autorestore .....	105
powermt set streamio_threshold.....	107

powermt set write_throttle .....	108
powermt set write_throttle_queue.....	109
powermt unmanage .....	110
powermt update lun_names.....	111
powermt version .....	112
emcpreg: Manage PowerPath license registration .....	113
emcphostid: Set the Host ID.....	115
emcpadm: List or rename PowerPath pseudo devices.....	118
Syntax.....	119
Commands.....	119
Examples .....	123
emcpcfg script: Save and restore a PowerPath configuration .....	134

**Chapter 2 PowerPath Migration Enabler CLI**

Migration Enabler new and changed commands.....	136
Migration Enabler technology type support .....	138
Migration Enabler syntax summary.....	139
Command syntax .....	140
Argument abbreviations .....	141
powermig abort.....	142
powermig commit .....	145
powermig getHandle .....	146
powermig help .....	147
powermig info.....	148
powermig options .....	150
powermig pause .....	151
powermig query .....	152
powermig recover.....	154
powermig resume .....	155
powermig selectSource .....	156
powermig selectTarget .....	157
powermig setup .....	158
powermig sync.....	162
powermig throttle .....	163
powermig undoRedirect .....	166
powermig version .....	167
powermigcl: Configure Migration Enabler cluster resources .....	168
powermigcl display .....	169
powermigcl config.....	171
powermigcl unconfig.....	172
powermigcl version .....	173
powermigcl help .....	174

**Chapter 3 PowerPath Encryption CLI**

PowerPath Encryption new and changed commands .....	166
ckmadm: configure and manage the key manager client.....	167
ckmadm help .....	168
ckmadm getconfig.....	169
ckmadm getkeyclass.....	170
ckmadm setconfig.....	171
ckmadm setup .....	172
ckmadm setkeyclass .....	173
ckmadm upgrade .....	174

cstadmin: manages the lockbox .....	175
cstadmin initialize.....	176
cstadmin new-passphrase .....	177
cstadmin sign-file .....	178
powervt: manages virtual logical units .....	179
powervt help .....	180
powervt resize.....	181
powervt update .....	182
powervt version.....	183
powervt xcrypt -on.....	184
powervt xcrypt -off.....	185
powervt xcrypt -info.....	186
xcrypt_config script, RKM_Config.bat batch file, and RKM_Config.exe file ..	188
Reconfigure the PowerPath Encryption host.....	189

## Part 2 PowerPath Family System Messages Reference

### Chapter 4 PowerPath Multipathing Messages

Audit and error messages .....	193
Audit log messages.....	193
Error log messages.....	195
Introduction to PowerPath error messages .....	196
Message prefixes .....	196
Messages for different storage systems.....	196
Licensing errors.....	196
Message sources .....	197
PowerPath error messages .....	197
Windows system event IDs.....	225
Windows application event IDs .....	227

### Chapter 5 PowerPath Migration Enabler Messages

Audit messages .....	229
Logging messages to a common file.....	229
Audit and error log messages .....	229
Responding to error messages .....	233

### Chapter 6 PowerPath Encryption Messages

Audit messages .....	241
Audit message format .....	241
Error messages .....	242
Error message sources .....	242
RSA Data Protection Manager (DPM) server messages.....	243
Key manager client messages .....	245
Lockbox service messages .....	245
ckmadm error messages .....	246
cstadmin error messages .....	247
emcp_xcryptd daemon error messages .....	249
powervt error messages .....	250
xcrypt_config error messages .....	253

## Index



# 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. Go to EMC Online Support (<https://support.emc.com>) to ensure that you are using the latest version of this document.

---

## Purpose

This document is part of the PowerPath Family documentation set, and is intended for use by storage administrators and other information system professionals responsible for using, installing, and maintaining PowerPath.

## Audience

Readers of this manual are expected to be familiar with the host operating system where PowerPath runs, storage system management, and the applications used with PowerPath.

## PowerPath documentation set

[Table 1 on page 7](#) lists the PowerPath documentation set.

**Table 1** PowerPath documentation set

Document title	Operating system	Description
PowerPath Family Product Guide	AIX, HP-UX, Linux, Solaris, and Windows	Provides an overview of the PowerPath Family of products.
PowerPath Family for <platform> Release Notes	Platform-specific	Provides the latest information on PowerPath multipathing, PowerPath Encryption with RSA, and PowerPath Migration Enabler including new and changed features, known and fixed problems, and limitations. Where <platform> is: AIX, HP-UX, Linux, Solaris, or Windows
PowerPath for <platform> Installation and Administration Guides	Platform-specific	Describes how to install and configure PowerPath. Where <platform> is: AIX, HP-UX, Linux, Solaris, or Windows
PowerPath Migration Enabler User Guide	AIX, HP-UX, Linux, Solaris, and Windows	Describes how to migrate data with Host Copy, Encapsulation, VPLEX™, Open Replicator, or TimeFinder/Clone® as the underlying technology.
PowerPath Encryption with RSA User Guide	Linux, Solaris, and Windows	Describes how to configure a PowerPath Encryption with RSA environment, how to enable and disable encryption, and how to encrypt data.

These PowerPath manuals are updated periodically. Electronic versions of the updated manuals are available on the EMC Online Support site: <http://support.emc.com>.

From the **Support by Product** pages, search for PowerPath using "Find a Product." These manuals are updated periodically, and the updated manuals are posted on the EMC Online Support site.

If your environment includes Symmetrix storage systems, refer also to the EMC host connectivity guides, which are available on the EMC Online Support site.

This manual does not include information for PowerPath/VE for VMware vSphere. Refer PowerPath/VE for VMware vSphere documentation on the EMC Online Support site.

If your environment includes VNX OE and CLARiiON storage systems, refer also to the following sources:

- ◆ EMC host connectivity guides
- ◆ CLARiiON Storage System Support website
- ◆ VNX OE Storage System Support website

Limited PowerPath functions are supported by the Unisphere™ application for VNX™ Operating Environment (OE) systems.

Limited PowerPath functions are supported by the Navisphere® and the Unisphere™ applications for CLARiiON systems.

If your environment includes other vendors' storage systems, refer to the appropriate documentation from your vendor.

## Revision history

The following table presents the revision history of this document.

**Table 2** Revision history (page 1 of 2)

Revision	Date	Description
10	August 27, 2014	Clarified powermigcl support for Windows Server 2003 and 2008.
09	June 9, 2014	Release of PowerPath 5.7 SP1 for AIX.
08	January 21, 2014	Release of PowerPath for Windows 5.7 SP2. <ul style="list-style-type: none"> <li>• Support for XtremIO firmware Version 2.2 and later</li> <li>• Native Loadable Array Module (LAM) support for NetApp storage arrays FAS2240.</li> </ul>
07	January 15, 2014	Release of PowerPath 5.7 for Linux on System z. <ul style="list-style-type: none"> <li>• Footnote added to <a href="#">Table 3</a> and <a href="#">Table 4</a>.</li> </ul>
06	December 16, 2013	Release of PowerPath 5.7 for AIX.
05	September 2, 2013	Note added to <a href="#">“powermt display unmanaged” on page 75</a> . Corrected the text in <a href="#">“powermt set path_latency_monitor”</a> .
04	June 28, 2013	Introduction of a new command <a href="#">“powermigcl: Configure Migration Enabler cluster resources” on page 168</a> to support the new MSCS cluster functionality beginning with PowerPath 5.7 SP1 for Windows.



**Table 2** Revision history (page 2 of 2)

Revision	Date	Description
03	January 25, 2013	Updated the description for <a href="#">“emcphostid: Set the Host ID” on page 115</a> , <a href="#">“powermt set reactive_autorestore” on page 105</a> . Release of PowerPath 5.7 Windows.
02	December 4, 2012	Release of PowerPath 5.7 SP1 for Linux.
01	September 6, 2012	First release of the EMC PowerPath Family CLI and System Messages Reference Guide Version 5.7.

## Conventions used in this document

EMC uses the following conventions for special notices:



**CAUTION**, used with the safety alert symbol, indicates a hazardous situation which, if not avoided, could result in minor or moderate injury.



**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>

<code>Courier</code>	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 :

[http:// support.emc.com](http://support.emc.com)

**Technical support** — For technical support, go to EMC online support and select **Service Center > Service Center**, 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)

If you have issues, comments, or questions about specific information or procedures, please include the title and, if available, the part number, the revision (for example, 01), the page numbers, and any other details that will help us locate the subject you are addressing.

# PART 1

## PowerPath Family CLI Reference

This section includes the following chapters:

### [Chapter 1, “PowerPath Multipathing CLI”](#)

Describes PowerPath multipathing CLI and includes a summary of new and changed commands since PowerPath 5.0.

### [Chapter 2, “PowerPath Migration Enabler CLI”](#)

Describes PowerPath Migration Enabler CLI and includes a summary of new and changed commands since PowerPath 5.0.

### [Chapter 3, “PowerPath Encryption CLI”](#)

Describes PowerPath Encryption with RSA CLI and includes a summary of new and changed commands since PowerPath 5.2.



# CHAPTER 1

## PowerPath Multipathing CLI

This chapter describes the PowerPath multipathing command line utilities.

- ◆ Multipathing new and changed commands ..... 14
- ◆ powermt: Monitor and manage a PowerPath environment ..... 17
- ◆ emcpreg: Manage PowerPath license registration ..... 113
- ◆ emcphostid: Set the Host ID..... 115
- ◆ emcpadm: List or rename PowerPath pseudo devices..... 118
- ◆ emcpcfg script: Save and restore a PowerPath configuration ..... 134

## Multipathing new and changed commands

Table 3 on page 14 lists the new PowerPath multipathing commands added in versions 5.0 and later, with the earliest supported version listed by platform. Support in later versions (including service packs) can be assumed if no exception in the table notes otherwise.

**Table 3** New and changed commands in versions 5.0 and later (page 1 of 3)

Command	Earliest supported version				
	AIX	HP-UX	Linux <sup>a</sup>	Solaris	Windows
emcpadm	5.3	N/A <sup>b</sup>	All <sup>c</sup>	All	N/A
XML export option is added	5.3	N/A	5.3	5.2 SP1	N/A
emcpreg	All	All	All	All	5.2
Removal of the -noinput option	N/A	N/A	5.3	N/A	5.3
Support for all emcpreg arguments	All	All	All	All	5.3 SP1
powermt check Support for reconfig argument	All	All	N/A	5.1	All
powermt display alua	5.7	5.2	5.7	5.5	5.7
powermt display [dev] Support for the wide   width option	5.0	5.0	5.0	5.0	5.0
Enhanced to include display of autostandby settings, reactive autorestore settings, VPLEX and VNX devices, and performance data collection settings	N/A	5.2	5.7	5.5	5.7
Addition of VNX and CLARiiON arrays failover mode fields	5.1	5.1	5.1	5.1	5.1
powermt disable hba	5.1 <sup>d</sup>	5.0 <sup>e</sup>	5.1 <sup>e</sup>	5.0 <sup>d</sup>	N/A
powermt display hba_mode	5.1	5.0	5.1	5.0	N/A
powermt display latency	5.1	5.0	5.0	5.0	5.3
powermt display nonvirtual	5.0	5.0	5.3	5.0	5.3
powermt display options Information on path latency monitoring displays	5.1	5.0	5.0	5.0	5.3
Information on autostandby and reactive autorestore	5.7	5.2	5.7	5.5	5.7
Removal of Default Storage System from the output and the PP_DEFAULT_STORAGE_SYS environment variable	5.7	5.2	5.7	5.5	5.7
powermt display paths	5.0	5.0	5.0	5.0	5.0
Change display paths to display bus in usage	5.7	5.2	5.7	5.5	5.7
powermt display bus Displays path performance metrics for buses	5.7	5.2	5.7	5.5	5.7
powermt display perf dev	5.7	5.2	5.7	5.5	5.7
powermt display perf bus	5.7	5.2	5.7	5.5	5.7

**Table 3** New and changed commands in versions 5.0 and later (page 2 of 3)

Command	Earliest supported version				
	AIX	HP-UX	Linux <sup>a</sup>	Solaris	Windows
powermt display port_mode	5.1	5.0	5.1	5.0	N/A
powermt enable hba	5.1	5.0	5.1	5.0	N/A
powermt help	5.7	5.2	5.7	5.5	5.7
powermt load At boot time, loads configuration data from an XML formatted file	5.7	5.2	5.7	5.5	5.7
powermt manage Removal of the all option	5.0	5.0	5.0	5.0	N/A
powermt save Saves configuration data into an XML formatted file	5.7	5.2	5.7	5.5	5.7
powermt set autostandby	5.7	5.2	5.7	5.5	5.7
powermt set autostandby agingperiod	5.7	5.2	5.7	5.5	5.7
powermt set autostandby iopflimit	5.7	5.2	5.7	5.5	5.7
powermt set mode Addition of the force option to override an autostandby setting	5.7	5.2	5.7	5.5	5.7
powermt set path_latency_monitor	5.1	5.0	5.0	5.0	5.3
powermt set path_latency_threshold	5.1	5.0	5.0	5.0	5.3
When powermt set path_latency_monitor=off is run, the latency values for the paths are reset to 0.	5.3	5.2	5.3	5.3	5.3
The range of valid latency threshold values has been changed to 0 to 3600.	5.5	5.1 SP2	5.3 SP1	5.3	5.5
powermt set perfmom	5.7	5.2	5.7	5.5	5.7
powermt set policy Addition of StreamIO (si) policy	5.3	5.1 SP1	5.1 SP2	5.2 SP1	5.3
Removal of BasicFailover (bf) and NoRedirect (nr) from powermt usage	5.5	5.2	5.5	5.5	5.5
powermt set port_disable port_disable={true false}	5.1	5.0	5.1	5.0	N/A
port_disable={on off}	5.7	5.2	5.7	5.5	N/A
powermt set priority General support is removed	5.5	5.2	5.5	5.5	5.5
powermt set reactive_autorestore	5.7	5.2	5.7	5.5	N/A
powermt set streamio_threshold	5.5	5.1 SP2	5.5	5.3	5.5
powermt set transparent General support is removed. Use powermt set policy instead.	4.x	4.x	4.x	4.x	N/A

**Table 3** New and changed commands in versions 5.0 and later (page 3 of 3)

Command	Earliest supported version				
	AIX	HP-UX	Linux <sup>a</sup>	Solaris	Windows
powermt unmanage Removal of the all option	5.0	5.0	5.0	5.0	N/A
powermt version The default output displays the full version, including hot fixes and the build number. Options exist to retain legacy-style output.	5.1 SP1	5.1 SP1	5.1 SP1	5.1 SP1	5.1 SP2
Removal of the format option	5.7	5.2	5.7	5.5	5.7

- a. 5.x related features and powermt CLI specific to Symmetrix are also applicable to PowerPath 5.7 for Linux on System z.
- b. *N/A* means the command or utility is not supported on the platform.
- c. *All* means all current PowerPath versions support the feature, including versions earlier than 5.0.
- d. Fibre Channel only.
- e. Fibre Channel and iSCSI only.



## powermt: Monitor and manage a PowerPath environment

Use the **powermt** commands listed in [Table 4 on page 17](#) to manage a PowerPath environment.

On Windows, PowerPath Administrator is the easiest way to access **powermt** functions. PowerPath Administrator is described in the *EMC PowerPath and PowerPath/VE for Windows Installation and Administration Guide* and the PowerPath Administrator online help.

**Table 4** powermt commands (page 1 of 3)

Command	Function	Storage systems	Platform <sup>a</sup>	Refer to page
powermt check	Checks each specified path and prompts you to remove from the PowerPath configuration any paths marked dead.	All	All	<a href="#">page 26</a>
powermt check_registration	Checks the state of the PowerPath license on a host.	All	All	<a href="#">page 28</a>
powermt config	Discovers and configures all working paths that are already configured by the operating system. On Windows, PowerPath automatically does this when needed.	All	UNIX, Linux	<a href="#">page 30</a>
powermt disable hba	Disables all paths configured through the specified HBA; I/O to the HBA and its devices is stopped.	All	AIX, Solaris (FC only) HP-UX, Linux (FC & iSCSI <sup>b</sup> )	<a href="#">page 32</a>
powermt display [dev]	Displays the state of HBAs configured for PowerPath. With <b>dev</b> , displays attached device specific information.	All	All	<a href="#">page 33</a>
powermt display alua	Displays the ALUA State for all the paths to an ALUA device.	ALUA devices only: EMC CLARiiON/VNX HP EVA 8000/8100 HP EVA 3000/5000 (with firmware VCS 4.x) IBM DS 8000/8100	All	<a href="#">page 45</a>
powermt display hba_mode	Displays information about all HBAs visible to the host, including the HBA mode (enabled or disabled).	All	UNIX, Linux	<a href="#">page 47</a>
powermt display latency	Displays storage-system path information, including metrics on path latency monitoring.	All	All	<a href="#">page 49</a>

**Table 4** powermt commands (page 2 of 3)

Command	Function	Storage systems	Platform <sup>a</sup>	Refer to page
powermt display nonvirtual	Displays information about the physical devices associated with InVista Virtual Volumes.	InVista	All	<a href="#">page 53</a>
powermt display options	Displays option settings for storage-system classes.	All	All	<a href="#">page 58</a>
powermt display bus	Displays the state of the buses for PowerPath.	All	All	<a href="#">page 60</a>
powermt display perf dev	Displays storage device performance metrics.	All	All	<a href="#">page 62</a>
powermt display perf bus	Displays path performance metrics for all the buses.	All	All	<a href="#">page 65</a>
powermt display port_mode	Displays storage-system port information, including the port mode (enabled or disabled).	All	UNIX, Linux	<a href="#">page 68</a>
powermt display ports	Displays attached array port specific information.	All	All	<a href="#">page 72</a>
powermt display unmanaged	Displays logical devices that are excluded from PowerPath management.	All	UNIX, Linux	<a href="#">page 75</a>
powermt enable hba	Enables an HBA port that was previously disabled, allowing I/O to the specified HBA and its devices to resume.	All	AIX, Solaris (FC only) HP-UX, Linux (FC & iSCSI <sup>a</sup> )	<a href="#">page 76</a>
powermt help	Displays the usage of all powermt commands.	All	All	<a href="#">page 77</a>
powermt load	Loads a previously saved PowerPath configuration.	All	All	<a href="#">page 78</a>
powermt manage	Places a specified logical device, or specified storage system class, under PowerPath management.	All	UNIX, Linux	<a href="#">page 80</a>
powermt release	Releases pseudo device names for later reuse.	All	Linux	<a href="#">page 81</a>
powermt remove	Removes a path from the PowerPath configuration.	All	All	<a href="#">page 82</a>
powermt restore	Tests and restores specified paths. Paths that pass the test remain in service or are returned to service. Paths that fail the test remain out of service or are taken out of service.	All	All	<a href="#">page 84</a>
powermt save	Saves the current settings as a custom PowerPath configuration.	All	All	<a href="#">page 86</a>
powermt set autostandby	Enables or disables the autostandby feature. Autostandby is enabled by default after installation or upgrade.	All	All	<a href="#">page 88</a>

**Table 4** powermt commands (page 3 of 3)

Command	Function	Storage systems	Platform <sup>a</sup>	Refer to page
powermt set autostandby agingperiod	Sets the length of time a path will stay in autostandby due to an <b>iopf</b> trigger before it is returned to active.	All	All	<a href="#">page 90</a>
powermt set autostandby iopflimit	Sets the average number of failed I/Os needed for a path error to automatically set that path to autostandby mode.	All	All	<a href="#">page 91</a>
powermt set mode	Sets paths to either active or standby mode.	All	All (but not with iSCSI)	<a href="#">page 92</a>
powermt set path_latency_monitor	Enables or disables path latency monitoring on all paths.	All	All	<a href="#">page 94</a>
powermt set path_latency_threshold	Sets on all paths a time limit for the completion of I/O.	All	All	<a href="#">page 95</a>
powermt set perfmon	Enables or disables performance monitoring for all devices.	All	All	<a href="#">page 97</a>
powermt set periodic_autorestore	Enables or disables PowerPath's periodic autorestore facility.	All	HP-UX, Linux, Solaris, AIX	<a href="#">page 98</a>
powermt set policy	Changes the load-balancing and failover policy for devices configured in PowerPath.	All	All	<a href="#">page 99</a>
powermt set port_disable	Disables the array port for the specified device. This command stops all I/O to the port and closes all devices configured through the port.	All	UNIX, Linux	<a href="#">page 103</a>
powermt set reactive_autorestore	Enables or disables PowerPath's reactive autorestore facility.	All	HP-UX, Linux, Solaris, AIX	<a href="#">page 105</a>
powermt set streamio_threshold	Sets the I/O threshold to switch paths while the Stream I/O (si) policy is in effect.	All	All	<a href="#">page 107</a>
powermt set write_throttle	Enables or disables I/O write throttling to a logical device.	Symmetrix	UNIX	<a href="#">page 108</a>
powermt set write_throttle_queue	Sets the write throttling queue depth for a storage system.	Symmetrix	UNIX	<a href="#">page 109</a>
powermt unmanage	Excludes a specified logical device or storage system class from PowerPath management.	All	UNIX, Linux	<a href="#">page 110</a>
powermt update lun_names	Causes PowerPath to retrieve the latest user-assignable LUN names for devices on a VNX and CLARiiON system.	VNX and CLARiiON	All	<a href="#">page 111</a>
powermt version	Displays the version of PowerPath that is installed on the host.	All	All	<a href="#">page 112</a>

a. 5.x related features and powermt CLI specific to Symmetrix are also applicable to PowerPath 5.7 for Linux on System z.

- b. On Linux, you can use **powermt disable**, **powermt enable**, and **powermt set port\_disable** in an iSCSI environment with HBAs; these commands are unsupported with NICs.

## Syntax summary

A summary of the syntax of all **powermt** commands follows:

```
powermt check [hba=<hba#>|all] [dev=<path>|<device>|all]
    [class=<class>|all] [force]

powermt check_registration

powermt config

powermt disable hba=<hba#>

powermt display [dev=<device>|all] [class=<class>|all]
    [every=<#seconds>] [wide|width=<#col>]

powermt display alua dev=<device>|all [class=<class>|all]
    [every=<#seconds>] [wide|width=<#col>]

powermt display hba_mode [class=<class>|all]

powermt display latency [dev=<device>|all] [class=<class>|all]
    [every=<#seconds>] [wide|width=<#col>]

powermt display nonvirtual dev=<device>|all
    [every=<#seconds>] [class=<class>|all] [wide|width=<#col>]

powermt display options

powermt display bus [class=<class>|all] [every=<#seconds>]
    [wide|width=<#col>]

powermt display perf dev=<device>|all [continuous] [verbose] [xml]
    [nowait]

powermt display perf bus [continuous] [verbose] [xml] [nowait]

powermt display port_mode [dev=<device>|all] [class=<class>|all]

powermt display ports [dev=<device>|all] [class=<class>|all]
    [every=<#seconds>] [wide|width=<#col>]

powermt display unmanaged

powermt enable hba=<hba#>

powermt help [<command>]

powermt load [file=<filename>]

powermt manage dev=<path>|<device> | class=<class>

powermt release

powermt remove hba=<hba#>|all | dev=<path>|<device>|all
    [class=<class>|all] [force]

powermt restore [hba=<hba#>|all] dev=<path>|<device>|all
    [class=<class>|all]

powermt save [file=<filename>]
```

```

powermt set autostandby={on|off|reinitialize}
           [trigger={prox|iopf}]

powermt set autostandby iopflimit=<value>

powermt set autostandby agingperiod=<#days>

powermt set mode={active|standby} [hba=<hba#>|all]
           [dev=<path>|<device>|all] [class=<class>|all] [force]

powermt set path_latency_monitor={on|off}

powermt set path_latency_threshold=<#seconds>|<#milliseconds>ms

powermt set perfmon={on [interval=<#seconds>] | off}

powermt set periodic_autorestore={on|off} [class=<class>|all]

powermt set policy={ad|co|lb|li|re|rr|so|si}
           [dev=<device>|all] [class=<class>|all]

powermt set port_disable={on|off} dev=<device>

powermt set reactive_autorestore={on|off} [class=<class>|all]

powermt set streamio_threshold=<threshold_count>
           [dev=<device>|all] [class=<class>|all]

powermt unmanage dev=<path>|<device> | class=<class>

powermt update lun_names

powermt version

```

## powermt install directory

In a default PowerPath installation, the **powermt** utility resides in the directory shown in [Table 5 on page 21](#).

**Table 5** powermt install directory

Platform	Directory
AIX	/usr/sbin
HP-UX	/sbin
Linux	/sbin
Solaris	/etc
Windows	%systemdrive%\Program Files\EMC\PowerPath

## Required access to run powermt

[Table 6 on page 22](#) shows the required access to run **powermt**.

**Table 6** Access to run powermt

Platform	Access
Windows	User Account Control, offered by Microsoft, helps regulate the quantum of administrator access.
UNIX, Linux	By default, the PowerPath installation process sets file permissions so only the superuser (root) can run <b>powermt</b> . You can allow other users to run <b>powermt</b> by changing access permissions on the <b>powermt</b> command, by setting the setuid/setgid bits.

## Running powermt

[Table 7 on page 22](#) shows the instructions to run **powermt**.

**Table 7** Running powermt

Platform	How to run powermt
Windows	Launch the Command Prompt window, and then run the command in that window.
UNIX, Linux	At the shell prompt, type any of the commands listed in this chapter.

## Command status

All **powermt** commands return 0 to the user process on success; 1 on error.

## Command syntax: dev argument specification

Several **powermt** commands have a **dev** argument. Depending on the command, this argument can specify a logical device or a path:

dev acts on logical devices in:	dev acts on paths in:
powermt display powermt display alua powermt display latency powermt display perf dev powermt manage powermt set path_latency_monitor powermt set perfmon powermt set policy powermt set port_disable powermt set write_throttle powermt unmanage	powermt check powermt display perf bus powermt remove powermt restore powermt set mode powermt set perfmon powermt set port_disable (acts on ports associated with paths) powermt set write_throttle_queue (acts on ports associated with paths)

You can specify either a PowerPath device or a native/path device in the **dev** argument.

The **powermt set perfmon** is a global command that turns performance monitoring on or off for all devices. It cannot be set for individual devices.

For Linux versions that have a dynamic /dev directory (such as, Linux SLES 11 and later), use pseudo device names with all **powermt dev** arguments, except when running these commands:

- ◆ powermt manage
- ◆ powermt set port\_disable

The **dev** argument is context specific. **powermt** interprets the meaning of the argument based on the context of the command being executed:

- ◆ For commands that act on logical devices, you can specify the logical device using either the PowerPath device or any of the native/path devices associated with that logical device.
- ◆ For commands that act on paths, specifying a PowerPath device causes the command to be applied to all paths associated with the device; specifying an associated native/path device causes the command to be applied only to that path.

Table 8 on page 24 shows the **dev** argument syntax for each host platform.

**Table 8** The dev argument syntax for powermt commands

Platform	To specify a device	To specify a path	Examples
AIX	hdiskpower# or #	hdisk#	hdiskpower3, 3, hdisk3
HP-UX legacy devices	N/A	c##d##	c1t2d2
HP-UX agile devices	disk#	c##d## or lunpath##	disk1010, c1t2d2
Linux	emcpower[a-z]	sdn (block device)	emcpowerc, sda
Solaris SPARC	emcpower#[a-h] or #	c##d#[s0-s7]	emcpower2a, 2, c1t1d1
Solaris Opteron	emcpower#[a-p] or #	c##d#[s0-s15] c##d#[p0-p4]	emcpower5m, c2t1d11p2
Windows	harddisk# or #	c##d##	harddisk3, 3, c1t2d3

## Command syntax: class argument specification

Several **powermt** commands have a **class** argument that specifies a storage system type. Table 9 on page 24 lists the valid values for the **class** argument added in versions 5.0 and later, with the earliest supported version listed by platform. Support in later versions (including service packs) can be assumed if no exception in the table notes otherwise. For the list of storage systems that EMC currently supports, see the *E-Lab Interoperability Navigator*.

**Table 9** Valid values for class argument (page 1 of 2)

Class	Storage arrays	AIX	HP-UX	Linux	Solaris	Windows
all	All storage arrays	4.x <sup>a</sup>	4.x	4.x	4.x	4.x
clariion	EMC CLARiiON arrays	4.x <sup>b</sup>	4.x <sup>b</sup>	4.x <sup>b</sup>	4.x <sup>b</sup>	4.x <sup>b</sup>
		5.7 <sup>c</sup>	5.2 <sup>c</sup>	5.7 <sup>c</sup>	5.5 <sup>c</sup>	5.7 <sup>c</sup>
ess	IBM arrays	4.x	4.x	4.x	4.x	4.x
generic	EMC Celerra and EMC VNXe arrays	5.7	N/A <sup>d</sup>	5.5	5.5	5.5
hitachi	HDS arrays	4.x	4.x	4.x	4.x	4.x
hpxp	HP arrays	4.x	4.x	4.x	4.x	4.x
invista	EMC Invista arrays	4.x <sup>e</sup>	4.x <sup>e</sup>	4.x <sup>e</sup>	4.x <sup>e</sup>	4.x <sup>e</sup>
		5.7 <sup>f</sup>	5.2 <sup>f</sup>	5.7 <sup>f</sup>	5.5 <sup>f</sup>	5.7 <sup>f</sup>
netapp	NetApp arrays	5.7	N/A	5.7 SP3	N/A	5.7 SP2
symm	EMC Symmetrix arrays	4.x	4.x	4.x	4.x	4.x



**Table 9** Valid values for class argument (page 2 of 2)

Class	Storage arrays	AIX	HP-UX	Linux	Solaris	Windows
vnx	EMC VNX arrays	5.7 <sup>g</sup>	5.2 <sup>g</sup>	5.7 <sup>g</sup>	5.5 <sup>g</sup>	5.7 <sup>g</sup>
vplex	EMC VPLEX arrays	5.7 <sup>h</sup>	5.2 <sup>h</sup>	5.7 <sup>h</sup>	5.5 <sup>h</sup>	5.7 <sup>h</sup>
xtremio	EMC XtremIO arrays	5.7 SP1	N/A	5.7 SP3	N/A	5.7 SP2

- a. 4.x means that this class was a valid value in PowerPath prior to version 5.0.
- b. Both VNX and CLARiiON arrays are supported under the **clarion** array class for the indicated PowerPath version.
- c. CLARiiON arrays are supported under the **clarion** array class for the indicated PowerPath version and platform, while VNX arrays are supported under the **vnx** array class.
- d. N/A (not applicable) means that class is not supported for that platform.
- e. Both Invista and VPLEX arrays are supported under the **invista** array class for the indicated PowerPath version.
- f. Invista arrays are supported under the **invista** array class for the indicated PowerPath version and platform, while VPLEX arrays are supported under the **vplex** array class.
- g. VNX arrays are supported under the **vnx** array class.
- h. VPLEX arrays are supported under the **vplex** array class.

## powermt check

**Purpose** Check a PowerPath configuration.

**Syntax**

```
powermt check
  [hba=<hba#>|all] [dev=<path>|<device>|all]
  [class=<class>|all] | [force] [reconfig]
```

**Description** **powermt check** checks specified paths and, if desired, removes from the PowerPath configuration any paths marked dead.

If a path is marked dead or the serial numbers encoded in the path configuration information do not match the serial numbers on the logical device, **powermt check** prompts you to remove the path:

```
Warning: <storage_system> I/O path <path_name> is dead.
Do you want to remove it (y/n/a/q)?
```

Valid responses are:

- y** Remove the dead path, and continue checking remaining paths.
- n** Do not remove the path, but continue checking remaining paths.
- a** Remove the dead path and any subsequent paths marked dead.
- q** Do not remove the dead path, and exit the command. Any paths that were already removed remain removed.

If the PowerPath license is invalid, **powermt check** issues a warning and continues its processing. Use **powermt check\_registration** to determine the problem with the PowerPath license.

### Reconfiguring mismatched paths (Windows and Solaris only)

Instead of removing dead paths, **powermt check reconfig** resets *mismatched* paths. Mismatched paths are paths that are declared dead because they are connected to unexpected LUNs, either:

- ◆ The connection is good but the LUN ID has an unexpected value; that is, the path points at the wrong LUN.
- ◆ The connection is good but the driver cannot interpret INQ data; that is, the path points at an unexpected array type.

If a new device is detected at the previously owned SCSI address of a dead path, the state of the new device is shown as mismatched in the GUI (on Windows) and dead in the **powermt** output (on Solaris and Windows).

Paths become mismatched occasionally on Fibre Channel networks as a result of SAN or array reconfiguration. Mismatching occurs more frequently on iSCSI networks, where iSCSI login order can cause the condition.

The **powermt display** output does not distinguish paths that are marked dead because of mismatch and paths that are marked dead because of a bad connection. On Windows hosts, the GUI does.

**CAUTION**

For operating systems on which PowerPath exposes native paths to the application level, `powermt check` or `powermt check reconfig` can cause data loss or data unavailability if you clean up dead paths configured for use by an application. Before confirming the operation, ensure any presented misconfigured or dead device path is not configured by an application for access to the old volume.

**Arguments**

`hba=<hba#>|all`

Limits the check to paths from the specified HBA. *hba#* is a number in the **Host Bus Adapters ###** column of `powermt display` output. **all** specifies all HBAs under PowerPath control. The default is **all**.

(**adapter** can be used instead of **hba**. **adapter** is included only for compatibility with prior PowerPath releases. Use **hba** instead.)

`dev=<path>|<device>|all`

Limits the check to the specified *path*, or all paths to the specified *device*. **all** specifies all paths to all devices. The default is **all**.

`class={symm|vnx|clariion|vplex|invista|netapp|ess|xtremio|hitachi|hpxp|generic|all}`

Limits the command to the specified type of storage system. **all** specifies all storage-system types. The default is **all**. For descriptions of all the class types, see [“Command syntax: class argument specification” on page 24](#).

If **class** is specified along with **hba** or **dev**, the command is applied to paths that meet all specified constraints.

`force`

Suppresses the confirmation prompt and silently removes all paths marked dead (or, if **reconfig** is specified, silently reconfigures all mismatched paths).

---

**Note:** EMC does not recommend using the [force] option in production environments.

---

`reconfig`

---

**Note:** The reconfig argument is supported only on Solaris and Windows platforms.

---

Reconfigures mismatched or misconfigured paths. Path mismatching occurs more frequently on iSCSI networks, where path login order can cause the condition. This option does not remove paths that are truly dead (use **powermt check** for this). Note that if you remove an iSCSI LUN from a VNX and CLARiiON storage group (but not from the fabric), `powermt check reconfig` sees the paths to the removed LUN as misconfigured and prompts you to reconfigure (that is, remove) them, as follows:

Warning: CLARiiON device path c3t1d0s0 is currently misconfigured.

Do you want to reconfigure it (y/n/a/q)?

If you choose to reconfigure them by selecting **y**, the paths are removed. (Under normal circumstances, the reconfig option does not remove dead paths.) The paths are seen as misconfigured rather than dead because the iSCSI LUN responds to inquiry data.

If you choose **n**, the paths to the removed volumes remain. You can use `check reconfig` to expose them.

### Environment variable

Running the **powermt check force** command saves the current PowerPath configuration. Use the **PP\_DISABLE\_AUTOSAVE** environment variable to disable any autosave of the configuration.

**PP\_DISABLE\_AUTOSAVE**=<value>

### Platform-specific notes

On Windows

On PowerPath for Windows 5.7 SP2 and later, the **PP\_DISABLE\_AUTOSAVE** environment variable is set to 1 by default.

To re-enable autosave, set the environment variable to **0** or **False**.

## powermt check\_registration

**Purpose** Check license information.

**Syntax** `powermt check_registration`

### Description

**powermt check\_registration** checks the state of the PowerPath license on a host. **powermt check\_registration** determines whether the host has:

- ◆ A valid permanent license or evaluation license
- ◆ An expired evaluation license

Evaluation licenses provide the same functionality as permanent licenses, but an evaluation license expires on the last day of a specified month.

- ◆ An invalid or corrupt license
- ◆ No license

---

**Note:** EMC recommends applying one license per server. When different devices use different licenses, the **powermt set** command with the **dev=all** option may apply the most restrictive license to all devices.

---

### Detailed description

**powermt check\_registration** displays output like the following:

```
Key <number>
Product: <PowerPath product>
Capabilities: <license> - <storage_systems>
```

where *<number>* is your 24-digit registration key, *<license>* is your license type, and *<storage\_systems>* is a list of the storage systems on which this license is valid (for example, VNX and CLARiION).

For example, output similar to the following appears on PowerPath/VE for Windows:

```
Key xxxx-xxxx-xxxx-xxxx-xxxx-xxxx
Product: PowerPath/VE for Windows
Capabilities: All
```

For an evaluation license, **powermt check\_registration** also outputs a line like this:

Expires: <month>, <year>

For an expired license, **powermt check\_registration** also outputs a line like this:

Expired: <month>, <year>

You cannot run certain **powermt** commands if the license is a Base license or is expired, or if there is no license (except on a host connected exclusively to CLARiiON AX-series arrays).

With an unlicensed version of PowerPath or with a PowerPath Base license:

- ◆ You cannot run these **powermt set** commands:
  - powermt set policy (with third-party arrays)
  - powermt set mode
- ◆ You can run other **powermt set** commands.

The *EMC PowerPath Family Product Guide* provides a summary of the platform, array, and feature support available with each PowerPath license type. [Table 21 on page 100](#) provides information about load-balancing and failover policies.

### CLARiiON AX-series arrays

Certain versions of PowerPath (as listed in the E-Lab Interoperability Navigator™ on EMC Online Support site) provide full functionality with or without a license when the host is connected exclusively to CLARiiON AX-series storage systems. Note that models earlier than the CLARiiON AX4-5 (for example, AX100 and AX150) are not supported with PowerPath for AIX or HP-UX.

---

**Note:** If a license for a CLARiiON AX-series storage system is present on the host as a result of an upgrade from a previous version of PowerPath, the **powermt check\_registration** command displays **CLARiiON Restricted** in the Capabilities field for this license.

---

## On Windows

The PowerPath Administrator GUI reports license status as either **Full** or **Restricted** in the Array Management Information box on the Options tab of the EMC PowerPathAdmin Properties (root node) dialog. **Full** indicates the host has full PowerPath capabilities for all supported storage system types. **Restricted** indicates PowerPath capabilities are limited for some supported storage system types. CLARiiON AX-series storage systems have full PowerPath capabilities regardless of the displayed license state.

## powermt config

**Purpose** Configure paths to logical devices.

**Syntax** `powermt config`

**Description** **powermt config** performs the following:

- ◆ configures all detected Symmetrix, VNX, and CLARiiON logical devices as PowerPath devices and adds these devices to the PowerPath configuration
- ◆ configures all detected third-party storage system logical devices as PowerPath devices if their storage system classes are set to managed
- ◆ creates devices as required
- ◆ creates devices on these platforms: AIX, Linux, Solaris and uses existing, native devices on these platforms: HP-UX, Linux, Solaris
- ◆ configures all detected paths to PowerPath devices and adds these paths to the PowerPath configuration
- ◆ adds paths to logical devices based on the storage-system frame serial number and the logical device serial number. Together, these values (shown in the output of **powermt display dev**) uniquely identify a logical device. By default, **powermt config** adds devices under PowerPath's control with the Symmetrix optimization, CLARiiON optimization, or Adaptive load-balancing and failover policy, write throttling set to off, and a write throttle queue depth of 256. It adds paths with the mode set to active. It adds storage systems with periodic autorestore set to on

**powermt config** does not remove previously configured paths when they become dead paths.

After **powermt config** completes, you can change the configuration with several **powermt** commands:

- ◆ `powermt load`
- ◆ `powermt set autostandby`
- ◆ `powermt set autostandby agingperiod`
- ◆ `powermt set autostandby iopflimit`
- ◆ `powermt set mode`
- ◆ `powermt set path_latency_monitor`
- ◆ `powermt set path_latency_threshold`
- ◆ `powermt set perfmon`
- ◆ `powermt set periodic_autorestore`
- ◆ `powermt set policy`
- ◆ `powermt set reactive_autorestore`
- ◆ `powermt set write_throttle`
- ◆ `powermt set write_throttle_queue`

If the PowerPath license is invalid, **powermt config** issues a warning and continues its processing. Use **powermt check\_registration** to determine the problem with the PowerPath license.

**Platform-specific notes**

On AIX

- ◆ If **powermt config** encounters an error, a general AIX error message is displayed, indicating there is a problem configuring a device. The text of the message depends on the underlying error. For more detailed PowerPath error information, it is recommended that AIX users run **cfgmgr -vl powerpath0** instead of **powermt config**. The **cfgmgr** command performs the same function as **powermt config** but generates more detailed error messages.
- ◆ **powermt config** hangs if either the / or the /tmp file system is full.
- ◆ PowerPath is configured at boot time using ODM (Object Data Manager) configuration rules. Run **powermt config** only if more paths are added after boot time.

On HP-UX and Linux, **powermt config** runs at boot time from /etc/init tab to configure PowerPath.

Also on Linux, before running **powermt config**, make sure the required special native SCSI disk block device files exist in the /dev directory.

On Solaris, **powermt config** runs at boot time from startup scripts under /etc/rcS.d.

On Windows, PowerPath automatically runs **powermt config** when needed.

## powermt disable hba

**Purpose** Disable an HBA port.

**Syntax** `powermt disable hba=<hba#>`

**Description** **powermt disable** disables all paths configured through the specified HBA. All I/O on the port is stopped, and devices configured through the port are closed.

This command is supported with the following platform configurations:

- ◆ AIX and Solaris platforms—Fibre Channel only.
- ◆ HP-UX platforms—iSCSI and Fibre Channel.
- ◆ Linux platforms—iSCSI (with HBAs only) and Fibre Channel.

Run **powermt disable** when, for example, you want to stop I/O to an HBA that needs to be replaced. When you disable an HBA port, all paths configured through the HBA are labeled **Dead** in the **powermt display dev=<device>** output. In the **powermt display hba\_mode** output, **Disabled** appears in the **Mode** column.

Note the following limitations:

- ◆ A disabled HBA is not persistent across host reboots.
- ◆ This command fails if the PowerPath load balancing and failover policy is **bf** or **nr**.

To check whether an HBA is enabled or disabled, use the **powermt display hba\_mode** command.

**Arguments** `hba=<hba#>`  
Disables the specified HBA port using the HBA ID on the host, for example, 1280.



## powermt display

**Purpose** Monitor HBAs and devices.

**Syntax**

```
powermt display
  [dev=<device>|all] [class=<class>|all]
  [every=<#seconds>] [wide|width=<#columns>]
```

**Description** **powermt display** displays information about HBAs or devices configured for and managed by PowerPath.

PowerPath configures a native path device the first time it is opened. If, before you run **powermt config**, you run some system utility that opens devices, **powermt display** shows the native path devices. (The associated pseudo devices are not displayed until you run **powermt config**).

The **powermt display dev** command includes an **Array failover mode** field, which displays the failover mode setting for VNX and CLARiiON arrays. This setting determines the failover mode for a server's HBA ports, and is configured through the VNX and CLARiiON Navisphere software. Failover mode values are described in [Table 11 on page 40](#).

---

**Note:** After changing devices to ALUA mode, rebooting the host is required for PowerPath to recognize the change.

---

### Arguments

**dev=<device>|all**

Displays information about the specified *device*. **all** specifies all devices. If **dev** is not specified, a summary HBA listing is displayed, with one line per HBA.

**class={symm|vnx|clariion|vplex|invista|netapp|ess|xtremio|hitachi|hpxp|generic|all}**

Limits the command to the specified type of storage system. **all** specifies all storage-system types. The default is **all**. For descriptions of all the class types, see [“Command syntax: class argument specification” on page 24](#).

If **class** is specified along with **dev**, the command is applied to HBAs or devices that meet all specified constraints.

**every=<#seconds>**

Integer in the range 1 to 86400 that specifies, in seconds, how often the display is updated. By default, the display is not updated.

The *seconds* value is the minimum time between refreshes; the actual time is affected by the overall system load. On busy systems, display updates can be less frequent than specified.

You can add a time stamp to the **powermt display every=<#seconds>** command output by defining the **PP\_DISPLAY\_TIME\_STAMP** environment variable and setting the time display format to **TIME\_VERBOSE** or **TIME\_SECONDS**. **TIME\_VERBOSE** displays complete details such as day, month, and year, time in hours, minutes, and seconds that increments until you exit the **powermt display** command. **TIME\_SECONDS** starts a counter from zero that increments until you exit the **powermt display** command.

If the environment variable is not defined or if it is defined but not assigned a proper value the time stamp is not displayed.

The following example shows **TIME\_VERBOSE** value output:

```
powermt display dev=1 every=2
Time: Wed Jul 28 08:49:18 2009
Pseudo name=emcpower1a
[...]
```

```
Time: Wed Jul 28 08:49:20 2009
Pseudo name=emcpower1a
[...]
```

The following example shows **TIME\_SECONDS** value output:

```
powermt display dev=1 every=2
Time: +0sec
Pseudo name=emcpower1a
[...]
```

```
Time: +2sec
Pseudo name=emcpower1a
[...]
```

`wide|width=<#columns>`

**Note:** The `wide` argument is not supported on Windows platforms.

**wide** sets the line length to the width of the window and **width** sets a range within which the **powermt display** output appears. The default is 80. The acceptable range is 80 to 160 columns. The output widens or narrows within the range depending on the width of the window. This value applies only to the specific output; it is not persistent. You can set a persistent value for the **width=<#columns>** option by using the environment variable **PP\_DISPLAY\_WIDTH**. This variable changes the width setting from 80 (the default) to the value you specify. If the environment variable value differs from the value you set with the command line, the command line value takes precedence.

If you set the **every=<#seconds>** option along with the **width=<#columns>** option, the **powermt display** output widens or narrows to accommodate changes you make to the window width.

## Detailed description

The first form of the command (**powermt display**) yields the “[HBA display](#)” on page 34. The second form of the command (**powermt display dev**) yields the “[Devices display](#)” on page 39.

## Format of examples in this section

All **powermt display** examples are for one of the platforms on which PowerPath runs. Output on all storage systems and platforms contains the same column headers, but the format of information in some fields differs according to platform.

## HBA display

The HBA display is presented if the **dev** option is not used. A sample of the basic HBA display (when neither the **paths** nor the **ports** argument is used) follows:

```
powermt display
Symmetrix logical device count=20
CLARiiON logical device count=21
Ess logical device count=10
Hitachi logical device count=9
HP xp logical device count=9
Invista logical device count=10
```

```

=====
----- Host Bus Adapters ----- I/O Paths ----- Stats -----
### HW Path                      Summary   Total   Dead   IO/Sec Q-IOs Errors
=====
1792 pci@16/fibre-channel@2       optimal    20     0     -     0     0
1793 pci@16/fibre-channel@4       optimal    20     0     -     0     0
2304 pci@17/fibre-channel@2       optimal   128     0     -     0     0
2305 pci@16/fibre-channel@1       optimal   128     0     -     0     0
=====

```

[Table 10 on page 37](#) provides more information about all fields in the HBA display. The values in the **HW Path** column are platform-specific.

For an HBA display with the paths argument, see [“powermt display perf bus” on page 65](#). For an HBA display with the ports or device argument, see [“powermt display ports” on page 72](#).

## Fabric zoning and displays of path information

Fabric zoning influences the way in which HBAs and the total number of paths are displayed. With the **paths** argument, the information under the **Host Bus Adapters ###**, **Host Bus Adapters HW Path**, and **I/O Paths Total** columns differs according to the operating system on which you are running PowerPath and whether your fabric is zoned.

### HBA display: Nonswitched environments or 1:1 zoning configurations

All the sample **powermt display** output shown previously is for a nonswitched topology or a configuration with 1:1 zoning. With 1:1 zoning, an HBA has access to 1 Fibre Channel interface. In addition, one logical device is assigned per Fibre Channel interface. Recall the previous example of **powermt display** (for Symmetrix systems):

```

Symmetrix logical device count=64
=====
- Host Bus Adapters - ---- I/O Paths --- ----- Stats -----
### HW Path          Summary Total Dead   IO/Sec Q-IOs Errors
=====
2302 port2\path0      optimal    64     0     -     0     0
2303 port3\path0      optimal    64     0     -     0     0
=====

```

In this example, there are 2 HBAs (numbered 2302 and 2303), each with 64 paths. Because of the 1:1 zoning restrictions, each HBA has access to one Fibre Channel interface. Consequently, there are 64 paths per HBA and a total storage-system logical device count of 64.

The values in the **I/O Paths Total** column are the number of logical devices visible to a particular HBA. For nonswitched topologies or topologies with 1:1 zoning, the value in the **I/O Paths Total** column is less than or equal to the value in the **logical device count** field.

### HBA display: Switched environments with multiple port zoning

In a switched topology, HBAs may be zoned to see multiple storage-system ports. **powermt display** cannot display virtual HBAs (multiple instances of an HBA, one per port). As a result, the sum of the values in the **I/O Paths Total** column may exceed the **logical device count** value. In addition, in the **Host Bus Adapters** column, only one instance of an HBA is shown for all paths in the **I/O Paths Total** column. These discrepancies are due to Fibre Channel switch zoning of HBAs to multiple storage-system ports.

The exception is HP-UX, on which virtual HBAs do display, one instance for each port, when you run **powermt display**. For HP-UX, continue reading at [“Multiple port zoning on HP-UX” on page 36](#). For all other platforms, continue reading below.

In the following sample output, the **I/O Paths Total** column shows the number of paths associated with each HBA. Although there are multiple storage-system ports visible to each HBA, the output displays one instance of each HBA and its hardware path (in the **Host Bus Adapters** column). In the figure, HBAs 2303 and 2304 are fibre connections that can access identical logical devices; both display a value of 128 in the **I/O Paths Total** column. In reality, there are 32 logical devices. Due to zoning, however, these 32 logical devices are shared among several ports (in this case, 4 ports); therefore, the number of logical devices in the output is 128. This is an example of 1:4 zoning; each HBA has access to 4 Fibre Channel interfaces:

```
powermt display
Symmetrix logical device count=32
=====
- Host Bus Adapters - ---- I/O Paths --- ----- Stats -----
### HW Path Summary Total Dead IO/sec Q-IOs Errors
=====
2303 port3\path0 optimal 128 0 - 0 0
2304 port4\path0 optimal 128 0 - 0 0
=====
```

To see individual paths for each HBA, use the **paths** argument. **powermt display bus** lists all paths between the HBAs on a host and all the storage-system ports that the HBAs can access. See the following sample output.

```
powermt display bus
Symmetrix logical device count=32
=====
- Host Bus Adapters - ---- Storage System ---- - I/O Paths ---
### HW Path ID Interface Total Dead
=====
2303 port3\path0 000000052360 FA 3aA 32 0
2303 port3\path0 000000052360 FA 3bA 32 0
2303 port3\path0 000000052360 FA 14aA 32 0
2303 port3\path0 000000052360 FA 14bA 32 0
2304 port4\path0 000000052360 FA 3aA 32 0
2304 port4\path0 000000052360 FA 3bA 32 0
2304 port4\path0 000000052360 FA 14aA 32 0
2304 port4\path0 000000052360 FA 14bA 32 0
=====
```

In the **powermt display bus** output, HBA 2303 is listed four times, once for each visible storage-system interface. There are 32 paths for each interface. Compare this output to the preceding output of **powermt display**, in which HBA 2303 is listed once, with a total of 128 paths (32 paths for each of 4 interfaces).

## Multiple port zoning on HP-UX

On HP-UX, **powermt display** shows multiple HBA instances; for example:

```
powermt display
Symmetrix logical device count=240
=====
- Host Bus Adapters - ---- I/O Paths --- ----- Stats -----
### HW Path Summary Total Dead IO/sec Q-IOs Errors
=====
3 0/8/0/0 optimal 120 0 - 0 0
4 1/4/0/0 optimal 120 0 - 0 0
13 0/2/0/0.1.16.0.0 optimal 120 0 - 0 0
15 0/2/0/0.1.17.0.0 optimal 120 0 - 0 0
17 0/4/0/0.1.16.0.0 optimal 120 0 - 0 0
19 0/4/0/0.1.17.0.0 optimal 120 0 - 0 0
21 0/12/0/0.1.16.0.0 optimal 120 0 - 0 0
23 0/12/0/0.1.17.0.0 optimal 120 0 - 0 0
25 1/10/0/0.1.26.0.0 optimal 120 0 - 0 0
=====
```

```

27  1/10/0/0.1.27.0.0  optimal  120    0    -    0    0
29  1/12/0/0.1.26.0.0  optimal  120    0    -    0    0
31  1/12/0/0.1.27.0.0  optimal  120    0    -    0    0

```

In the sample HP-UX output previously, note the following:

- ◆ There are a total of 240 devices.
- ◆ HBAs 3, 13, 15, 17, 19, 21, and 23 connect to the same 120 devices.
- ◆ HBAs 4, 25, 27, 29, and 31 connect to the same 120 devices.
- ◆ HBAs 13 and 15 are the same HBA; however, because of the SAN configuration imposed by switches, they appear as two separate HBAs. The HBA's **HW Path** information verifies that these pairs are the same HBA. For example, the **Host Bus Adapters ###** column displays the HBA as two separate instances, number 13 and 15. The **Host Bus Adapters HW Path** column shows that these two instances are the same HBA, differing only in the target ID given by the switch.

Similarly, each of the following virtual HBA pairs actually are single HBAs: 17 and 19, 21 and 23, 25 and 27, and 29 and 31.

## Fields in the HBA display

[Table 10 on page 37](#) explains the fields in the HBA display.

**Table 10** Fields in the HBA display (page 1 of 2)

Field	Value	Description
<code>&lt;storage_ system_type&gt;</code> logical device count	Non-negative integer	Total number of unique logical devices from all storage devices of a given type that are configured by PowerPath and that this host can access. The maximum value is platform-specific. Each <b>powermt display</b> table contains information about one type of storage system, such as Symmetrix.
Host Bus Adapters ###	Non-negative integer	PowerPath number for the HBA. When the <b>dev</b> option is used, the output of <b>powermt display</b> identifies the HBA by this same HBA number. This number is preserved across boots but is not preserved after configuration changes.
Host Bus Adapters HW Path	Alphanumeric string	Displays the path to the host bus adapter. <b>Windows:</b> <code>port#\path#</code> , where # is the integer displayed in Properties in Disk Administrator or Disk Management; for example, <code>port1\path0</code> . <b>UNIX:</b> Operating-system-specific hardware descriptive name for the bus: <b>AIX:</b> The name of the device that is the parent of the <code>hdisk</code> ; for example, <code>scsi1</code> . <b>HP-UX:</b> Hardware path as defined by the <code>ioscan</code> command; for example, <code>10/4.2.0</code> . <b>Solaris:</b> Modification of the <code>/dev/rdisk</code> symbolic link target; for example, <code>sbus@1f,0/fcaw@3,0</code> . <b>Linux:</b> HBA name registered by the HBA driver, for example, <code>lpfc</code> . Registration of the HBA name is optional. For those drivers that do not register a name, Linux displays the string information returned from the driver. This string is truncated to 24 characters without parsing. To view the entire name of the HBA, under the <code>/proc/scsi</code> directory, change to the subdirectory containing copies of the HBA drivers for the HBA SCSI bus instance you are investigating. Use the <code>cat</code> command to view the full name of the HBA.
I/O Paths Summary	optimal degraded failed	Status of the paths originating from this HBA: <ul style="list-style-type: none"> <li>• <i>optimal</i> means all paths are alive (usable).</li> <li>• <i>degraded</i> means one or more, but not all, paths from this HBA are dead (not usable).</li> <li>• <i>failed</i> means all paths are dead and no data is passing through this HBA.</li> </ul>

**Table 10** Fields in the HBA display (page 2 of 2)

Field	Value	Description
I/O Paths Total	Non-negative integer	Total number of paths that originate from this HBA. The maximum number of logical devices supported by one HBA is platform specific. The total number of paths may exceed the number of logical devices in a complex SAN topology, due to zoning.
I/O Paths Dead	Integer in the range 0 - <i>I/O Paths Total</i>	Total number of paths originating from this HBA that are dead (not usable).
Stats IO/sec	Non-negative integer	This field is blank for <b>powermt display</b> , unless it is used with the <b>every</b> parameter. Subsequent <b>powermt display</b> iterations display the average number of I/Os sent across this bus each second.
Stats Q-IOs	Non-negative integer	Total number of I/O operations under way to this path. This is the total number of I/O requests to this device that have not completed. The sum of in-progress I/Os for all paths should equal the number of in-progress I/Os for the PowerPath device.  <b>Note:</b> Some I/O may occasionally go down paths that are: dead, in manual standby or autostandby mode, or non-optimal ALUA.  This is because Q-IOs is inclusive of all types of I/Os filtered by PowerPath, not just application reads and writes. Some platforms and products have daemons or services that occasionally "ping" all paths with ioctls - these are generally unconditionally routed to the "request" path, and not redirected to optimal or active paths.
Stats Errors	Non-negative integer	Total number of times this path transitioned from alive to dead. This is always equal to or less than the total number of HBA I/O path errors. Note that the error count may be different for different LUNs using the same physical paths. This value is cleared at boot time, when <b>powermt restore</b> executes, when <b>powermt set port_disable=on</b> executes, or when <b>powermt unmanage</b> executes.  PowerPath periodic autorestore, which is performed automatically whenever Plug and Play takes a device offline or brings a device back online, does not clear the error count.
Storage System ID	Hexadecimal value or alphanumeric string	Identification number for the storage system on which the logical device is located. This is the array serial number/ID typically set by EMC manufacturing. Note that the formatting or naming scheme of the array serial number/ID will look different in production environments.
Storage System Interface	Alphanumeric string	Storage-system interface. For Symmetrix systems, this has three parts: <ul style="list-style-type: none"> <li>• Interface type: Fibre Channel (FA) or SCSI (SA)</li> <li>• Interface address: integer in the range 1 to 16</li> <li>• Interface port: [abcd][AB]</li> </ul> For VNX and CLARiiON systems, this has one part: <ul style="list-style-type: none"> <li>• Interface port: SP [A-B][0-3]</li> </ul> For HP StorageWorks EVA this has one part: <ul style="list-style-type: none"> <li>• Interface port: [A-B][1-2]</li> </ul> A, B denotes the controller. 1, 2 denotes the port. For Hitachi Lightning, Hitachi TagmaStore and HP xp systems, this is the port name on the array. For IBM ESS systems, this is a two-digit hexadecimal port ID that can be decoded to a Bay/Adapter/Port on the ESS system. (IBM ESS command-line interface documentation describes how to decode the port ID.) For EMC Invista and VPLEX systems, a two-digit hexadecimal number that can be decoded to a virtual target on the Invista system.
Storage System Wt_Q	Non-negative integer	Maximum number of write I/O requests that will be serviced before the I/O queue checks for any outstanding read I/O.

## Devices display

The second form of the command (**powermt display dev**) displays devices, as in the following example, for one of the platforms on which PowerPath runs):

```
# powermt display dev=emcpoweraa
Pseudo name=emcpoweraa
CLARiiON ID=FNM00083800185 [lppb100]
Logical device ID=60060160F21A22002A0E84DD2610DF11 [LUN 2302]
state=alive; policy=CLAROpt; queued-IOs=0
Owner: default=SP A, current=SP A      Array failover mode: 4
=====
----- Host ----- - Stor - -- I/O Path -- -- Stats ---
### HW Path          I/O Paths   Interf.  Mode   State  Q-IOs  Errors
=====
  1 qla2xxx           sdf       SP A0   active  alive   0       0
  2 qla2xxx           sdk       SP A1   active  alive   0       0
  1 qla2xxx           sdw       SP B0   active  alive   0       0
  2 qla2xxx           sdbu     SP B1   active  alive   0       0
=====
```

The following is second example of the devices display output:

```
powermt display dev=all
Pseudo name=emcpower1a
Symmetrix ID=000192601719
Logical device ID=015F
state=alive; policy=SymmOpt; queued-IOs=0
=====
----- Host ----- -Stor- -I/O Path- -Stats-
### HW Path          I/O Paths   Interf.  Mode  State  Q-IOs  Errors
=====
3074 pci@1e,600000/SUNW,qlc@3/fp@0,0 c2t50000972081ADD18d4s0 FA 7eA active  alive  0  0
3072 pci@1e,600000/SUNW,qlc@3,1/fp@0,0 c3t50000972081ADD18d4s0 FA 7eA asb:iopf alive 0  0
3072 pci@1e,600000/SUNW,qlc@3,1/fp@0,0 c3t50000972081ADD1Cd4s0 FA 8eA asb:prox alive 0  0
3074 pci@1e,600000/SUNW,qlc@3/fp@0,0 c2t50000972081ADD1Cd4s0 FA 8eA active  alive  0  0
=====
```

Table 11 on page 40 describes the fields in the devices display.

**Table 11** Fields in the devices display (page 1 of 4)

Field	Value	Description
Pseudo name	Alphanumeric string	Platform-specific value assigned by PowerPath to the PowerPath device. See <a href="#">“Command syntax: dev argument specification” on page 23</a> for pseudo device naming formats for each platform.
<storage_system_type> ID	Hexadecimal value or alphanumeric string	<p>Identification number for the storage system on which the logical device is located. Each <b>powermt display</b> table contains information about one type of storage system, such as Symmetrix.</p> <p>For VNX and CLARiiON systems only, the user-configurable storage group name, if available, is by default displayed in brackets after this identification number. To suppress display of the name, set the <b>PP_SHOW_CLAR_LUN_NAMES</b> environment variable to <b>false</b> or <b>FALSE</b>. (To re-enable the default behavior, either remove the variable or set it to <b>true</b> or <b>TRUE</b>.) An example using Bourne shell syntax to suppress the display of LUN names is:</p> <pre>PP_SHOW_CLAR_LUN_NAMES=false export PP_SHOW_CLAR_LUN_NAMES</pre> <p>See also <a href="#">“powermt display options” on page 58</a>.</p>
Logical device ID	Hexadecimal value or alphanumeric string	<p>Identification number for the logical device. Each logical device on each storage system has a unique ID. Each storage system, however, uses the same storage-system ID. Together, storage-system ID and logical-device ID create a unique ID for every logical device in the world. On a VNX and CLARiiON system, the logical-device ID is a 32-digit number. This is <i>not</i> the same as the standard device identifier used by Navisphere.</p> <p>For VNX and CLARiiON systems only, the user-assignable LUN name, if available, is by default displayed in brackets after this identification number. To suppress display of the LUN name, set the <b>PP_SHOW_CLAR_LUN_NAMES</b> environment variable to <b>false</b> or <b>FALSE</b>. (To re-enable the default behavior, either remove the variable or set it to <b>true</b> or <b>TRUE</b>.) An example using Bourne shell syntax to suppress the display of LUN names is:</p> <pre>PP_SHOW_CLAR_LUN_NAMES=false export PP_SHOW_CLAR_LUN_NAMES</pre> <p>See also <a href="#">“powermt display options” on page 58</a>.</p>
mode	active standby asb:prox asb:iopf unlic	<ul style="list-style-type: none"> <li>• <b>active</b> specifies a path that will be actively used for I/O for multipathing.</li> <li>• <b>standby</b> specifies a path that has been manually placed in standby mode by the user. This path will not be used unless all other active paths are dead. However, note that with ALUA-enabled and active-passive devices, standby paths will be used in favor of active/non-optimized and passive paths, respectively.</li> <li>• <b>asb</b> (for autostandby) specifies a path that has been placed into standby automatically either using the VPLEX proximity autostandby algorithm (<b>asb:prox</b>) or the intermittent I/O failure algorithm (<b>asb:iopf</b>). This path will not be used unless all other active paths are dead. <b>asb</b> path mode settings do not persist across reboots.</li> <li>• <b>unlic</b> indicates that unlicensed PowerPath is running for a Symmetrix, VNX, and CLARiiON storage system (no license key has been installed). In this scenario, all paths are marked unlicensed except one path to each SP. Unlicensed paths cannot become candidates for path failover.</li> </ul>
state	alive dead	<p>State of the PowerPath device:</p> <ul style="list-style-type: none"> <li>• PowerPath marks a PowerPath device <i>dead</i> if it returns certain types of I/O errors and is judged unusable. Once a PowerPath device is marked dead (and until it is restored), PowerPath returns subsequent I/O requests with a failure status, without forwarding them to the associated logical device. This prevents further, unrecoverable corruption and allows the user to perform data recovery if needed.</li> <li>• If a PowerPath device is <i>alive</i>, either PowerPath never marked it dead, or it was marked dead but restored with <b>powermt restore</b> or an automatic path test.</li> </ul>
Owner	default=SP <i>x</i> , current=SP <i>x</i>	Default and current owners of the logical device (VNX and CLARiiON systems only).



**Table 11** Fields in the devices display (page 2 of 4)

Field	Value	Description
policy	Adaptive (ad) CLAROpt (co) LeastBlocks (lb) Least I/O (li) Request (re) RoundRobin (rr) Stream I/O (si) SymmOpt (so)	Current load-balancing and failover policy for the device. Use <b>powermt set policy</b> to change the policy. Refer to “ <a href="#">powermt set policy</a> ” on page 99.  <b>Note:</b> As of version 5.5, the BasicFailover (BF) and NoRedirect (NR) policies have been removed from the <b>powermt set policy</b> command usage. In subsequent releases, it will not be possible to manually set these policies. However, on hosts without a valid PowerPath license, BF and NR will continue to be the default policies for EMC and third-party arrays, respectively.
queued- IOs	Non-negative integer	Number of I/O requests queued to this PowerPath device.
write_ throttle	<b>on</b> or <b>off</b>	Write-throttling setting for this logical device.
Array failover mode	1, 3, or 4	Failover mode set on a VNX and CLARiiON array. This field displays only when PowerPath is managing VNX and CLARiiON LUNs. You set this mode through the Navisphere software. Stop all I/O on the host before changing the failover mode. After changing the failover mode, rebooting the host is required for PowerPath to recognize the change.  <b>Note:</b> The array failover mode must be the same for all paths that access a single LUN. If two paths access the same LUN, and one path is set to PNR (passive not ready) mode and one to ALUA (asymmetric logical unit access) mode, PowerPath behavior is undefined for that LUN.  Modes supported with PowerPath include: <ul style="list-style-type: none"> <li>• 1—passive not ready (PNR) mode—Supported with all PowerPath platforms.</li> <li>• 3—passive always ready (PAR) mode—Supported only with PowerPath for AIX in specific configurations.</li> <li>• 4—asymmetric logical unit access (ALUA) mode—Supported with PowerPath on some platforms, as described in the E-Lab Interoperability Navigator.</li> </ul> You can suppress the <b>powermt</b> failover mode display to ensure compatibility with existing scripts. To suppress the display, set the <b>PP_SHOW_ALUA_FAILOVER_MODE</b> environment variable to <b>false</b> . Refer to the VNX and CLARiiON documentation for more information about the array failover mode.
Host ###	Non-negative integer	PowerPath number for the HBA. This number is preserved across boots but is not preserved after configuration changes. When the <b>dev</b> option is used, the output of <b>powermt display</b> identifies the HBA by this same HBA number.
Host HW Path	Alphanumeric string	Platform-specific hardware descriptive name for the path. <b>Windows:</b> <b>port#\path#\tgt#\lun#</b> , where # is the integer displayed in Properties in Disk Administrator or Disk Management. For example, port2\path0\tgt6\lun7. (A LUN, or Logical Unit Number, is a logical device.) <b>AIX:</b> The name of the device that is the parent of the hdisk; for example, fscsi1. <b>HP-UX:</b> Hardware path as defined by the <b>ioscan</b> command; for example, 10/4.2.0. <b>Solaris:</b> Modification of the /dev/rdisk symbolic link larger; for example, /pci@8,700000/QLGC,qla@3/sd@0,2. <b>Linux:</b> HBA name registered by the HBA driver. Registration of the HBA name is optional. For those drivers that do not register a name, Linux displays the string information returned from the driver, truncated to 24 characters without parsing. To view the entire HBA name, under the /proc/scsi directory, change to the subdirectory containing copies of the HBA drivers for the HBA SCSI bus instance you are investigating. Use the <b>cat</b> command to view the HBA’s full name. On Linux SLES 11 and later, UNKNOWN may appear in this column following a path failure. This occurs because the corresponding native device has been removed from the operating system kernel. Refer to the Release Notes for more information.

**Table 11** Fields in the devices display (page 3 of 4)

Field	Value	Description
Host I/O Path	Alphanumeric string	<p>The platform-specific device name for the path. Refer to <a href="#">“Command syntax: dev argument specification” on page 23</a>.</p> <p>On HP-UX and Solaris, if new devices are added, <b>powermt display every=&lt;#seconds&gt;</b> may display <b>unknown</b> in place of actual device names, due to the creation of device-name mappings at startup. If this occurs, restart <b>powermt display every=&lt;#seconds&gt;</b>.</p> <p>On Linux SLES 11 and later, UNKNOWN may appear in this column following a path failure. This occurs because the corresponding native device has been removed from the operating system kernel. Refer to the Release Notes for more information.</p>
Stor Interf.	Alphanumeric string	<p>Storage-system interface.</p> <p>For Symmetrix systems, this has three parts:</p> <ul style="list-style-type: none"> <li>• Interface type: Fibre Channel (FA) or SCSI (SA)</li> <li>• Interface address: integer in the range 1 to 16</li> <li>• Interface port: [abcd][AB]</li> </ul> <p>For VNX and CLARiiON systems, this has one part:</p> <ul style="list-style-type: none"> <li>• Interface port: SP [A-B][0-3]</li> </ul> <p>For HP StorageWorks EVA this has one part:</p> <ul style="list-style-type: none"> <li>• Interface port: [A-B][1-2]</li> </ul> <p>A, B denotes the controller. 1, 2 denotes the port.</p> <p>For Hitachi Lightning, Hitachi TagmaStore, and HP xp systems, this is the port name on the array.</p> <p>For IBM ESS systems, this is a two-digit hexadecimal port ID that can be decoded to a Bay/Adapter/Port on the ESS system. (IBM ESS command-line interface documentation describes how to decode the port ID.)</p> <p>For EMC Invista and VPLEX systems, a two-digit hexadecimal number that can be decoded to a virtual target on the Invista system.</p>
I/O Path Mode	active standby asb:prox asb:iopf unlic	<p>Current path mode:</p> <ul style="list-style-type: none"> <li>• <b>active</b> indicates this path can accept I/O. Load balancing is performed for a device with more than one active path, based on the load-balancing and failover policy set for the device. On active-passive storage systems, I/O is load balanced across paths to a single SP for each logical device.</li> <li>• <b>standby</b> indicates this path is held in reserve. Being set to standby does not mean a path will not be used. It only means the weight of the path is heavily adjusted to preclude its use in normal operations. A standby path still can be selected if it is the best path for a request. For example, if the active paths are so heavily loaded that the weighting factor in favor of the active path is overcome, the I/O will be sent to the standby path.</li> <li>• <b>asb</b> (for autostandby) specifies a path that has been placed into standby automatically either using the VPLEX proximity autostandby algorithm (<b>asb:prox</b>) or the intermittent I/O failure algorithm (<b>asb:iopf</b>). This path will not be used unless all other active paths are dead. <b>asb</b> path mode settings do not persist across reboots.</li> <li>• <b>unlic</b> indicates that unlicensed PowerPath is running for a Symmetrix, VNX, and CLARiiON storage system (no license key has been installed). In this scenario, all paths are marked unlicensed except one path to each SP. Unlicensed paths cannot become candidates for path failover.</li> </ul> <p>Use <b>powermt set mode</b> to change the mode. Refer to <a href="#">“powermt set mode” on page 92</a>.</p>

**Table 11** Fields in the devices display (page 4 of 4)

Field	Value	Description
I/O Path State	alive dead	Current path state: <ul style="list-style-type: none"> <li><b>alive</b> indicates the path is usable: PowerPath can direct I/O to this path.</li> <li><b>dead</b> indicates the path is not usable: PowerPath will not direct I/O to this path. After marking the path dead and incrementing the Errors count, PowerPath tests the path to see whether it is usable. If the test succeeds, PowerPath marks the path alive; the path is then available for I/O. If the test fails, the path remains dead, and PowerPath ignores it for subsequent I/O operations. If all the paths to a logical device are dead, PowerPath retests each path.</li> </ul>
Stats Q-IOs	Non-negative integer	Total number of I/O operations under way to this path. This is the total number of I/O requests to this device that have not completed. The sum of in-progress I/Os for all paths should equal the number of in-progress I/Os for the PowerPath device.  <b>Note:</b> Some I/Os may occasionally go down paths that are: dead, in manual standby or autostandby mode, or non-optimal ALUA.  This is because Q-IOs is inclusive of all types of I/Os filtered by PowerPath, not just application reads and writes. Some platforms and products have daemons or services that occasionally "ping" all paths with ioctls - these are generally unconditionally routed to the "request" path, and not redirected to optimal or active paths.
Stats Errors	Non-negative number	Total number of times this path transitioned from alive to dead. This is always equal to or less than the total number of HBA I/O path errors. Note that the error count may be different for different LUNs using the same physical paths. This value is cleared at boot time, when <b>powermt restore</b> executes, when <b>powermt set port_disable=on</b> executes, or when <b>powermt unmanage</b> executes.  PowerPath periodic autorestore, which is performed automatically whenever Plug and Play takes a device offline or brings a device back online, does not clear the error count.

## Specifying the wide or width option

**Note:** The wide option is not supported on Windows platforms.

You can change the width of the output by specifying the **wide** or **width** option. The following two examples show how setting the width option changes the output display. The first example uses the default column width (85) while the second example shows the width set to 100:

```
# powermt display dev=a width=85
Pseudo name=emcpowera
Symmetrix ID=000190300848
Logical device ID=05A6
state=alive; policy=SymmOpt; queued-IOs=0
=====
----- Host ----- - Stor -- -- I/O Path -- -- Stats ---
### HW Path          I/O Paths   Interf.  Mode   State  Q-IOs Errors
=====
   1 qla2xxx          sdbd      FA  2cB   active  alive   0    0
   1 qla2xxx          sdbp      FA  1cB   active  alive   0    0
   2 qla2xxx          sdcz      FA  2cB   active  alive   0    0
   2 qla2xxx          sddk      FA  1cB   active  alive   0    0

# powermt display dev=0 width=100
Pseudo name=emcpower0a
CLARiiON ID=WRE00021400504
Logical device ID=600601760B09000054F4356E2B17D711
state=alive; policy=BasicFailover; priority=0; queued-IOs=0
Owner: default=SP B, current=SP B
```

```

=====
----- Host ----- - Stor ----- -- I/O Path -- -- Stats ---
### HW Path          I/O Paths    Interf.    Mode    State    Q-IOs Errors
=====
2308 pci@4,4000/lpfc@4      c2t0d3s0    SP A1      unlic   alive    0      0
2308 pci@4,4000/lpfc@4      c2t1d3s0    SP A0      unlic   alive    0      0
2305 pci@4,2000/fibre-channel@1  c3t0d3s0    SP B1      active  alive    0      0

```

## powermt display alua

**Purpose** Display the ALUA (Asymmetric Logical Unit Access) state for all the paths to an ALUA device.

**Syntax** `powermt display alua dev=<device>|all [class=<class>|all] [every=<#seconds>] [wide|width=<#columns>]`

**Description** `powermt display alua` displays the current asymmetric access state for each path to an ALUA LUN:

- ◆ Active/optimized
- ◆ Active/non-optimized
- ◆ Standby
- ◆ Unavailable
- ◆ Offline
- ◆ Transitioning

An ALUA LUN is reachable across all available storage processors/directors, but only one storage processor/director owns the LUN at any given time. Paths associated with the owning storage processor/director have an *active/optimized* asymmetric access state, while paths associated with the non-owning storage processors/directors have an *active/non-optimized* asymmetric access state.

Paths in an active/non-optimized asymmetric access state typically provide lower performance than those in an active/optimized asymmetric access state. During normal load balancing, PowerPath will distribute I/O among paths in the active/optimized asymmetric access state only for optimal performance. Paths in the active/non-optimized asymmetric access state will be used only when all paths in the active/optimized asymmetric access state have failed.

### Arguments

`dev=<device>|all`  
Specifies a valid ALUA device or all ALUA devices.

`class={symm|vnx|clariion|vplex|invista|netapp|ess|xtremio|hitachi|hpxp|generic|all}`

Limits the command to the specified type of storage system. **all** specifies all storage-system types. The default is **all**. For descriptions of all the class types, see [“Command syntax: class argument specification” on page 24](#).

`every=<#seconds>`  
Specifies a valid interval to display the data. If a value is not specified, then the information is displayed once. The value can be an integer in the range of 1 to 86400.

The *seconds* value is the minimum time between refreshes; the actual time is affected by the overall system load. On busy systems, display updates can be less frequent than specified.

`wide|width=<#columns>`

**Note:** The wide option is not supported on Windows platforms.

**wide** sets the line length to the width of the window and **width** sets a range within which the `powermt display` output appears. The default is 80. The acceptable range is 80 to 160 columns. The output widens or narrows within the range depending on the

width of the window. This value applies only to the specific output; it is not persistent. You can set a persistent value for the **width=<#columns>** option by using the environment variable **PP\_DISPLAY\_WIDTH**. This variable changes the width setting from 80 (the default) to the value you specify. If the environment variable value differs from the value you set with the command line, the command line value takes precedence.

If you set the **every=<#seconds>** option along with the **width=<#columns>** option, the **powermt display** output widens or narrows to accommodate changes you make to the window width.

## Examples

```
# powermt display alua dev=8
Pseudo name=emcpower8a
CLARiion ID=APM00090700078 [lclb039]
Logical device ID=600601601C102400AC335B99B796E011 [Thin LUN 166]
state=alive; policy=CLAROpt; queued-Ios=0
Owner: default=SP A, current=SP A      Array failover mode: 4
=====
---- Host ----- - Stor - ----- I/O Path ----- Stats -
### I/O Paths      Interf. ALUA State      Mode State      Errors
=====
3072 c2t5006016041E0A22Ed9s0 SP A0      Active/optimized      active alive 0
3072 c2t5006016841E0A22Ed9s0 SP B0      Active/non-optimized  active alive 0
3074 c3t5006016141E0A22Ed9s0 SP A1      Active/optimized      active alive 0
3074 c3t5006016941E0A22Ed9s0 SP B1      Active/non-optimized  active alive 0

# powermt display alua dev=8 class=clariion every=2
Pseudo name=emcpower8a
CLARiion ID=APM00090700078 [lclb039]
Logical device ID=600601601C102400AC335B99B796E011 [Thin LUN 166]
state=alive; policy=CLAROpt; queued-Ios=0
Owner: default=SP A, current=SP A      Array failover mode: 4
=====
---- Host ---- - Stor - ----- I/O Path ----- -- Stats ---
### I/O Paths  Interf.  ALUA State      Mode  State  Errors
=====
3072 c2t5006016041E0A22Ed9s0 SP A0      Active/optimized      active alive 0
3072 c2t5006016841E0A22Ed9s0 SP B0      Active/non-optimized  active alive 0
3074 c3t5006016141E0A22Ed9s0 SP A1      Active/optimized      active alive 0
3074 c3t5006016941E0A22Ed9s0 SP B1      Active/non-optimized  active alive 0

Pseudo name=emcpower8a
CLARiion ID=APM00090700078 [lclb039]
Logical device ID=600601601C102400AC335B99B796E011 [Thin LUN 166]
state=alive; policy=CLAROpt; queued-Ios=0
Owner: default=SP A, current=SP A      Array failover mode: 4
=====
---- Host ---- - Stor - ----- I/O Path ----- -- Stats ---
### I/O Paths  Interf.  ALUA State      Mode  State  Errors
=====
3072 c2t5006016041E0A22Ed9s0 SP A0      Active/optimized      active alive 0
3072 c2t5006016841E0A22Ed9s0 SP B0      Active/non-optimized  active alive 0
3074 c3t5006016141E0A22Ed9s0 SP A1      Active/optimized      active alive 0
3074 c3t5006016941E0A22Ed9s0 SP B1      Active/non-optimized  active alive 0
```

<repeated every two seconds>

## powermt display hba\_mode

**Purpose** Monitor HBA modes.

**Syntax** `powermt display hba_mode`

**Description** `powermt display hba_mode` displays information about all HBAs configured on the host. The information displayed includes the HBA mode, which is either enabled or disabled.

When the mode is enabled, the HBA is active and working as it normally would within a PowerPath environment. When the mode is disabled, it means the HBA was disabled previously using the `powermt disable` command.

To enable or disable an HBA, refer to [“powermt disable hba” on page 32](#) or [“powermt enable hba” on page 76](#).

**Detailed description** The following sample output for `powermt display hba_mode` shows two enabled HBAs.

The summary display information does not list any arrays if there are no devices for an array type. In the following example display, the Hitachi, Invista, HP, and Ess summary counts of zero does not appear.

### PowerPath 5.7 Linux Example Output:

```
# powermt display hba_mode
Symmetrix logical device count=10
CLARiiON logical device count=11
VNX logical device count=6
=====
----- Host Bus Adapters -----          I/O Paths -----  Stats
###  HW Path                               Summary   Total   Dead  Q-IOs Mode
=====
      1 qla2xxx                             optimal    54     0     0  Enabled
      2 qla2xxx                             optimal    54     0     0  Enabled
```

### PowerPath 5.3 Solaris Example Output:

```
powermt display hba_mode
Symmetrix logical device count=21
CLARiiON logical device count=4
Hitachi logical device count=8
HP xp logical device count=0
Ess logical device count=0
VPLEX logical device count=0
Invista logical device count=0
HP HSx logical device count=0
=====
----- Host Bus Adapters -----          I/O Paths -----  Stats --
###  HW Path                               Summary   Total   Dead  Q-IOs Mode
=====
3072 pci@1f,700000/fibre-channel@2/fp@0,0  optimal    66     0     1  Enabled
3073 pci@1f,700000/fibre-channel@2,1/fp@0,0  optimal    66     0     2  Enabled
```

### PowerPath 5.5 Solaris Example Output:

```
powermt display hba_mode
Symmetrix logical device count=21
CLARiiON logical device count=4
Hitachi logical device count=8
=====
----- Host Bus Adapters -----          I/O Paths -----  Stats --
```

```

### HW Path Summary Total Dead Q-IOs Mode
=====
3072 pci@1f,700000/fibre-channel@2/fp@0,0 optimal 66 0 1 Enabled
3073 pci@1f,700000/fibre-channel@2,1/fp@0,0 optimal 66 0 2 Enabled

```

Table 12 on page 48 describes the fields in the `hba_mode` display.

**Table 12** Fields in the `hba_mode` display

Field	Value	Description
Host Bus Adapters ###	Non-negative integer	PowerPath number for the HBA. When the <b>dev</b> option is used, the output of <b>powermt display</b> identifies the HBA by this same HBA number. This number is preserved across boots but is not preserved after configuration changes.
Host Bus Adapters HW Path	Alphanumeric string	Displays the path to the host bus adapter. <b>Solaris:</b> Modification of the <code>/dev/rdsk</code> symbolic link target; for example, <code>sbus@1f,0/fcaw@3,0</code> .
I/O Paths Summary	optimal degraded failed	Status of the paths originating from this HBA: <ul style="list-style-type: none"> <li>• <b>optimal</b> means all paths are alive (usable).</li> <li>• <b>degraded</b> means one or more, but not all, paths from this HBA are dead (not usable).</li> <li>• <b>failed</b> means all paths are dead and no data is passing through this HBA.</li> </ul>
I/O Paths Total	Non-negative integer	Total number of paths that originate from this HBA. The maximum number of logical devices supported by one HBA is platform-specific. The total number of paths may exceed the number of logical devices in a complex SAN topology, due to zoning.
I/O Paths Dead	Integer in the range 0– <i>I/O Paths Total</i>	Total number of paths originating from this HBA that are dead (not usable).
Stats Q-IOs	Non-negative integer	Total number of I/O operations under way to this path. This is the total number of I/O requests to this device that have not completed. The sum of in-progress I/Os for all paths should equal the number of in-progress I/Os for the PowerPath device.
Mode	Enabled Disabled	The mode of the HBA: <ul style="list-style-type: none"> <li>• <b>Enabled</b> means that the HBA is alive (usable)</li> <li>• <b>Disabled</b> means the HBA was disabled previously with the <b>powermt disable</b> command. While an HBA is disabled, I/O is not sent to it and devices configured through it are closed.</li> </ul>



## powermt display latency

**Purpose** Monitor path latency.

**Syntax** `powermt display latency`  
`[dev=<device>|all] [class=<class>|all]`  
`[every=<#seconds>] [wide|width=<#columns>]`

**Description** `powermt display latency` displays information on path latency monitoring, including the last and longest I/O completion times for each path.

The `powermt display latency` output includes the following two columns:

- ◆ **Current** — Time in microseconds it took the last I/O on the path to complete.
- ◆ **Max** — Time in microseconds it took the longest I/O on the path to complete.

Refer [“Detailed description” on page 34](#) to see sample output.

The **Current** and **Max** values can help determine an appropriate threshold value, which you can then set with the `powermt set path_latency_threshold` command.

To turn on path latency monitoring, use the `powermt set path_latency_monitor` command. To view whether path latency monitoring is enabled, use the `powermt display options` command.

### Arguments

`dev=<device>|all`  
 Displays latency information for the specified device or for all devices.

`class={symm|vnx|clariion|vplex|invista|netapp|ess|xtremio|hitachi|hpxp|generic|all}`  
 Limits the command to the specified type of storage system. **all** specifies all storage-system types. The default is **all**. For descriptions of all the class types, see [“Command syntax: class argument specification” on page 24](#).

`every=<#seconds>`  
 Specifies the frequency with which path latency monitoring information displays. If a value is not specified, then the information is displayed once. The value can be an integer in the range of 1 to 86400.

The *seconds* value is the minimum time between refreshes; the actual time is affected by the overall system load. On busy systems, display updates can be less frequent than specified.

`wide|width=<#columns>`

**Note:** The `wide` argument is not supported on Windows platforms.

**wide** sets the line length to the width of the window and **width** sets a range within which the `powermt display` output appears. The default is 80. The acceptable range is 80 to 160 columns. The output widens or narrows within the range depending on the width of the window. This value applies only to the specific output; it is not persistent. You can set a persistent value for the `width=<#columns>` option by using the environment variable `PP_DISPLAY_WIDTH`. This variable changes the width setting from 80 (the default) to the value you specify. If the environment variable value differs from the value you set with the command line, the command line value takes precedence.

If you set the **every=<#seconds>** option along with the **width=<#columns>** option, the **powermt display** output widens or narrows to accommodate changes you make to the window width.

## Fields in the latency display

[Table 13 on page 50](#) describes the fields in the **latency** display.

**Table 13** Fields in the latency display

Field	Value	Description
Host ###	Non-negative integer	PowerPath number for the HBA. This number is preserved across boots but is not preserved after configuration changes. When the <b>dev</b> option is used, the output of <b>powermt display</b> identifies the HBA by this same HBA number.
Host Bus Adapters HW Path	Alphanumeric string	Displays the path to the host bus adapter. <b>Solaris:</b> Modification of the /dev/rdisk symbolic link target; for example, sbus@1f,0/fcaw@3,0.
Storage System ID	Hexadecimal value or alphanumeric string	Identification number for the storage system on which the logical device is located.
Storage System Interface	Alphanumeric string	Storage-system interface. For Symmetrix systems, this has three parts: <ul style="list-style-type: none"> <li>• Interface type: Fibre Channel (FA) or SCSI (SA)</li> <li>• Interface address: integer in the range 1 to 16</li> <li>• Interface port: [abcd][AB]</li> </ul> For VNX and CLARiiON systems, this has one part: <ul style="list-style-type: none"> <li>• Interface port: SP [A-B][0-3]</li> </ul> For HP StorageWorks EVA this has one part: <ul style="list-style-type: none"> <li>• Interface port: [A-B][1-2]</li> </ul> A, B denotes the controller. 1, 2 denotes the port. For Hitachi Lightning, Hitachi TagmaStore and HP xp systems, this is the port name on the array. For IBM ESS systems, this is a two-digit hexadecimal port ID that can be decoded to a Bay/Adaptor/Port on the ESS system. (IBM ESS command-line interface documentation describes how to decode the port ID.) For EMC Invista and VPLEX systems, a two-digit hexadecimal number that can be decoded to a virtual target on the Invista system.
Latency Current	Non-negative integer	In microseconds, the amount of time it took the last I/O on the path to complete.
Latency Max	Non-negative integer	In microseconds, the longest time it took any I/O on the path to complete.

The following output shows the latency information for the specified device.

```
powermt display latency dev=harddisk5
Pseudo name=harddisk5
CLARiiON ID=FNMO0022300718 [Storage Group 1]
Logical device ID=6006016038E02400BCCBAF9428AADE1 [LUN 8]
state=alive; policy=CLAROpt; ; queued-IOs=0
Owner: default=SP A, current=SP A                Array failover mode: 4
=====
----- Host ----- -Stor - I/O Path - - Latency (us)
### HW Path          I/O Paths   Interf.   Mode     State  Current Max
=====
5 port5\path0\tgt0\lun2  c5t0d2     SP A0    active   alive  0 0
5 port5\path0\tgt1\lun2  c5t1d2     SP B1    active   alive  0 0
6 port6\path0\tgt0\lun2  c6t0d2     SP A1    active   alive  0 0
6 port6\path0\tgt1\lun2  c6t1d2     SP B0    active   alive  0 0
7 port7\path0\tgt0\lun2  c7t0d2     SP A0    active   alive  0 0
7 port7\path0\tgt1\lun2  c7t1d2     SP B1    active   alive  0 0
8 port8\path0\tgt0\lun2  c8t0d2     SP A1    active   alive  0 0
8 port8\path0\tgt1\lun2  c8t1d2     SP B0    active   alive  0 0
```

Table 14 on page 51 explains the fields in the `latency dev=device` display.

**Table 14** Fields in the `latency dev=device` display (page 1 of 2)

Field	Value	Description
Host ###	Alphanumeric string	PowerPath number for the HBA. This number is preserved across boots but is not preserved after configuration changes. <b>Solaris:</b> Modification of the <code>/dev/rdisk</code> symbolic link larger; for example, <code>/pci@8,700000/QLGC,qla@3/sd@0,2</code> .
Host Bus Adapters HW Path	Alphanumeric string	Displays the path to the host bus adapter. <b>Solaris:</b> Modification of the <code>/dev/rdisk</code> symbolic link target; for example, <code>sbus@1f,0/fcaw@3,0</code> .
Host I/O Paths	Alphanumeric string	The platform-specific device name for the path. Refer <a href="#">“Command syntax: dev argument specification” on page 23</a> . On HP-UX and Solaris, if new devices are added, <code>powermt display every=&lt;#seconds&gt;</code> may display <b>unknown</b> in place of actual device names, due to the creation of device-name mappings at startup. If this occurs, restart <code>powermt display every=&lt;#seconds&gt;</code> On Linux SLES 11 and later, UNKNOWN may appear in this column following a path failure. This occurs because the corresponding native device has been removed from the operating system kernel. Refer to the Release Notes for more information.
Stor Interf.	Alphanumeric string	Storage-system interface. For Symmetrix systems, this has three parts: <ul style="list-style-type: none"> <li>Interface type: Fibre Channel (FA) or SCSI (SA)</li> <li>Interface address: integer in the range 1 to 16</li> <li>Interface port: <code>[abcd][AB]</code></li> </ul> For VNX and CLARiiON systems, this has one part: <ul style="list-style-type: none"> <li>Interface port: <code>SP [A-B][0-3]</code></li> </ul> For HP StorageWorks EVA this has one part: <ul style="list-style-type: none"> <li>Interface port: <code>[A-B][1-2]</code></li> </ul> A, B denotes the controller. 1, 2 denotes the port. For Hitachi Lightning, Hitachi TagmaStore, and HP xp systems, this is the port name on the array. For IBM ESS systems, this is a two-digit hexadecimal port ID that can be decoded to a Bay/Adaptor/Port on the ESS system. (IBM ESS command-line interface documentation describes how to decode the port ID). For EMC Invista and VPLEX systems, a two-digit hexadecimal number that can be decoded to a virtual target on the Invista system.

**Table 14** Fields in the latency dev=device display (page 2 of 2)

Field	Value	Description
I/O Path Mode	active standby asb:prox asb:iopf unlic	<p>Current path mode:</p> <ul style="list-style-type: none"> <li>• <b>active</b> indicates this path can accept I/O. Load balancing is performed for a device with more than one active path, based on the load-balancing and failover policy set for the device. On active-passive storage systems, I/O is load balanced across paths to a single SP for each logical device.</li> <li>• <b>standby</b> indicates this path is held in reserve. Being set to standby does not mean a path will not be used. It only means the weight of the path is heavily adjusted to preclude its use in normal operations. A standby path still can be selected if it is the best path for a request. For example, if the active paths are so heavily loaded that the weighting factor in favor of the active path is overcome, the I/O will be sent to the standby path.</li> <li>• <b>asb</b> (for autostandby) specifies a path that has been placed into standby automatically either using the VPLEX proximity autostandby algorithm (<b>asb:prox</b>) or the intermittent I/O failure algorithm (<b>asb:iopf</b>). This path will not be used unless all other active paths are dead. <b>asb</b> path mode settings do not persist across reboots.</li> <li>• <b>unlic</b> indicates that unlicensed PowerPath is running for a Symmetrix, VNX, and CLARiiON storage system (no license key has been installed). In this scenario, all paths are marked unlicensed except one path to each SP. Unlicensed paths cannot become candidates for path failover.</li> </ul> <p>Use <b>powermt set mode</b> to change the mode. Refer <a href="#">“Purpose” on page 92</a>.</p>
I/O Path State	alive dead	<p>Current path state:</p> <ul style="list-style-type: none"> <li>• <b>alive</b> indicates the path is usable: PowerPath can direct I/O to this path.</li> <li>• <b>dead</b> indicates the path is not usable: PowerPath will not direct I/O to this path. After marking the path dead and incrementing the Errors count, PowerPath tests the path to see whether it is usable. If the test succeeds, PowerPath marks the path alive; the path is then available for I/O. If the test fails, the path remains dead, and PowerPath ignores it for subsequent I/O operations. If all the paths to a logical device are dead, PowerPath retests each path.</li> </ul>
Latency Current	Non-negative integer	In microseconds, the amount of time it took the last I/O on the path to complete.
Latency Max	Non-negative integer	In microseconds, the longest time it took any I/O on the path to complete.

## powermt display nonvirtual

**Purpose** Display physical volumes for Invista.

**Syntax** `powermt display nonvirtual dev=<device>|all [class=<class>|all]  
[every=<#seconds>] [wide|width=<#columns>]`

**Description** **powermt display nonvirtual** displays, along with the standard **powermt display** output, information on the physical devices associated with Virtual Volumes.

At this time, the information on physical devices associated with Virtual Volumes appears only for Invista in the NOT READY state. For non-Invista devices and Invista devices in the READY state, the standard **powermt display** output appears without additional nonvirtual information.

---

**Note:** Devices set up for Virtual encapsulation with PowerPath Migration Enabler must be in the NOT READY state.

---

Refer to the EMC Invista documentation for more information on setting up Invista devices for PowerPath Migration Enabler. For more information on PowerPath Migration Enabler, refer to the *EMC PowerPath Migration Enabler User Guide*.

### Arguments

`dev=<device>|all`

Displays information about the specified *device*. **all** specifies all devices.

`class={symm|vnx|clariion|vplex|invista|netapp|ess|xtremio|hitachi|hpxp  
|generic|all}`

Limits the command to the specified type of storage system. **all** specifies all storage-system types. The extra information this command displays applies only to Invista storage systems. For all other storage systems, the output is the same as the output for **powermt display**. The default is **all**. For descriptions of all the class types, see [“Command syntax: class argument specification” on page 24](#).

`every=<#seconds>`

Integer in the range 1 to 86400 that specifies, in seconds, how often the display is updated. By default, the display is not updated. The *seconds* value is the minimum time between refreshes; the actual time is affected by the overall system load. On busy systems, display updates can be less frequent than specified.

`wide|width=<#columns>`

---

**Note:** The wide argument is not supported on Windows platforms.

---

**wide** sets the line length to the width of the window and **width** sets a range within which the **powermt display** output appears. The default is 80. The acceptable range is 80 to 160 columns. The output widens or narrows within the range depending on the width of the window. This value applies only to the specific output; it is not persistent. You can set a persistent value for the `width=<#columns>` option by using the environment variable **PP\_DISPLAY\_WIDTH**. This variable changes the width setting from 80 (the default) to the value you specify. If the environment variable value differs from the value you set with the command line, the command line value takes precedence.

If you set the **every=<#seconds>** option along with the **width=<#columns>** option, the **powermt display** output widens or narrows to accommodate changes you make to the window width.

**Detailed description**

The following sample output for **powermt display nonvirtual** shows the output for two Invista Virtual Volumes, one that is associated with a Symmetrix device, and one that is associated with a VNX and CLARiiON device. The information related to the physical device associated with the Virtual Volume appears in bold font.

```
powermt display nonvirtual dev=all class=invista
```

The following output is truncated:

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

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

Table 15 on page 55 describes the fields in the **powermt display nonvirtual** output.

**Table 15** Fields in the nonvirtual display (page 1 of 3)

Field	Value	Description
Pseudo name	Alphanumeric string	Platform-specific value assigned by PowerPath to the PowerPath device. When using PowerPath Migration Enabler, this device is the target device in a migration.
Invista ID	Hexadecimal value or alphanumeric string	Identification number for the Invista or VPLEX system on which the logical device is located.
Logical device ID	Hexadecimal value or alphanumeric string	<p>Identification number for the logical device. Each logical device on each storage system has a unique ID. Each storage system, however, uses the same storage-system ID. Together, storage-system ID and logical-device ID create a unique ID for every logical device in the world. On a VNX and CLARiiON system, the logical-device ID is a 32-digit number. This is <i>not</i> the same as the standard device identifier used by Navisphere.</p> <p>For VNX and CLARiiON systems only, the user-assignable LUN name, if available, is by default displayed in brackets after this identification number. To suppress display of the LUN name, define the environment variable <b>PP_SHOW_CLAR_LUN_NAMES</b> and set it to <b>false</b> or <b>FALSE</b>. (To re-enable the default behavior, either remove the variable or set it to <b>true</b> or <b>TRUE</b>.) An example using Bourne shell syntax to suppress the display of LUN names is:</p> <pre>PP_SHOW_CLAR_LUN_NAMES=false export PP_SHOW_CLAR_LUN_NAMES</pre> <p>See also <a href="#">“powermt display options” on page 58</a>.</p>
NonVirtual Symmetrix ID	Hexadecimal value or alphanumeric string	For Symmetrix systems, the identification number for the Symmetrix system on which the physical device is located.
NonVirtual Device ID	Hexadecimal value or alphanumeric string	For VNX and CLARiiON systems, the identification number associated with the physical device and the VNX and CLARiiON storage system. For Symmetrix systems, the identification number associated with the physical device.
NonVirtual pseudo name	Alphanumeric string	For the physical device associated with the Invista or VPLEX Virtual Volume, the platform-specific value assigned by PowerPath to the device. When using PowerPath Migration Enabler (Virtual encapsulation) to set up a migration, this is the source device name corresponding to the target device name given in the <b>Pseudo name</b> field (in this command's output).
Host ###	Non-negative integer	PowerPath number for the HBA. This number is preserved across boots but is not preserved after configuration changes. When the <b>dev</b> option is used, the output of <b>powermt display</b> identifies the HBA by this same HBA number.
Host HW Path	Alphanumeric string	<p>Platform-specific hardware descriptive name for the path.</p> <p><b>Windows:</b> port#\path#\tgt#\lun#, where # is the integer displayed in Properties in Disk Administrator or Disk Management. For example, port2\path0\tgt6\lun7. (A LUN, or Logical Unit Number, is a logical device.)</p> <p><b>AIX:</b> The name of the device that is the parent of the hdisk; for example, fscsi1.</p> <p><b>HP-UX:</b> Hardware path as defined by the <b>ioscan</b> command; for example, 10/4.2.0.</p> <p><b>Solaris:</b> Modification of the /dev/rdsk symbolic link larger; for example, /pci@8,700000/QLGC,qla@3/sd@0,2.</p> <p><b>Linux:</b> HBA name registered by the HBA driver. Registration of the HBA name is optional. For those drivers that do not register a name, Linux displays the string information returned from the driver, truncated to 24 characters without parsing. To view the entire HBA name, under the /proc/scsi directory, change to the subdirectory containing copies of the HBA drivers for the HBA SCSI bus instance you are investigating. Use the <b>cat</b> command to view the HBA's full name.</p> <p>On Linux SLES 11 and later, UNKNOWN may appear in this column following a path failure. This occurs because the corresponding native device has been removed from the operating system kernel. Refer to the Release Notes for more information.</p>

**Table 15** Fields in the nonvirtual display (page 2 of 3)

Field	Value	Description
Host I/O Path	Alphanumeric string	<p>The platform-specific device name for the path. Refer to <a href="#">“Command syntax: dev argument specification” on page 23</a>.</p> <p>On HP-UX and Solaris, if new devices are added, <b>powermt display every=&lt;#seconds&gt;</b> may display <b>unknown</b> in place of actual device names, due to the creation of device-name mappings at startup. If this occurs, restart <b>powermt display every=&lt;#seconds&gt;</b>.</p> <p>On Linux SLES 11 and later, UNKNOWN may appear in this column following a path failure. This occurs because the corresponding native device has been removed from the operating system kernel. Refer to the Release Notes for more information.</p>
Stor Interf.	Alphanumeric string	<p>Storage-system interface.</p> <p>For Symmetrix systems, this has three parts:</p> <ul style="list-style-type: none"> <li>Interface type: Fibre Channel (FA) or SCSI (SA)</li> <li>Interface address: integer in the range 1 to 16</li> <li>Interface port: [abcd][AB]</li> </ul> <p>For VNX and CLARiiON systems, this has one part:</p> <ul style="list-style-type: none"> <li>Interface port: SP [A-B][0-3]</li> </ul> <p>For HP StorageWorks EVA this has one part:</p> <ul style="list-style-type: none"> <li>Interface port: [A-B][1-2]</li> </ul> <p>A, B denotes the controller. 1, 2 denotes the port.</p> <p>For Hitachi Lightning, Hitachi TagmaStore, and HP xp systems, this is the port name on the array.</p> <p>For IBM ESS systems, this is a two-digit hexadecimal port ID that can be decoded to a Bay/Adapter/Port on the ESS system. (IBM ESS command-line interface documentation describes how to decode the port ID.)</p> <p>For EMC Invista and VPLEX systems, a two-digit hexadecimal number that can be decoded to a virtual target on the Invista system.</p>
I/O Paths Mode	active standby asb:prox asb:iopf unlic	<p>Current path mode:</p> <ul style="list-style-type: none"> <li><b>active</b> indicates this path can accept I/O. Load balancing is performed for a device with more than one active path, based on the load-balancing and failover policy set for the device. On active-passive storage systems, I/O is load balanced across paths to a single SP for each logical device.</li> <li><b>standby</b> indicates this path is held in reserve. Being set to standby does not mean a path will not be used. It only means the weight of the path is heavily adjusted to preclude its use in normal operations. A standby path still can be selected if it is the best path for a request. For example, if the active paths are so heavily loaded that the weighting factor in favor of the active path is overcome, the I/O will be sent to the standby path.</li> <li><b>asb</b> (for autostandby) specifies a path that has been placed into standby automatically either using the VPLEX proximity autostandby algorithm (<b>asb:prox</b>) or the intermittent I/O failure algorithm (<b>asb:iopf</b>). This path will not be used unless all other active paths are dead. <b>asb</b> path mode settings do not persist across reboots.</li> <li><b>unlic</b> indicates that unlicensed PowerPath is running for a Symmetrix, VNX, and CLARiiON storage system (no license key has been installed). In this scenario, all paths are marked unlicensed except one path to each SP. Unlicensed paths cannot become candidates for path failover.</li> </ul> <p>Use <b>powermt set mode</b> to change the mode. Refer to <a href="#">“Purpose” on page 92</a>.</p>



**Table 15** Fields in the nonvirtual display (page 3 of 3)

Field	Value	Description
I/O Paths State	alive dead	Current path state: <ul style="list-style-type: none"> <li>• <b>alive</b> indicates the path is usable: PowerPath can direct I/O to this path.</li> <li>• <b>dead</b> indicates the path is not usable: PowerPath will not direct I/O to this path. After marking the path dead and incrementing the Errors count, PowerPath tests the path to see whether it is usable. If the test succeeds, PowerPath marks the path alive; the path is then available for I/O. If the test fails, the path remains dead, and PowerPath ignores it for subsequent I/O operations. If all the paths to a logical device are dead, PowerPath retests each path.</li> </ul>
Stats Q-IOs	Non-negative integer	Total number of I/O operations under way to this path. This is the total number of I/O requests to this device that have not completed. The sum of in-progress I/Os for all paths should equal the number of in-progress I/Os for the PowerPath device.
Stats Errors	Non-negative number	Total number of times this path transitioned from alive to dead. This is always equal to or less than the total number of HBA I/O path errors. Note that the error count may be different for different LUNs using the same physical paths. This value is cleared at boot time, when <b>powermt restore</b> executes, when <b>powermt set port_disable=on</b> executes, or when <b>powermt unmanage</b> executes. PowerPath periodic autorestore, which is performed automatically whenever Plug and Play takes a device offline or brings a device back online, does not clear the error count.

## powermt display options

**Purpose** Display PowerPath options.

**Syntax** `powermt display options`

**Description** `powermt display options` displays the following option settings for each storage system class:

- ◆ Whether VNX and CLARiiON user-assignable LUN names are displayed
- ◆ Whether path latency monitoring is enabled and the path latency threshold, if one is set
- ◆ Whether periodic autorestore and reactive autorestore is enabled
- ◆ Autostandby settings
- ◆ Performance data collection settings
- ◆ Status of the storage system class (managed, unmanaged, manage\_incomplete, or unmanage\_incomplete)

The following is sample output for `powermt display options`:

```

Default storage system class: all
Show CLARiiON LUN names:      true

Path Latency Monitor: Off

Path Latency Threshold: 0 milliseconds

Performance Monitor: enabled
Interval: 1 min(s)

Autostandby: Proximity          (prox): enabled
              IOs per failure (iopf): enabled
              iopf Reset period   : 7 d
              iopf limit          : 6000

Storage
System Class      Attributes
-----
Symmetrix         periodic autorestore = on
                  reactive autorestore = on
                  status = managed

CLARiiON          periodic autorestore = on
                  reactive autorestore = on
                  status = managed

Invista           periodic autorestore = on
                  reactive autorestore = on
                  status = managed

Hitachi           periodic autorestore = on
                  reactive autorestore = on
                  status = managed

HP xp             periodic autorestore = on
                  reactive autorestore = on
                  status = managed

```

Ess	periodic autorestore = on reactive autorestore = on status = managed
Generic	periodic autorestore = on reactive autorestore = on status = managed
VPLEX	periodic autorestore = on reactive autorestore = on status = managed
VNX	periodic autorestore = on reactive autorestore = on status = managed
NetApp	periodic autorestore = on reactive autorestore = on status = unmanaged
XtremIO	periodic autorestore = on reactive autorestore = on status = managed

## Platform-specific notes

PowerPath for Windows 5.7 SP2 and later supports NetApp and EMC XtremIO storage system class.

To change the settings, use the appropriate **powermt set** or **powermt load** commands to change the default storage system class.

VNX and CLARiiON user-assignable LUN names are displayed by default. To suppress the display of LUN names, define the environment variable **PP\_SHOW\_CLAR\_LUN\_NAMES** and set it to **false** or **FALSE**. (To re-enable the default behavior, either remove the variable or set it to **true** or **TRUE**.) An example using Bourne shell syntax to suppress the display of LUN names is:

```
PP_SHOW_CLAR_LUN_NAMES=false
export PP_SHOW_CLAR_LUN_NAMES
```

You can add a time stamp to the **powermt display every=<#seconds>** command output by defining the **PP\_DISPLAY\_TIME\_STAMP** environment variable and setting the time display format to **TIME\_VERBOSE** or **TIME\_SECONDS**. **TIME\_VERBOSE** displays complete details such as day, month, and year, time in hours, minutes, and seconds that increments until you exit the **powermt display** command. **TIME\_SECONDS** starts a counter from zero that increments until you exit the **powermt display** command. If the environment variable is not defined or if it is defined but not assigned a proper value the time stamp is not displayed.

## powermt display bus

**Purpose** Displays PowerPath bus-port specific information.

**Syntax** `powermt display bus [class={<class>|all}]  
[every=<#seconds>] [wide|width=<#columns>]`

**Description** **powermt display bus** displays all paths between the HBAs on a host and the storage-system ports that the HBAs can access. The output has one line per path. It displays information one line per bus. Alternatively, you can type **display paths** to display path information.

There is a table for each type of storage system. The **logical device count** above each table indicates the type of storage system to which that table applies. The summary display information does not list any arrays if there are no devices for an array type.

**Arguments** `class={symm|vnx|clariion|vplex|invista|netapp|ess|xtremio|hitachi|hpxp|generic|all}`  
Limits the command to the specified type of storage system. **all** specifies all storage-system types. The default is **all**. For descriptions of all the class types, see [“Command syntax: class argument specification” on page 24](#).

`every=<#seconds>`  
Integer in the range 1 to 86400 that specifies, in seconds, how often the display is updated. By default, the display is not updated.

The *#seconds* value is the minimum time between refreshes; the actual time is affected by the overall system load. On busy systems, display updates can be less frequent than specified.

You can add a time stamp to the **powermt display every=<#seconds>** command output by defining the **PP\_DISPLAY\_TIME\_STAMP** environment variable and setting the time display format to **TIME\_VERBOSE** or **TIME\_SECONDS**. **TIME\_VERBOSE** displays complete details such as day, month, and year, time in hours, minutes, and seconds that increments until you exit the **powermt display** command. **TIME\_SECONDS** starts a counter from zero that increments until you exit the **powermt display** command.

If the environment variable is not defined or if it is defined but not assigned a proper value the time stamp is not displayed.

`wide|width=<#columns>`

**Note:** The wide argument is not supported on Windows platforms.

**wide** sets the line length to the width of the window and **width** sets a range within which the **powermt display** output appears. The default is 80. The acceptable range is 80 to 160 columns. The output widens or narrows within the range depending on the width of the window. This value applies only to the specific output; it is not persistent. You can set a persistent value for the `width=<#columns>` option by using the environment variable **PP\_DISPLAY\_WIDTH**. This variable changes the width setting from 80 (the default) to the value you specify. If the environment variable value differs from the value you set with the command line, the command line value takes precedence.

If you set the **every=<#seconds>** option along with the **width=<#columns>** option, the **powermt display** output widens or narrows to accommodate changes you make to the window width.

**Example** Sample output of **powermt display bus** follows:

```
powermt display bus
Symmetrix logical device count=20
=====
----- Host Bus Adapters ----- Storage System ----- - I/O Paths -
### HW Path ID Interface Total Dead
=====
2304 pci@17/fibre-channel@2 000183500491 FA 1aA 20 0
2304 pci@17/fibre-channel@2 000183500491 FA 1bA 20 0
2305 pci@16/fibre-channel@1 000183500491 FA 1aA 20 0
2305 pci@16/fibre-channel@1 000183500491 FA 1bA 20 0

CLARiiON logical device count=21
=====
----- Host Bus Adapters ----- Storage System ----- - I/O Paths -
### HW Path ID Interface Total Dead
=====
2304 pci@17/fibre-channel@2 F80004600258 SP A1 21 0
2304 pci@17/fibre-channel@2 F80004600258 SP B0 21 0
2305 pci@16/fibre-channel@1 F80004600258 SP A0 21 0
2305 pci@16/fibre-channel@1 F80004600258 SP B1 21 0

Ess logical device count=10
=====
----- Host Bus Adapters ----- Storage System ----- - I/O Paths -
### HW Path ID Interface Total Dead
=====
2304 pci@17/fibre-channel@2 15748 84 10 0
2305 pci@16/fibre-channel@1 15748 24 10 0

Hitachi logical device count=9
=====
----- Host Bus Adapters ----- Storage System ----- - I/O Paths -
### HW Path ID Interface Total Dead
=====
2304 pci@17/fibre-channel@2 9CAD 1C 9 0
2304 pci@17/fibre-channel@2 9CAD 2C 9 0
2305 pci@16/fibre-channel@1 9CAD 1C 9 0
2305 pci@16/fibre-channel@1 9CAD 2C 9 0
```

## powermt display perf dev

**Purpose** Display storage device performance metrics for all logical devices.

**Syntax** `powermt display perf dev=<device>|all [continuous] [verbose] [xml]`

[nowait] **Description** **powermt display perf dev** displays storage device performance metrics for all logical devices. The output has one line per path. There is a table for each type of storage device. The storage device count above each table includes the type of storage device to which that table applies.

If there is no data for the device or if there is no device specified, the following error message is thrown: “Device(s) not found”

---

**Note:** A *path* is the ITL nexus (Initiator, Target, LUN) and represents the connection between the HBA, Storage Port and LUN.

---

If performance monitoring has just been turned on and continuous is specified, then the first output will wait until the first interval has completed, and subsequent output will also wait the interval time.

### Arguments

`dev=<device>|all`

Displays performance data for a specific storage device or all storage devices.

`continuous`

Displays the metrics over every interval at the end of each interval. If **continuous** is not selected, displays the single most recently calculated metrics for all storage devices.

`verbose`

Displays latency distribution buckets for each path, which is the count of I/Os in different latency ranges for reads and writes. The count for the most recent interval of monitoring is displayed. The **verbose** option adds Read and Write I/Os per second to the read and write buckets to each storage device metrics.

`xml`

Displays performance data in XML formatting after the normal output. To redirect the performance data to a text file, add **xml 2** <filename>.xml, and optionally, a full path name, to the **powermt display perf dev** command. For example, running **powermt display perf dev=emcpower22 xml 2> perf.xml** displays the performance data on the screen and creates a file named perf.xml that contains the data in XML formatting. If you redirect the XML metrics to a file, then the XML data does not display on the screen.

`nowait`

If this option is specified, **powermt** will not wait for the interval time to display the next updated PMI statistics. It returns immediately to report the remaining time interval for the next updated PMI data. Detailed description

The following output for the **powermt display perf dev** command shows performance metrics of a CLARiiON storage system.

```
# powermt display perf dev=all verbose
Timestamp = 22:25:54 UTC, 12 Mar 2012
Sample Interval = 300
Pseudo name=emcpower32a
```

```

state=alive; policy=SymmOpt; queued-I/Os=10
          KB<=4K      4K<KB<=8K    8K<KB<=128K  KB>128K      All
Read bytes/s      4.50M      3.25M      13.1M      -      20.8M
Write bytes/s     4.50M      3.26M      13.1M      -      20.8M
Total bytes/s                    41.6M

Read Avg Response ms    1.03      0.985      1.01      -      1.02
Write Avg Response ms   1.05      0.989      1.02      -      1.04
All Avg Response ms    1.03

Read I/Os/s          2.80K      396      796      -      3.99K
Write I/Os/s         2.79K      397      798      -      3.99K
Total I/Os/s                    7.98K

```

```

=====
----- Host ----- Latency ----- Retry Error
### HW Path          I/O Paths      Reads--Writes  delta  delta
=====
3073 pci@0/pci@0/pci@8/pci@0/pci@1/QLGC,qlc@0,1/fp@0,0 c5t50000972081A4D95d384s0
          Low (ms)  0.139  0.197  0  0
          High (ms)  34.5  34.7
          Lat <= 1ms (#) 76004 71588
          1ms < Lat <= 10ms (#) 73418 76847
          10ms < Lat <= 100ms (#) 25 35
          Lat > 100ms (#) 0 0

3073 pci@0/pci@0/pci@8/pci@0/pci@1/QLGC,qlc@0,1/fp@0,0 c5t50000972081A4D91d384s0
          Low (ms)  0.139  0.198  0  0
          High (ms)  28.6  38.8
          Lat <= 1ms (#) 73201 69408
          1ms < Lat <= 10ms (#) 75819 77195
          10ms < Lat <= 100ms (#) 18 20
          Lat > 100ms (#) 0 0

3072 pci@0/pci@0/pci@8/pci@0/pci@1/QLGC,qlc@0/fp@0,0 c4t50000972081A4D95d384s0
          Low (ms)  0.138  0.198  0  0
          High (ms)  34.6  34.6

```

[Table 16 on page 63](#) describes the fields in the output for the command.

**Table 16** Fields in the perf dev display (page 1 of 2)

Field	Value	Description
Timestamp	Hours:Minutes:Seconds	Period since when the last metrics were computed.
Sample Interval	Seconds	Interval period when the metrics are gathered.
Pseudo name	Alphanumeric string	The storage device pseudo name for each storage device from which the metrics are gathered. The metrics for each storage devices are listed separately.
State	alive dead	The state of the storage device from which the metrics are gathered. The state can be either <b>dead</b> or <b>alive</b> .
Policy	Adaptive (ad) CLAROpt (co) LeastBlocks (lb) Least I/O (li) Request (re) RoundRobin (rr) Stream I/O (si) SymmOpt (so)	Current load-balancing and failover policy for the device.
queued-I/Os	Non-negative integer	Number of queued I/Os to the storage device.

**Table 16** Fields in the perf dev display (page 2 of 2)

Field	Value	Description
Sz<=4K 4K<Sz<=8K 8K<Sz<=128K Sz>128K All	Alphanumeric string	The I/O size buckets for each storage device and for all storage devices put together. Where Sz (size) is in KB or MB.
Read bytes/s Write bytes/s Total bytes/s	Kilobytes/second Megabytes/second	The Read and Write per second for different I/O size buckets for each storage device. The total throughput per second of all reads and writes for all storage devices.
Read Avg Response (ms) Write Avg Response (ms) All Avg Response (ms)	Milliseconds	The average Read and Write response time in the Sample Interval for different I/O size buckets for each storage device. The total average response time for all Read and Writes in the Sample Interval for different I/O size buckets for each storage device.
Read IOs/s Write IOs/s Total IOs/s	Kilobytes/second	The read and write I/Os per second rate value for each consolidated buckets. The total I/Os per second for all reads and writes.
Host HW Path	Alphanumeric string	Platform-specific hardware descriptive name for the path.
Host I/O Path	Alphanumeric string	Platform-specific device name for the I/O path.
Latency		When the verbose option is specified, the Read and Write I/Os per second count for each device path is shown under each bins. The latency bins shown in the verbose output are a good place to verify that dead, standby (both manual and autostandby), and non-optimal paths are not handling application reads and writes under normal circumstances. Keep in mind, however, that even standby and non-optimal paths can have application reads and writes routed to them under extremely heavy loads, and such "overflow" ios would show up in the latency bins.
Latency Reads	Milliseconds	The high and low water marks for Read counts per each device path in the interval. In milliseconds.
Latency Writes	Milliseconds	The high and low water marks for Write counts per each device path in the interval. In milliseconds.
Retry delta	Non-negative integer	Number of retry attempts to restore a failed I/O through another path to the storage device during the interval.
Error delta	Non-negative integer	Number of alive to dead transitions during the interval.



## powermt display perf bus

**Purpose** Display path performance metrics for each bus.

**Syntax** `powermt display perf bus [continuous] [verbose] [xml] [nowait]`

**Description** **powermt display perf bus** displays path performance metrics for each buses. If performance monitoring has just been turned on and continuous is specified, then the first output will wait until the first interval has completed and the subsequent output will also wait the interval time. The summary display information does not list any arrays if there are no devices for an array type.

If there is no device specified, the following error message is thrown: “Device(s) not found”

---

**Note:** A *bus* is the IT nexus (Initiator, Target) and represent the connection between the HBA and the Storage Port.

---

### Arguments

`continuous`

Displays the metrics over every interval at the end of each interval. If **continuous** is not selected, displays the single most recently calculated metrics for all storage devices.

`verbose`

Displays latency distribution buckets for each path, which is the count of I/Os in different latency ranges for reads and writes. The count for the most recent interval of monitoring is displayed. The **verbose** option adds Read and Write I/Os per second to the read and write buckets to each storage device metrics.

`xml`

Displays performance data in XML formatting after the normal output. To redirect the performance data to a text file, add **xml 2> <filename>.xml**, and optionally, a full path name, to the **powermt display perf dev=emcpower22 xml 2> perf.xml** displays the performance data on the screen and creates a file named perf.xml that contains the data in XML formatting. If you redirect the XML metrics to a file, then the XML data does not display on the screen.

`nowait`

If this option is specified, **powermt** will not wait for the interval time to display the next updated PMI statistics and will return immediately reporting remaining time interval for the next updated PMI data.

**Detailed description** The following output for the **powermt display perf bus** command shows performance metrics of a Symmetrix storage system.

```
# powermt display perf bus verbose
Timestamp = 14:25:54 UTC, 13 Mar 2012
Sample Interval = 300

Symmetrix logical device count=418
=====
----- Host Bus Adapters ----- Storage System ----- Retry Error
### HW Path ID Interface delta delta
=====
3072 pci@0/pci@0/pci@8/pci@0/pci@1/QLGC,qlc@0/fp@0,0
      000192601683 FA 6gB 0 0
      Latency---Reads--Writes
      Low (ms) 0.138 0.196
      High (ms) 34.6 34.6
      Lat <= 1ms (#) 75876 71571
      1ms < Lat <= 10ms (#) 72599 77647
      10ms < Lat <= 100ms (#) 39 39
      Lat > 100ms (#) 0 0
3072 pci@0/pci@0/pci@8/pci@0/pci@1/QLGC,qlc@0/fp@0,0
      000192601683 FA 5gB 0 0
      Latency---Reads--Writes
      Low (ms) 0.138 0.197
      High (ms) 38.4 30.6
      Lat <= 1ms (#) 72033 69174
      1ms < Lat <= 10ms (#) 75792 78296
      10ms < Lat <= 100ms (#) 26 32
      Lat > 100ms (#) 0 0
3073 pci@0/pci@0/pci@8/pci@0/pci@1/QLGC,qlc@0,1/fp@0,0
      000192601683 FA 6gB 0 0
      Latency---Reads--Writes
      Low (ms) 0.139 0.192
      High (ms) 34.5 34.7
      Lat <= 1ms (#) 75561 71178
      1ms < Lat <= 10ms (#) 73703 76272
      10ms < Lat <= 100ms (#) 40 55
      Lat > 100ms (#) 0 0
```

[Table 17 on page 66](#) describes the fields in the output of the command.

**Table 17** Fields in the perf bus display (page 1 of 2)

Field	Value	Description
Timestamp	Hours:Minutes:Seconds	Period since when the last metrics are computed.
Sample Interval	Seconds	Interval period when the metrics are gathered.
Device class name count	Alphanumeric string	Storage device class name and the total number of devices accessed through the buses.
HW Path	Alphanumeric string	Platform-specific hardware descriptive name of every path for the bus.
ID	Non-negative integer	Platform-specific ID for the bus.
Latency		When the verbose option is specified, the Read and Write I/Os per second count for each bus is shown under each latency bins per interface.
Latency Reads	Milliseconds	The high and low water marks for Read counts per each interface for the bus in the interval. In milliseconds.

**Table 17** Fields in the perf bus display (page 2 of 2)

Field	Value	Description
Latency Writes	Milliseconds	The high and low water marks for Write counts per each interface for the bus in the interval. In milliseconds.
Retry delta	Non-negative integer	Number of retry attempts to restore a failed I/O through another path to the bus during the interval.
Error delta	Non-negative integer	Number of alive to dead transitions during the interval.

## powermt display port\_mode

**Purpose** Display storage-system port modes.

**Syntax** `powermt display port_mode [dev=<device>|all] [class=<class>|all]`

**Description** `powermt display port_mode` displays the mode of all ports on all storage systems configured on the host. The port mode is either enabled or disabled.

An enabled port is active and operates as it normally would within a PowerPath environment. A disabled port was disabled previously using the `powermt set port_disable` command. When a port is disabled, the devices configured through it are closed.

To enable or disable a storage system port, use the `powermt set port_disable` command.

The summary display information does not list any arrays if there are no devices for an array type.

### Arguments

`dev=<device>`

Specifies the device port whose information you want to view. The device you enter must be a native device name.

`class={symm|vnx|clariion|vplex|invista|netapp|ess|xtremio|hitachi|hpxp|generic|all}`

Limits the command to the specified type of storage system. **all** specifies all storage-system types. The default is **all**. For descriptions of all the class types, see [“Command syntax: class argument specification” on page 24](#).

**Note:** CLI parser is enhanced to process the `powermt` command regardless of its arguments and options’ location. For example, same output would be shown for both the `powermt display port_mode dev=4` and `powermt display dev=4 port_mode` commands.

### Detailed description

The following output for `powermt display port_mode` shows a Symmetrix, CLARiiON, and Hitachi storage system with both ports enabled.

```
#powermt display port_mode
Storage class = Symmetrix
=====
----- Storage System ----- -- I/O Paths -- Stats
      ID          Interface      Wt_Q      Total    Dead    Q-IOs  Mode
=====
      000192601683    FA  5gB        256        42        0        0  Enabled
      000192601683    FA  6gB        256        42        0        0  Enabled

Storage class = Clariion
=====
----- Storage System ----- -- I/O Paths -- Stats
      ID          Interface      Wt_Q      Total    Dead    Q-IOs  Mode
=====
      APM00053001551  SP  A0        256         8         0         0  Enabled
      APM00053001551  SP  B1        256         8         0         0  Enabled

Storage class = Hitachi
=====
----- Storage System ----- -- I/O Paths -- --- Stats ---
      ID          Interface      Wt_Q      Total    Dead    Q-IOs  Mode
=====
      10406          7R            256        16         0         0  Enabled
```

10406                      8R                      256                      16                      0                      0                      Enabled

The following output for **powermt display port\_mode dev=all** shows all the port modes of the Symmetrix device configured on the host.

```
# powermt display port_mode dev=all
Pseudo name=emcpowera
Symmetrix ID=000190300848
Logical device ID=05A6
state=alive; policy=SymmOpt; queued-IOS=0
=====
----- Host ----- - Stor - -- I/O Path -- -- Stats ---
### HW Path          I/O Paths   Interf.  Mode   State  Q-IOS Errors
=====
  1 qla2xxx           sdbd       FA  2cB  active  alive    0    0
  1 qla2xxx           sdbp       FA  1cB  active  alive    0    0
  2 qla2xxx           sdcz       FA  2cB  active  alive    0    0
  2 qla2xxx           sddk       FA  1cB  active  alive    0    0

Pseudo name=emcpowerb
Symmetrix ID=000190300848
Logical device ID=05A7
state=alive; policy=SymmOpt; queued-IOS=0
=====
----- Host ----- - Stor - -- I/O Path -- -- Stats ---
### HW Path          I/O Paths   Interf.  Mode   State  Q-IOS Errors
=====
  1 qla2xxx           sdbf       FA  2cB  active  alive    0    0
  1 qla2xxx           sdbp       FA  1cB  active  alive    0    0
  2 qla2xxx           sdda       FA  2cB  active  alive    0    0
  2 qla2xxx           sddl       FA  1cB  active  alive    0    0

Pseudo name=emcpowerc
Symmetrix ID=000190300848
Logical device ID=05A3
state=alive; policy=SymmOpt; queued-IOS=0
=====
----- Host ----- - Stor - -- I/O Path -- -- Stats ---
### HW Path          I/O Paths   Interf.  Mode   State  Q-IOS Errors
=====
  1 qla2xxx           sdba       FA  2cB  active  alive    0    0
  1 qla2xxx           sdbm       FA  1cB  active  alive    0    0
  2 qla2xxx           sdcw       FA  2cB  active  alive    0    0
  2 qla2xxx           sddh       FA  1cB  active  alive    0    0
```

Table 18 on page 70 explains the fields in the port\_mode display.

**Table 18** Fields in the port\_mode display

Field	Value	Description
Storage System ID	Hexadecimal value or alphanumeric string	Identification number for the storage system on which the logical device is located.
Storage System Interface	Alphanumeric string	Storage-system interface. For Symmetrix systems, this has three parts: <ul style="list-style-type: none"> <li>• Interface type: Fibre Channel (FA) or SCSI (SA)</li> <li>• Interface address: integer in the range 1 to 16</li> <li>• Interface port: [abcd][AB]</li> </ul> For VNX and CLARiiON systems, this has one part: <ul style="list-style-type: none"> <li>• Interface port: SP [A-B][0-3]</li> </ul> For HP StorageWorks EVA this has one part: <ul style="list-style-type: none"> <li>• Interface port: [A-B][1-2]</li> </ul> A, B denotes the controller. 1, 2 denotes the port. For Hitachi Lightning, Hitachi TagmaStore and HP XP systems, this is the port name on the array. For IBM ESS systems, this is a two-digit hexadecimal port ID that can be decoded to a bay/adaptor/port on the ESS system. (IBM ESS command-line interface documentation describes how to decode the port ID.) For EMC Invista and VPLEX systems, a two-digit hexadecimal number that can be decoded to a virtual target on the Invista system.
Storage System Wt_Q	Non-negative integer	Maximum number of write I/O requests that will be serviced before the I/O queue checks for any outstanding read I/O.
I/O Paths Total	Non-negative integer	Total number of paths that originate from this HBA. The maximum number of logical devices supported by one HBA is platform specific. The total number of paths may exceed the number of logical devices in a complex SAN topology, due to zoning.
I/O Paths Dead	Integer in the range 0 – <i>I/O Paths Total</i>	Total number of paths originating from this HBA that are dead (not usable).
Stats Q-IOs	Non-negative integer	Total number of I/O operations under way to this path. This is the total number of I/O requests to this device that have not completed. The sum of in-progress I/Os for all paths should equal the number of in-progress I/Os for the PowerPath device.
Mode	Enabled Disabled	The mode of the storage-system port: <ul style="list-style-type: none"> <li>• <b>Enabled</b> means the port is active and operating as it normally would in a PowerPath environment.</li> <li>• <b>Disabled</b> means the port has been disabled using the <b>powermt set port_disable</b> command.</li> </ul>

## powermt\_display\_dead

**Purpose** Lists the volumes which have dead paths and also the total number of dead paths.

**Syntax** `powermt_display_dead`

**Description** This script lists all the volumes which have dead paths in them. It also lists the total number of dead paths on a host. If there are no dead paths on a host, then it will just display the Total dead paths as 0.

**Platform-specific notes** This command is supported only on Linux.

## powermt display ports

**Purpose** Monitor ports.

**Syntax** `powermt display ports`  
`[dev=<device>|all] [class=<class>|all]`  
`[every=<#seconds>] [wide|width=<#columns>]`

**Description** `powermt display ports` displays attached array port specific information.

The `powermt display ports` command includes the **Array failover mode** field, which displays the failover mode setting for VNX and CLARiiON arrays. This setting determines the failover mode for a server's HBA ports, and is configured through the VNX and CLARiiON Navisphere software. Failover mode values are described in [Table 11 on page 40](#).

---

**Note:** After changing devices to ALUA mode, rebooting the host is required for PowerPath to recognize the change.

---

### Arguments

`ports`

Displays attached port specific information, one line per port.

`dev=<device>|all`

Displays information about the specified device. `all` specifies all devices.

`class={symm|vnx|clariion|vplex|invista|netapp|ess|xtremio|hitachi|hpxp|generic|all}`

Limits the command to the specified type of storage system. `all` specifies all storage-system types. The default is `all`. For descriptions of all the class types, see [“Command syntax: class argument specification” on page 24](#).

If `class` is specified along with `dev`, the command is applied to the devices that meet all specified constraints.

`every=<#seconds>`

Integer in the range 1 to 86400 that specifies, in seconds, how often the display is updated. By default, the display is not updated.

The `seconds` value is the *minimum* time between refreshes; the actual time is affected by the overall system load. On busy systems, display updates can be less frequent than specified.

You can add a time stamp to the `powermt display every=<#seconds>` command output by defining the `PP_DISPLAY_TIME_STAMP` environment variable and setting the time display format to `TIME_VERBOSE` or `TIME_SECONDS`. `TIME_VERBOSE` displays complete details such as day, month, and year, time in hours, minutes, and seconds that increments until you exit the `powermt display` command. `TIME_SECONDS` starts a counter from zero that increments until you exit the `powermt display` command.

If the environment variable is not defined or if it is defined but not assigned a proper value the time stamp is not displayed.



`wide | width=<#columns>`

**Note:** The `wide` argument is not supported on Windows platforms.

**wide** sets the line length to the width of the window and **width** sets a range within which the **powermt display** output appears. The default is 80. The acceptable range is 80 to 160 columns. The output widens or narrows within the range depending on the width of the window. This value applies only to the specific output; it is not persistent. You can set a persistent value for the `width=<#columns>` option by using the environment variable `PP_DISPLAY_WIDTH`. This variable changes the width setting from 80 (the default) to the value you specify. If the environment variable value differs from the value you set with the command line, the command line value takes precedence.

If you set the `every=<#seconds>` option along with the `width=<#columns>` option, the **powermt display** output widens or narrows to accommodate changes you make to the window width.

**Description** The **powermt display ports** command displays the state of the storage-system ports that are visible to the HBAs under PowerPath control. The output has one line per port. Sample output of **powermt display ports** is shown as follows:

```
powermt display ports
Storage class = Symmetrix
=====
----- Storage System -----          -- I/O Paths --          --- Stats ---
      ID              Interface           Wt_Q           Total      Dead      Q-I/Os  Errors
=====
      000190300848      FA 2cB             256            20         0         0         0
      000190300848      FA 1cB             256            20         0         0         0

Storage class = CLARiiON
=====
----- Storage System -----          -- I/O Paths --          --- Stats ---
      ID              Interface           Wt_Q           Total      Dead      Q-I/Os  Errors
=====
      FNM00083800185    SP A0              256            11         0         0         0
      FNM00083800185    SP B0              256            11         0         0         0
      FNM00083800185    SP A1              256            11         0         0         0
      FNM00083800185    SP B1              256            11         0         0         0

Storage class = VNX
=====
----- Storage System -----          -- I/O Paths --          --- Stats ---
      ID              Interface           Wt_Q           Total      Dead      Q-I/Os  Errors
=====
      APM00120701240    SP B8              256            12         0         0         0
      APM00120701240    SP A8              256            12         0         0         0
```

## Specifying the dev option with ports

If the **dev** option is specified with **ports**, the following displays:

```
# powermt display ports dev=emcpowera
Pseudo name=emcpowera
Symmetrix ID=000195700363
Logical device ID=7743
state=alive; policy=SymmOpt; queued-I/Os=0; write_throttle=off
=====
----- Host ----- - Storage Sys - - Path - -- Stats --
### HW Path          I/O Paths   Interf.   Wt_Q   Status   Q-I/Os Errors
=====
 4 lpfc              sdat      FA  8gB   256    a,v     0       0
 4 lpfc              sdbg      FA  9gB   256    a,v     0       0
 5 lpfc              sdcz      FA  8gB   256    a,v     0       0
 5 lpfc              sddm      FA  9gB   256    a,v     0       0
```

All fields are as described previously for the HBA and devices displays (see the Display command), except one new column, **Path Status**. **Path Status** is a combination of the **I/O Paths Mode** and **I/O Paths State** columns of the normal devices display. **Path Status** is a two-part field, with valid values as shown in [Table 19 on page 74](#).

**Table 19** Path status

Value	Meaning	Description
a s u	active standby unlicensed	Refer to the description of <b>Path Mode</b> in <a href="#">Table 11 on page 40</a> . Standby in this situation includes <b>asb:prox</b> and <b>asb:iopf</b> standby created automatically by PowerPath algorithms.
v d	alive dead	Refer to the description of <b>Path State</b> in <a href="#">Table 11 on page 40</a> .

## powermt display unmanaged

**Purpose** Display information about unmanaged devices.

**Syntax** `powermt display unmanaged`

**Description** `powermt display unmanaged` displays information about logical devices that have been excluded from PowerPath management through the `powermt unmanage` command (refer to “[powermt unmanage](#)” on page 110). You can use the native device as the `dev` argument to the `powermt manage dev` command (refer to “[powermt manage](#)” on page 80).

Following is sample output of `powermt display unmanaged`:

```
Hitachi ID=40109
Logical device ID=HITACHI R4009CAD0144
Vendor ID=HITACHI ; Product ID=OPEN-9
=====
----- Host -----
HW Path                               I/O Paths
=====
pci@1f/fibre-channel@4/sd@1           c2t1d10s0
pci@1f/fibre-channel@1/sd@1           c3t1d10s0

Hitachi ID=40109
Logical device ID=HITACHI R4009CAD0150
Vendor ID=HITACHI ; Product ID=OPEN-9
=====
----- Host -----
HW Path                               I/O Paths
=====
pci@1f/fibre-channel@4/sd@1           c2t1d11s0
pci@1f/fibre-channel@1/sd@1           c3t1d11s0
```

**Note:** Devices unmanaged with `powermt unmanage class = <class>` will not be listed.

## powermt enable hba

**Purpose** Enable an HBA port.

**Syntax** `powermt enable hba=<hba#>`

**Description** **powermt enable** enables an HBA port that was previously disabled with the **powermt disable** command. Running **powermt enable** resumes I/O to the specified HBA port and its devices.

This command is supported with the following platform configurations:

- ◆ AIX and Solaris platforms—Fibre Channel only.
- ◆ HP-UX platforms—iSCSI and Fibre Channel.
- ◆ Linux platforms—iSCSI (with HBAs only) and Fibre Channel.

To display whether an HBA is enabled or disabled, refer to [“powermt display hba\\_mode” on page 47](#).

Note the following limitation: This command fails if the PowerPath load balancing and failover policy is **bf** or **nr**.

**Arguments** `hba=<hba#>`  
Enables the specified HBA port, which resumes I/O to the HBA and its devices.

**Platform-specific notes** On a Solaris 10 host, you may see the following error message when you try to enable a disabled port:

```
Bad dev value <device>, or not under PowerPath control.
```

If this occurs, run the Solaris commands shown in the following procedure, and then enable the HBA:

1. Type **cfgadm -vl** to learn the controller number for the cards on each host.
2. For the appropriate card, type **cfgadm -vc configure cX**  
where *X* is the controller number for the card  
Optionally, type **devfsadm**.
3. Type **powermt enable hba=<hba#>**  
where *hba#* is the PowerPath HBA you want to enable.

## powermt help

**Purpose** Display the usage of all **powermt** commands.

**Syntax** `powermt help [<command>]`

**Description** **powermt help** displays a summary of the command syntax. The usage syntax is also displayed when invalid commands or options are entered.

**Arguments** `<command>`  
If a specific command is provided, displays the usage only of the command specified.

**Example** `# powermt help display`

`display - Display multipathing information.`

`Usage:`

`powermt <command>`

`[class={all|symm|vnx|clariion|vplex|invista|xtremio|netapp|hitachi|hpx  
p|ess|generic}]`

`powermt display [dev=<device>|all] [class=<class>|all]  
[every=<#seconds>] [wide|width=<#col>]`

`powermt display alua dev=<device>|all [class=<class>|all]  
[every=<#seconds>] [wide|width=<#col>]`

`powermt display hba_mode [class=<class>|all]`

`powermt display latency [dev=<device>|all] [class=<class>|all]  
[every=<#seconds>] [wide|width=<#col>]`

`powermt display nonvirtual dev=<device>|all  
[every=<#seconds>] [class=<class>|all] [wide|width=<#col>]`

`powermt display options`

`powermt display bus [class=<class>|all]  
[every=<#seconds>] [wide|width=<#col>]`

`powermt display perf dev=<device>|all [continuous] [verbose] [xml]  
[nowait]`

`powermt display perf bus [continuous] [verbose] [xml] [nowait]`

`powermt display port_mode [dev=<device>|all] [class=<class>|all]`

`powermt display ports [dev=<device>|all] [class=<class>|all]  
[every=<#seconds>] [wide|width=<#col>]`

`powermt display unmanaged`

## powermt load

**Purpose** Load a custom configuration.

**Syntax** `powermt load [file=<filename>]`

**Description** **powermt load** applies to the current configuration the following settings from a saved configuration file. These include PowerPath device policies and priorities, write throttle setting and write throttle queue depth, periodic autorestore setting, and path modes.

The configuration file is saved when **powermt save** is run manually and, on some platforms, at boot time. See “[powermt save](#)” on page 86 for additional information. (On Windows, PowerPath Administrator prompts to save the file after configuration changes are made in the GUI.)

Configuration file saved only when <b>powermt save</b> is run manually	Configuration file also saved when host is booted
Windows AIX	HP-UX Solaris Linux

If a configuration is saved with **powermt save** under the default filename, the configuration loads automatically at boot time. The exception is AIX, which always loads default settings at boot time. If the default filename is not used, the configuration does not load automatically at boot time; in this case, you must load the configuration manually with **powermt load**.

On Solaris, HP-UX, and Windows, **powermt load** also causes devices to be configured:

- ◆ Paths added since the last save remain as is.
- ◆ If the configuration file has information about dead paths, they are added in a failed state, allowing them to be restored later.
- ◆ If the file has information about paths that were removed since the last save, they are added.

If the PowerPath license is invalid, **powermt load** issues a warning and continues its processing. Use **powermt check\_registration** to determine the problem with the license.

### Arguments

`file=<filename>`

This is an optional argument. The system will load the default configuration file if one is not specified.

---

**Note:** The default configuration file name and location must be used during boot. For the default filenames on each platform, see “[Default Configuration filename and pathname](#)” on page 86.

---

If you specify a filename and it is not in your current directory, then you must include the full file path as well as the filename in the file specification. For example: `powermt load file=/etc/emc/powermt_custom.xml`.

**Platform-specific  
notes**

On Solaris, **powermt load** behavior is undefined when invoked after the host has booted and while paths are in a state other than alive. Do not use **powermt load** to restore dead paths online.

On Linux, after you run **powermt load**, you must run **powermt config** to configure pseudo devices; otherwise, only native devices are configured.

## powermt manage

**Purpose** Place devices or storage system classes under PowerPath management.

**Syntax** `powermt manage dev=<path>|<device> | class=<class>`

**Description** **powermt manage** places a specified logical device or storage system class under PowerPath management.

By default, all storage system classes are under PowerPath management. You can use **powermt unmanage** to exclude a storage system class from PowerPath management. For more information, refer to “[powermt unmanage](#)” on page 110. Use **powermt manage class** to return that storage system class to PowerPath control.

A **powermt manage class** operation is legal only if the class’s status is **unmanaged** or **manage\_incomplete**. If the class status is **unmanage\_incomplete**, the unmanage operation must be completed—or the host rebooted— before **powermt manage** succeeds. The **powermt display options** command displays the status of a class.

Likewise, by default, when a storage system class is under PowerPath management, all logical devices of that class are under PowerPath management. You can use **powermt unmanage** to exclude an individual device from PowerPath management. Use **powermt manage** to return that device to PowerPath control. **powermt manage** returns a device to PowerPath control only if its storage system class is currently managed by PowerPath.

### Arguments

`dev=<path>|<device>`

Specifies the pathname of a logical device to be managed by PowerPath.

---

**Note:** **powermt manage** cannot accept a pseudo device as an argument, since pseudo devices are not assigned to unmanaged devices/classes. Note, however, that this restriction does not apply to the **powermt unmanage** command (refer to “[powermt unmanage](#)” on page 110), as pseudo device names *are* assigned to managed devices/classes.

---

`class={symm|vnx|clariion|vplex|invista|netapp|ess|xtremio|hitachi|hpxp|generic}`

Limits the command to the specified type of storage system. For descriptions of all the class types, see “[Command syntax: class argument specification](#)” on page 24.



## powermt release

**Purpose** Release pseudo device names for later reuse.

**Syntax** `powermt release`

**Description** **powermt release** releases PowerPath pseudo device (emcpower) names for later reuse. Run **powermt release** when, after numerous reconfigurations (adding and removing LUNs), the Linux host and PowerPath can discover all devices, but PowerPath has run out of emcpower names to use for devices.

## powermt remove

**Purpose** Remove paths from PowerPath management.

**Syntax** `powermt remove hba=<hba#>|all | dev=<path>|<device>|all  
[class=<class>|all] [force]`

**Description** **powermt remove** deletes the specified path (or paths) from PowerPath's list of configured paths. It does not delete the logical device that the paths refer to. As long as the logical device remains visible on the system, the logical device can be reconfigured via **powermt config**.

**⚠ CAUTION**

**Do not remove the last path to a logical device unless you plan to remove the logical device entirely, as data access will be interrupted.**

On all platforms, **powermt remove** does not remove a path if that native path is open—for example, if a file system is mounted on the native path device or an application has opened the native path device. Instead, **powermt remove** removes all specified paths that are not currently open, issues an error message, and exits.

If a PowerPath device is open, **powermt remove** does not remove that device or the last path to that device.

In addition, **powermt remove** operates on paths, so when **powermt remove** is executed with an HBA or PowerPath device argument, **powermt remove** iterates through all the associated paths and attempts to remove them. The HBA or PowerPath device is removed only if all the associated paths are successfully removed.

To permanently remove a path from the PowerPath configuration:

1. Run **powermt check** or **powermt remove**.
2. Physically remove the path.
3. Remove the operating system objects associated with the path or device.
4. Run **powermt save**.

If the PowerPath license is invalid, **powermt remove** issues a warning and continues its processing. Use **powermt check\_registration** to determine the problem with the license.

### Arguments

`force`

Suppresses the confirmation prompt and silently removes all specified paths, including the last path to a logical device.

`hba=<hba#>|all`

Limits removal to paths from the specified HBA. <hba#> is a number in the **Host Bus Adapters ###** column of **powermt display dev** output. **all** specifies all HBAs under PowerPath control. The default is **all**.

(**adapter** can be used instead of **hba**. **adapter** is included only for compatibility with prior PowerPath releases. Use **hba** instead.)

`dev=<path>|<device>|all`

Limits removal to the specified path, or all paths to the specified device. **all** specifies all paths to all devices. The default is **all**.

`class={symm|vnx|clariion|vplex|invista|netapp|ess|xtremio|hitachi|hpxp|generic|all}`

Limits the command to the specified type of storage system. **all** specifies all storage-system types. The default is **all**. For descriptions of all the class types, see [“Command syntax: class argument specification” on page 24](#).

### Platform-specific notes

On all platforms except Solaris, if you run **powermt remove** on a live path with active I/O, PowerPath returns the message “Cannot remove alive device <device\_name>” and exits. On Windows, running **powermt remove** will remove only dead paths, not alive paths. PowerPath returns the message “Cannot remove device that is in use: <device\_name>” is displayed even without active I/O on the alive path.

On Linux, even after all I/O has failed over from a dead path, **powermt remove** cannot remove devices until I/O is finished and devices are closed.

On Linux, devices removed using **powermt remove** will be automatically reconfigured when I/O is sent to a native device. In order to remove a device from PowerPath control, use **powermt unmanage dev=<pseudo>** to prevent it from being automatically reconfigured when IO is sent to one of its native device paths.

## powermt restore

**Purpose** Restore paths.

**Syntax** `powermt restore [hba=<hba#>|all] [dev=<path>|<device>|all]  
[class=<class>|all]`

**Description** **powermt restore** tests and restores specified paths. It issues test I/Os and responds to the test results as follows:

- ◆ If a *live* path passes the test, **powermt restore** does nothing.
- ◆ If a *dead* path passes the test, **powermt restore** marks it alive; PowerPath can use it. In addition, other dead paths on the same HBA or storage system port may be tested.
- ◆ If a *live* path fails the test, **powermt restore** marks it dead and prints a warning (every time the path fails the test). In addition, other paths that share the same HBA and port may be marked dead, and other paths that share only the HBA or only the port may be tested.
- ◆ If a *dead* path fails the test, **powermt restore** prints a warning (every time the path fails the test).

There may be a delay in accessing a recovered path. To avoid this delay, run **powermt restore** after a path is physically restored.

In addition to testing and restoring paths, **powermt restore** attempts to resurrect dead storage devices. A storage device may be marked dead if write errors occur that could jeopardize the integrity of the data structures, and if subsequent writes could aggravate the problem.

When applied to VNX and CLARiiON storage systems, **powermt restore** also relocates LUNs to their default storage processor.

If the PowerPath license is invalid, **powermt restore** issues a warning and continues its processing. To determine the problem with the PowerPath license, use **powermt check\_registration**.

### Arguments

`hba=<hba#>|all`

Limits restoration to paths from the specified HBA. *hba#* is a number in the **Host Bus Adapters ###** column of the **powermt display dev** output. **all** specifies all HBAs under PowerPath control. The default is **all**.

(**adapter** can be used instead of **hba**.)

**adapter** is included only for compatibility with prior PowerPath releases. Use **hba** instead.)

`dev=<path>|<device>|all`

Limits restoration to the specified *path*, or all paths to the specified *device*. **all** specifies all paths to all devices. The default is **all**.

```
class={symm|vnx|clariion|vplex|invista|netapp|ess|xtremio|hitachi|hpxp  
|generic|all}
```

Limits the command to the specified type of storage system. **all** specifies all storage-system types. The default is **all**. For descriptions of all the class types, see [“Command syntax: class argument specification” on page 24](#). If class is specified along with hba or dev, the command is applied to paths that meet all specified constraints.

## powermt save

**Purpose** Save a custom configuration.

**Syntax** `powermt save [file=<filename>]`

**Description** By default, **powermt save** saves the PowerPath settings in the PowerPath configuration file. These settings are important PowerPath driver information required for the PowerPath state persistence across host reboots.

For each configured logical device, **powermt save** records its serial numbers, policy, write throttle setting, write throttle queue depth, and pseudo device name (for platforms that support pseudo devices). For each configured path (alive or dead), **powermt save** records its mode. For each storage-system class, **powermt save** records its periodic autorestore setting.

For performance tuning purposes, you can use **powermt save** to save different configurations under different filenames.

A saved configuration can be reloaded with **powermt load**.

**Argument** `file=<filename>`  
This is an optional argument to override the default filename and location. If you do not specify this argument, PowerPath loads the default configuration file from the default location. If you specify a filename and it is not in your current directory, then you should include the full file path as well as the name of the file in which to save the configuration. PowerPath creates the file if it does not exist.

---

**Note:** The default filename and location must be using during a PowerPath boot.

---

**Table 20** Default Configuration filename and pathname

Platform	Default filename and pathname
Windows	The current default location of powermt custom file is: <PowerPath_Installation_Directory>\PowerPath\powermt.custom.xml  For example, if the PowerPath installation directory is default, then this path will be: C:\Program Files\EMC\PowerPath\powermt.custom.xml
AIX	/etc/powermt_custom.xml
HP-UX, Linux Solaris	/etc/powermt_custom.xml /etc/powermt.custom (a read only file)

If a configuration is saved with **powermt save** under the default filename, the configuration loads automatically at boot time. The exception is AIX, which always loads default settings at boot time. If the default filename is not used, the configuration does not load automatically at boot time; in this case, you must load the configuration manually with **powermt load**.

---

**Note:** If a configuration is saved under the default filename, any previously existing configuration file is overwritten without notice.

---

Run **powermt save** to update your PowerPath configuration whenever you make changes to hardware on the host. If you do not capture these changes in the **powermt** configuration file and then need to restore the system, the file that is restored will be incorrect.



**Do not manually modify the configuration file.**

## Platform-specific notes

On Windows, the file inherits the access permissions of the directory where it is saved. Thus, on Windows, if you save your configuration file to a directory other than the default directory, the file could inherit broader access permission than what was intended.

On Windows, without persistent binding, the SCSI ID of a LUN may change, leading to ghost paths appearing during a host reboot with a saved custom file configuration. This may cause powermt to display misleading information. You may also see dead paths in case these paths have different SCSI IDs. In such cases, you will also see more paths per LU than what was actually configured in your SAN environment.

On UNIX and Linux, the configuration file is saved with read-only access for root and no access for others.

## powermt set autostandby

**Purpose** Enables or disables autostandby. Autostandby is enabled by default.

**Syntax** `powermt set autostandby={on|off|reinitialize} [trigger={prox|iopf}]`

**Description** Autostandby specifies a path that has been placed into standby by PowerPath automatically either using the proximity-based autostandby algorithm (**asb:prox**) or the IOsPerFailure-based autostandby algorithm (**asb:iopf**). Autostandby enables you to automatically avoid paths that can lead to performance issues.

- ◆ Proximity-based autostandby (**asb:prox**) applies to path modes for multi-frame storage appliances, such as VPLEX Metro configurations, where a host is connected to and manages a distributed volume across multiple VPLEX clusters and frames within a given VPLEX system. PowerPath detects the most optimal paths (local paths) and places the remote paths in autostandby mode.

In mixed environments (VPLEX clusters with paths configured to cross-connected distributed volumes as well as paths configured to non-distributed volumes), the paths to the non-distributed volume inherit the cluster mode allocation that the proximity algorithm applies to its VPLEX cluster. For example, the PowerPath proximity algorithm designates a VPLEX cluster as local and preferred for a cross-connected, distributed volume on the host. After the distributed volume is designated as local, all paths to that cluster will be set as active. In addition, even if there is a non-distributed volume in the cluster, paths to that non-distributed volume are also set as active. This setting is inherited from the global setting set on the VPLEX cluster.

The default for the **prox** trigger is **on**. While autostandby is enabled or if autostandby is re-initialized with the **prox** trigger, PowerPath performs a test on all paths, sets the appropriate paths to active or autostandby, and determines the preferred paths to a VPLEX distributed volume.

- ◆ IOsPerFailure-based autostandby (**asb:iopf**) applies to paths with intermittent I/O failures (also known as *flakypaths*). When the average number of I/Os per path error falls between certain limits, the path changes to autostandby mode. Set the **iopflimit** using the `powermt set autostandby iopflimit=<value>` command. Additionally, you can set a time period of when the path changes back to active mode. Set the time period using the `powermt set autostandby agingperiod=<#days>` command.

The default for the **iopf** trigger is **on**. If autostandby is re-initialized with the **iopf** trigger, PowerPath begins keeping track of I/O failures for each path. If the running average of I/Os per failure for a path falls below the global iopf limit, PowerPath places that path in autostandby. A higher limit makes the trigger more sensitive to errors and conversely a lower limit makes it more tolerant of errors.

This command turns autostandby on or off globally within PowerPath and does not pertain to a specific storage array or volume. When you turn autostandby **off**, and then run `powermt save`, then this setting persists across reboot. The autostandby modes for each path do not persist across reboot. For example, if path set A is in autostandby mode and path set B is in active mode, then on reboot path set A can change to active mode and path set B can change to autostandby mode. On each reboot (assuming the global autostandby setting is on), the system performs latency measurements and re-runs the algorithms for selecting which path or path set to put in autostandby mode and which path or path set to put in active mode.



If a path is in standby mode, then PowerPath cannot move that path into autostandby mode.

---

**Note:** For scenarios in which all paths to a volume are in **asb:prox** and/or **asb:iopf** mode, PowerPath selects paths for I/O as if all paths are active.

---

You can override autostandby for a path by running **powermt set mode={standby|active} dev=<path> force**. If the path is already in the requested mode, then the command is ignored.

## Arguments

`autostandby=on`

Enables autostandby. There is no output. If you do not select a trigger, both **prox** and **iopf** are enabled. For example, **powermt set autostandby=on** globally turns on autostandby with both the proximity-based and the IOsPerFailure-based autostandby features. The default is **on** for both triggers.

`autostandby=off`

Disables autostandby. There is no output. PowerPath clears the global setting for enabling autostandby (or the selected type of autostandby), and persists that setting if you subsequently run the **powermt save** command. PowerPath also changes all paths in autostandby (or the selected type of autostandby) to active mode. If you do not select a trigger, then both **prox** and **iopf** are enabled. For example, **powermt set autostandby=off** globally turns off autostandby with both the proximity-based and the IOsPerFailure-based autostandby features.

`autostandby=reinitialize`

---

**Note:** Applies to the proximity trigger only. Re-initialize has no effect on an enabled **iopf** trigger.

---

If the proximity trigger is enabled, this option recalibrates the trigger and determines anew the preferred paths to VPLEX distributed volumes — this can be useful after a cluster has been physically moved, for example. There is no output. If autostandby is not enabled, then an error message displays. If you do not select a trigger, then **prox** is enabled. For example, **powermt set autostandby=reinitialize** re-initializes autostandby with the proximity-based autostandby feature. The minimum latencies for each cluster and the mode changes for each path are logged in audit messages.

`trigger=prox`

Selects only the autostandby proximity autostandby feature. This determines and selects the preferred paths to a VPLEX distributed volume and places the nonpreferred paths into autostandby mode (**asb:prox**).

---

**Note:** For scenarios in which all paths to a volume are in **asb:prox** mode, review your configuration and ensure that it is set up as intended.

---

`trigger=iopf`

Selects only the IOsPerFailure-based autostandby feature. This determines whether or not a path that has intermittent I/O failures should be placed into autostandby mode (**asb:iopf**).

## Example

The following command turns autostandby on and with the proximity autostandby feature:

```
powermt set autostandby=on trigger=prox
```

## powermt set autostandby agingperiod

**Purpose** Sets the length of time a path stays in autostandby due to an **iopf** (IOsPerFailure) trigger before returns back to active.

**Syntax** `powermt set autostandby agingperiod=<#days>`

**Description** Sets the length of time a path stays in autostandby due to an **iopf** (IOsPerFailure) trigger before it returns back to active. When a path is placed into autostandby mode through the IOsPerFailure trigger, PowerPath adds the aging amount to the current time. The aging period timeout allows PowerPath to return the paths to full use without user intervention. Error monitoring is resumed when the paths are returned to an active state. Each path ages separately. This may result in each path having different times at which they return to active mode.

If the aging period is modified, the modified aging period will be reflected on all paths even if the paths are already in **asb:iopf** mode.

This command sets the aging period for autostandby iopf globally within PowerPath and does not pertain to a specific storage array or volume. When you set aging period and then run **powermt save**, this setting persists across reboot.

**Arguments** `agingperiod=<#days>`  
 Specifies the number of days after which a path in autostandby mode due to intermittent path failure automatically resets to active mode. The default number of days is **7**. The range of acceptable values is 0 to 24854 days. There is no output.  
 Specifying **0** turns off the periodic reset of a path in autostandby due to intermittent path failure. These paths in autostandby (**asb:iopf**) never return to active mode without manual intervention.

**Example** The following command resets *flaky* paths (**asb:iopf** paths) to active after 90 days:

```
powermt set autostandby agingperiod=90
```

## powermt set autostandby iopflimit

**Purpose** Specifies the average number of I/Os between failures to automatically change that path to autostandby mode.

**Syntax** `powermt set autostandby iopflimit=<value>`

**Description** Sets the average number of I/Os needed between failures to automatically change that path to autostandby mode.

This command sets autostandby iopflimit globally within PowerPath and does not pertain to a specific storage array or volume. When you set the autostandby iopflimit and then run **powermt save**, the setting persists across reboot.

**Arguments** `iopflimit=<value>`  
Sets the average number of I/Os needed for a path error to automatically change that path into autostandby mode. The default is **6000** I/Os per I/O failure. The range of acceptable values is 2 to 100,000,000. This setting means that PowerPath will mark a path as having intermittent I/O errors if there are 6000 or fewer I/Os for every I/O error on the path. If there are greater than 6000 I/Os per I/O error on a path, PowerPath will not mark the path as having intermittent I/O errors. A higher limit makes the trigger more sensitive to errors and conversely a lower limit makes it more tolerant of errors.

## powermt set mode

**Purpose** Set the path mode.

**Syntax** `powermt set mode={active|standby}  
[hba=<hba#>|all] [dev=<path>|<device>|all]  
[class=<class>|all] [force]`

**Description** **powermt set mode** sets paths to active or standby mode. For most applications, the best performance is achieved by designating all paths to a PowerPath device as active. By using this command selectively, however, the path usage of devices can be controlled. This is helpful if you do not want I/O for one device to affect the performance of another device.

Paths in autostandby mode can be overridden, and moved to active or standby mode with this command. The **force** option is needed to override an autostandby setting. If the path is already in the requested mode, no change or error occurs. If the path is in autostandby and the **force** option is not specified, then an error is returned.

### Arguments

`mode=active`

User I/O is delivered to a path. If you designate a path as active, it is continuously scheduled for I/O according to the load-balancing and failover policy in use.

`mode=standby`

The path is held in reserve. Being set to standby does not mean a path will not be used. It only means the weight of the path is heavily adjusted to preclude its use in normal operations. A standby path still can be selected if it is the best path for a request. For example, if the active paths are so heavily loaded that the weighting factor in favor of the active path is overcome, the I/O will be sent to the standby path. A standby path returns to the default setting or exits standby mode if the host is rebooted without saving the settings using `powermt save`. Using `powermt save` persists standby mode settings across reboots.

`hba=<hba#>|all`

Sets the mode for paths from the specified HBA. *hba#* is a number in the **Host Bus Adapters ###** column of **powermt display dev** and **powermt display bus** output. **all** specifies all HBAs under PowerPath control. The default is **all**.

(**adapter** can be used instead of **hba**.

**adapter** is included only for compatibility with prior PowerPath releases. Use **hba** instead.)

`dev=<path>|<device>|all`

Sets the mode for the specified *path*, or all paths to the specified *device*. **all** specifies all paths to all devices. The default is **all**.

`class={symm|vnx|clariion|vplex|invista|netapp|ess|xtremio|hitachi|hpxp  
|generic|all}`

Limits the command to the specified type of storage system. **all** specifies all storage-system types. The default is **all**. For descriptions of all the class types, see [“Command syntax: class argument specification” on page 24](#).

If class is specified along with hba or dev, the command is applied to paths that meet all specified constraints.

`force`

This argument manually overrides autostandby. If a path is in autostandby mode, you must use the force argument to change it to active or standby.

## powermt set path\_latency\_monitor

**Purpose** Enable or disable path latency monitoring.

**Syntax** `powermt set path_latency_monitor={on|off}`

**Description** **powermt set path\_latency\_monitor** enables or disables path latency monitoring. When you enable or disable path latency monitoring, the setting applies globally to all paths.

When path latency monitoring is enabled, PowerPath measures the amount of time each I/O request takes to complete on the selected path.

---

**Note:** Before enabling latency monitoring for the first time, you should not only read this section but also the next section on setting a threshold for path latency. See [“powermt set path\\_latency\\_threshold” on page 95](#).

---

The **powermt display options** command shows if path latency monitoring is enabled or disabled. When path latency monitoring is enabled, **powermt display latency** shows the most recent (Current) and the high watermark (Max) I/O completion time for each path on the host.

By sampling outputs from **powermt display latency** you can infer the expected path latencies in your environment, and thus be able to set an appropriate threshold for the generation of meaningful latency events in your system log. To set a path latency threshold, use the **powermt set path\_latency\_threshold** command.

### Arguments

`path_latency_monitor=on`  
Enables I/O latency measurement on each path, resulting in meaningful outputs from **powermt display latency**, which are the Current (most recent) and Max (high watermark) latencies for each path.

`path_latency_monitor=off`  
Disables I/O latency measurement for all paths, and resets all paths' Current and Max latencies to zero. The default setting is **off**.

Disabling latency monitoring does not change a previously set threshold. On re-enabling, the last threshold entered goes back into effect.

## powermt set path\_latency\_threshold

**Purpose** Set the path latency threshold.

**Syntax** `powermt set path_latency_threshold=<#seconds>|<#milliseconds>ms`

**Description** `powermt set path_latency_threshold` sets a time interval in seconds within which I/Os should complete. This threshold value applies to all paths in your environment.

When a threshold has been set, PowerPath generates system log messages indicating each threshold crossing that results in a new Max latency (high watermark) for a path.

For example, if

- ◆ the global `path_latency_threshold` is set to 2, and
- ◆ the current Latency Max for the path (as seen in `powermt display latency`) is 2.5 seconds (2500ms), and
- ◆ an I/O request on the path takes 2.6 seconds to complete,

then the threshold-crossing event is sent to the system log.

---

**Note:** For environments where 24x7 latency monitoring is in effect, periodically disable, then re-enable latency monitoring so that high watermarks are zero-ed and threshold crossings are captured in the system log with regularity.

---

When changing the threshold, use the following three steps:

1. Disable path latency monitoring so that high watermarks are zeroed.
2. Set the new threshold.
3. Re-enable path latency monitoring.

To enable path latency monitoring, use the `powermt set path_latency_monitor` command. To view information on I/O completion times, use the `powermt display latency` command.

**Log message example** When an I/O completion time exceeds the threshold value, PowerPath sends to the log file a message similar to the following:

Message Format:

```
<date> <time> <host_name> emcp: [ID 801593 kern.notice] Error: Latency Threshold
exceeded by UserDev Bus <bus_number> Tgt <target> Lun <LUN> Vol <volume> with a
new high latency of <latency_value> msec
```

Message:

```
Mar 30 12:32:40 abc01.lss.emc.com emcp: [ID 801593 kern.notice] Error: Latency
Threshold exceeded by UserDev Bus 3072 Tgt 50001fe15000b02c Lun 2 Vol
600508B4000147160002E00000370000 with a new high latency of 10 msec
```

**Arguments** `path_latency_threshold=<#seconds>|<#milliseconds>ms`  
 Sets a time interval in seconds within which I/Os should complete. The value applies to all paths. The default value is 0. The range of acceptable values is 0 to 3600 seconds or 0 to 3600000 milliseconds. Seconds is the default.

When the threshold is set to zero (also known as Discovery Mode), PowerPath logs every new Max latency for each path. The resulting system log messages can be helpful in determining an appropriate threshold for the system.

---

**Note:** Discovery Mode goes into effect if path latency monitoring is enabled before a threshold has been set. If the increased volume of system log messages that results from Discovery Mode is not desired, be sure to set an exceptionally high threshold (for example, 30 seconds) before enabling latency monitoring for the first time. You can then sample **powermt display latency** output to estimate expected latencies for your environment.

---



## powermt set perfmon

<b>Purpose</b>	Enables or disables performance monitoring for all devices.
<b>Syntax</b>	<code>powermt set perfmon={on [interval=&lt;#seconds&gt;]   off}</code>
<b>Description</b>	<p><b>powermt set perfmon</b> enables or disables performance monitoring for all devices. Performance monitoring helps characterize I/O patterns and possibly aide in diagnosing I/O problems.</p> <p>When performance monitoring is enabled, <b>powermt display perf dev=all</b> and <b>powermt display perf bus</b> displays performance metrics for all paths to all devices and bus ports. See <a href="#">“powermt display perf dev” on page 62</a> and <a href="#">“powermt display perf bus” on page 65</a> for more information. To view whether performance monitoring is enabled, use the <b>powermt display options</b> command.</p>
<b>Arguments</b>	<p><code>perfmon=on</code> Enables performance monitoring for all devices. There is no output. This command initializes the counters, clears all performance measurements including high and low watermarks for latency, and accepts a sampling interval for data collection in seconds with a default interval of 900 seconds (15 minutes). The legacy path latency monitoring threshold measurements are not affected.</p> <p><code>perfmon=off</code> Disables performance monitoring and clears all measurements.</p> <p><code>interval=&lt;#seconds&gt;</code> The sampling interval in seconds for data collection. The sampling interval range is from a minimum of 60 seconds (1 minute) to a maximum of 86,400 seconds (1,440 minutes or 1 day) with a default of 900 seconds (15 minutes).</p> <p>To change the performance monitoring sampling interval, you must first <b>set perfmon=off</b> and then <b>set perfmon=on</b>.</p>
<b>Examples</b>	<p>This command example enables performance monitoring on all devices:</p> <pre># powermt set perfmon=on</pre> <p>This command example disables performance monitoring on all devices:</p> <pre># powermt set perfmon=off</pre>

## powermt set periodic\_autorestore

**Purpose** Enable periodic autorestore.

**Syntax** `powermt set periodic_autorestore={on|off} [class=<class>|all]`

**Description** `powermt set periodic_autorestore` enables or disables the periodic autorestore feature.

When periodic autorestore is **on**, PowerPath periodically tests dead paths and, if they pass the test, restores them to service. Even when periodic autorestore is **off**, path testing continues to be done under certain conditions as described in the *EMC PowerPath Family Product Guide*, and automatic path restoration continues to occur based on the results of that testing.

To determine the periodic autorestore setting, use `powermt display options`.

`periodic_autorestore` can be used along with `reactive_autorestore` to define when paths are restored after path failure. [“powermt set reactive\\_autorestore” on page 105](#) provides more information on reactive autorestore.

### Arguments

`periodic_autorestore={on|off}`

Enables (**on**) or disables (**off**) periodic autorestore. The default is **on**.

`class={symm|vnx|clariion|vplex|invista|netapp|ess|xtremio|hitachi|hpxp|generic|all}`

Limits the command to the specified type of storage system. **all** specifies all storage-system types. The default is **all**. For descriptions of all the class types, see [“Command syntax: class argument specification” on page 24](#).

## powermt set policy

**Purpose** Set the load-balancing and failover policy.

**Syntax** `powermt set policy={ad|co|lb|li|re|rr|si|so}  
[dev=<device>|all] [class=<class>|all]`

**Description** `powermt set policy` sets the load-balancing and failover policy for devices configured in PowerPath.

---

**Note:** Unlicensed versions of PowerPath support EMC arrays only. This configuration is supported if the host has a single HBA only. This configuration is also referred to as PowerPath/SE. The *EMC PowerPath Family Product Guide* provides more information. With third-party arrays in an unlicensed PowerPath environment, either unmanage the third-party array class (`powermt unmanage class=<class>`) or upgrade to a licensed version of PowerPath.

---

**Arguments** `policy={ad|co|lb|li|re|rr|si|so}`  
Sets the load-balancing and failover policy to one of the values, as shown in [Table 21 on page 100](#).

The *EMC PowerPath Family Product Guide* provides a summary of the platform, array, and feature support available with each type of PowerPath license.

**Table 21** powermt set policy (page 1 of 2)

Policy	Available storage classes	Load balancing enabled?	Failover enabled?	Description
ad (Adaptive)	vplex invista ess hitachi hpxp xtremio netapp	✓	✓	I/O requests are assigned to paths based on an algorithm that takes into account path load . With a valid PowerPath license, this policy is the default for the storage classes listed in the “ <a href="#">Available storage classes</a> ” column.
bf (Basic failover)	symm vnx clariion vplex invista		✓ (only the SP fails over)	<p><b>Note:</b> The basic failover policy is supported only in configurations with one HBA connected to a storage system through a switch and one HBA port zoned to each SP on a VNX and CLARiiON system, or to a port on two separate FAs on a Symmetrix system. Using this policy in other configurations is not supported and may result in loss of access to data in the event of path failures.</p> <p>Load balancing is not in effect. I/O routing on failure is limited to one HBA and one port on each storage system interface. When a host boots, it designates one path (through one interface) for all I/O. If an I/O is issued to a logical device that cannot be reached via that path (that is, the I/O cannot reach that logical device through the device's assigned interface), the logical device is assigned to the other interface.</p> <p>This policy protects against VNX and CLARiiON SP failures, Symmetrix FA port failures, and back-end failures, and it allows non-disruptive upgrades to work when running PowerPath without a license key. It does not protect against HBA failures.</p> <p>This policy is the default policy without a PowerPath license for EMC storage arrays.</p> <p><b>HP-UX Note:</b> HP-UX hosts designate the first two paths to a volume for use as basic failover paths; the remaining paths are unlicensed.</p>
co (CLARiiON optimization)	vnx clariion	✓	✓	I/O requests are assigned to paths based on an algorithm that takes into account path load . This policy is valid only for VNX and CLARiiON storage classes and is the default policy for them, on platforms with a valid PowerPath license.
lb (Least blocks)	All	✓	✓	Load balance is based on the number of blocks in pending I/Os. I/O requests are routed to the path with the fewest queued blocks, regardless of the number of requests involved.
li (Least I/O)	All	✓	✓	Load balance is based on the number of pending I/Os. I/O requests are routed to the path with the fewest queued requests, regardless of total block volume.

**Table 21** powermt set policy (page 2 of 2)

Policy	Available storage classes	Load balancing enabled?	Failover enabled?	Description
nr (No redirect)	ess hitachi hpxp			<p><b>Note:</b> Do not use this policy in production environments. Use only for diagnostic purposes.</p> <p>Neither load balancing nor path failover is in effect. If <b>nr</b> is set on a pseudo device and the I/O path fails, data errors can occur. If <b>nr</b> is set on a native device and the I/O path fails, data errors will occur.</p> <p>This policy is the default policy without a PowerPath license for third-party storage arrays.</p> <p><b>Note:</b> If the policy is <b>nr</b> and there is only one path, I/O to that path is not guaranteed.</p>
re (Request)	All (PowerPath license), CLARiiON only (PowerPath Base license)		✓	For native devices, this policy uses the path that would have been used if PowerPath were not installed. For pseudo devices, it uses one arbitrary path for all I/O. For all devices, path failover is in effect, but load balancing is not.
rr (Round robin)	All	✓	✓	I/O requests are assigned to each available path in rotation.
si (StreamIO)	All	✓	✓	<p>For each I/O to a volume, streamio policy uses the same path that was used for the previous I/O to the volume. However, if the volume's I/O count exceeds the volume's stream I/O threshold since the last path selection, the policy selects a path using the adaptive policy algorithm (see <a href="#">“ad (Adaptive)” on page 100</a>). The volume I/O count is rezeroed on each path change.</p> <p><a href="#">“powermt set streamio_threshold” on page 107</a> provides additional information on setting the threshold values for when to switch paths.</p> <p><a href="#">“Multipathing new and changed commands” on page 14</a> lists the versions for each platform that support this policy.</p>
so (Symmetrix optimization)	symm	✓	✓	<p>I/O requests are routed to paths based on an algorithm that takes into account path load. Load is a function of the number, size, and type of I/O queued on each path.</p> <p>This policy is valid only for Symmetrix storage classes and is the default policy for them, on platforms with a valid PowerPath license.</p>

```
class={symm|vnx|clariion|vplex|invista|netapp|ess|xtremio|hitachi|hpxp|generic|all}
```

Limits the command to the specified type of storage system. **all** specifies all storage-system types. The default is **all**. For descriptions of all the class types, see [“Command syntax: class argument specification” on page 24](#).

If **class** is specified along with **dev**, the command is applied to devices that meet all specified constraints.

```
dev=<device>|all
```

Limits the policy change to the specified device. **all** specifies all devices. The default is **all**.

If neither `class` nor `dev` is specified, and the specified *policy* does not apply to any storage system in the environment (for example, if the `bf` policy is specified with a PowerPath license, but there are no VNX and CLARiiON storage systems), no policy setting occurs, and no error message appears.

### **powermt set policy changes**

Changes to **powermt set policy** include:

- ◆ As of version 5.5, the BasicFailover (BF) and NoRedirect (NR) policies have been removed from the **powermt set policy** command usage. In subsequent releases, it will not be possible to manually set these policies. However, on hosts without a valid PowerPath license, BF and NR will continue to be the default policies for EMC and third-party arrays, respectively.
- ◆ If **powermt set policy=<co|so>** is specified without the **class** option and you do not have the appropriate devices then no policy change occurs. You will not receive an error message.

## powermt set port\_disable

**Purpose** Enable or disable a storage system port.

**Syntax** `powermt set port_disable={on|off} dev=<device>`

**Description** `powermt set port_disable` enables or disables a storage system port, and closes all devices configured through the port. To disable a port, type the name of a native-named device configured through the port. Once disabled, all paths configured through a port are marked **dead** and no I/O is sent to the port.

Use `powermt set port_disable` when, for example, you want to replace a disk array controller. To display information on storage-system ports, use the `powermt display port_mode` command.

Note the following limitations:

- ◆ A disabled port state persists across host reboots; however, the complete suppression of I/O from the host to the disabled array port is not guaranteed during early boot (of the host).
- ◆ In an R1/R2 boot failover configuration, a disabled port does not persist through a failover or fallback.
- ◆ This command fails if the PowerPath load balancing and failover policy is **bf** or **nr**.

**Arguments** `port_disable=on|off`  
Disables (**on**) or enables (**off**) a port on the specified device. The default is **off**.

`dev=<device>`  
Specifies the device port you want to enable or disable. The device you enter must be a native-named device. This **dev** argument applies to all paths using the port.

**Example** The following example shows the command to disable a port and the subsequent `powermt display port_mode` output.

To disable the specified devices, type `powermt set port_disable=on dev=c5t50060482CAFD7A5Dd1s0`

To view the disabled port, type `powermt display port_mode dev=c5t50060482CAFD7A5Dd1s0`

Output similar to the following appears:

```
Pseudo name=emcpower27a
Symmetrix ID=000187430377
Logical device ID=0001
state=alive; policy=SymmOpt; queued-I/Os=0
=====
----- Host ----- Stor-- I/O Path-- Stats-
### HW Path          I/O Paths          Interf. Mode State Q-I/Os Errors
=====
3072 pci@1c,600000/SUNW,qlc@1/fp@0,0 c3t50060482CAFD7A7Cd1s0 FA 13dB active alive 0 0
3074 pci@1d,700000/SUNW,qlc@1/fp@0,0 c5t50060482CAFD7A5Dd1s0 FA 14bA active dead 0 0
```

## Platform-specific notes

On a Solaris 10 host, you may see the following error message when you try to enable a disabled port:

```
Bad dev value <device>, or not under PowerPath control.
```

If this occurs, run the Solaris commands shown in the following procedure, and then enable the port.

1. Type **cfgadm -vl** to learn the controller number for the card on each host.
2. For a disabled port, type **cfgadm -vc configure cX**  
where *X* is the controller number for the card  
Optionally, type **devfsadm**.
3. Type **powermt set port\_disable=off dev=<device>**  
where *device* is the offline device configured through the port you want to disable.



## powermt set reactive\_automorestore

**Purpose** Enable or disable reactive automorestore.

**Syntax** `powermt set reactive_automorestore={on|off} [class=<class>|all]`

**Description** `powermt set reactive_automorestore` enables or disables PowerPath’s reactive automorestore facility. Reactive automorestore is the automatic restoration in the context of an I/O of physically alive paths that are still in the dead (unusable) state.

When reactive\_automorestore is **on**, PowerPath reactively tests dead paths and, if they pass the test, restores them to service. For non-disruptive upgrades (NDU), reactive\_automorestore should be set to **on**.

To determine the reactive\_automorestore setting, use **powermt display options**.

**reactive\_automorestore** can be used along with **periodic\_automorestore** to define when paths are restored after path failure. “[powermt set periodic\\_automorestore](#)” on page 98 provides more information on periodic automorestore. [Table 22 on page 105](#) lists the automorestore scenarios.

**Table 22** Automorestore scenarios

If	and	then
reactive_automorestore=on	periodic_automorestore=on	paths will be automatically restored at the next opportunity in either context. This is the default for these settings.
reactive_automorestore=on	periodic_automorestore=off	a physically alive path will be automatically restored in the context of an I/O if there is no other path in the alive state available to complete the I/O.  <b>Note:</b> Any qualified path from either SPA or SPB can be restored to service. Therefore, if HBAs are getting enabled in this scenario, the I/Os will go through any path that is brought alive first. The storage processor of the serving path might take ownership of the LUN resulting in a reassign, restore, or follow event.
reactive_automorestore=off	periodic_automorestore=on	the path will be restored when the path is selected for path test.
reactive_automorestore=off	periodic_automorestore=off	<b>Note:</b> It is not recommended this setting in your normal running environment.  paths that move to the dead state due to path test failures will never be automatically restored. However, physically alive paths that are proactively moved to the dead state as a result of another path failure on the same bus may be automorestored and returned to the alive state during the next path test. This path test may be either periodic or in the context of completing an I/O (reactive). To restore physically dead paths in this scenario run <b>powermt restore</b> at the command line or from the PowerPath Administrator console on Windows.

**Note:** `powermt set reactive_automorestore` is not supported on Windows.

**Arguments**`reactive_autorestore=on`

Enables reactive autorestore, which restores the path if I/Os are issued over the path. The default is **on**.

`reactive_autorestore=off`

Suppresses path testing and restoration in the context of I/O, thereby quickening I/O failure when all paths to a volume are dead. Turning reactive autorestore off can be useful in failover clusters, or with host-based disk mirroring, if you place a higher priority on these failover mechanisms than you place on PowerPath failover.

`class={symm|vnx|clariion|vplex|invista|netapp|ess|xtremio|hitachi|hpxp|generic|all}`

Limits the command to the specified type of storage system. **all** specifies all storage-system types. The default is **all**. For descriptions of all the class types, see [“Command syntax: class argument specification” on page 24](#).

## powermt set streamio\_threshold

**Purpose** Set the I/O threshold value for when to switch paths.

**Syntax** `powermt set streamio_threshold=<threshold_count>  
[dev=<device>|all] [class=<class>|all]`

**Description** `powermt set streamio_threshold` sets the I/O threshold values for when to switch paths while the Stream I/O (si) load-balancing and failover policy is in effect. [“si \(StreamIO\)” on page 101](#) provides more information.

You can check whether Stream I/O threshold is set by viewing the `powermt display dev` output. For example:

```
state=alive; policy=StreamIo; streamio_threshold=1000; queued-IOS=0
```

If the PowerPath license is invalid, setting the `streamio_threshold` will generate an error message.

**Arguments**

`streamio_threshold=<threshold_count>`  
Integer in the range from 64 to 2048 that specifies, in number of I/Os, when to switch paths. The default for EMC arrays is 1024 and the default for third-party arrays is 128.

`dev=<device>|all`  
Limits the change to the specified logical device. **all** specifies all logical devices. The default is **all**.

`class={symm|vnx|clariion|vplex|invista|netapp|ess|xtremio|hitachi|hpxp|generic|all}`  
Limits the command to the specified type of storage system. **all** specifies all storage-system types. The default is **all**. For descriptions of all the class types, see [“Command syntax: class argument specification” on page 24](#).

## powermt set write\_throttle

**Purpose** Enable write throttling.

**Syntax** `powermt set write_throttle={on|off}  
[class={symm|all}] [dev=<device>|all]`

**Description** `powermt set write_throttle` enables or disables I/O write throttling to the specified logical devices.

Write throttling is enabled to limit the number of queued writes to the common I/O queue in the HBA driver; instead, the writes are queued in PowerPath. As a result, read requests do not get delayed behind a large number of write requests. Write throttling is disabled by default.

**Arguments**

`write_throttle={on|off}`  
Enables (on) or disables (off) write throttling to the specified logical device. For new logical devices, write throttling is set to **off** by default.

`class={symm|all}`  
Limits the command to the specified type of storage system. **all** specifies all storage-system types eligible for write throttling. The default is **all**. For descriptions of all the class types, see [“Command syntax: class argument specification” on page 24](#).

If class is specified along with **dev**, the command is applied to logical devices that meet all specified constraints.

`dev=<device>|all`  
Limits the change to the specified logical device. **all** specifies all logical devices eligible for write throttling. The default is **all**.

`<volume_dev>` can be used instead of `<dev>`. `<volume_dev>` is included only for compatibility with prior PowerPath releases. Use `<dev>` instead.

## powermt set write\_throttle\_queue

**Purpose** Set write throttling queue depth.

**Syntax** `powermt set write_throttle_queue=<queue_depth>  
[class={symm|all}] [dev=<path>|<device>|all]`

**Description** `powermt set write_throttle_queue` sets the write throttling queue depths for a storage-system port connected to a specified device. The queue-depth setting limits the number of writes to all devices enabled for write throttling that can be outstanding (from PowerPath's perspective) on the storage-system port. The queues are allocated within PowerPath, one per storage-system port.

**Arguments**

`write_throttle_queue=<queue_depth>`  
Limits the number of I/Os sent or written to a port. The intent of setting this threshold value is to prevent write operations from delaying the completion of read operations. The value of *queue\_depth* must be in the range 0 to 2048. For new ports, the queue depth is set to 256 by default.

`class={symm|all}`  
Limits the command to the specified type of storage system. **all** specifies all storage-system types eligible for write throttling. The default is **all**. For descriptions of all the class types, see [“Command syntax: class argument specification” on page 24](#).  
  
If class is specified along with `dev`, the command is applied to ports that meet all specified constraints.

`dev=<path>|<device>|all`  
Limits the change to the storage-system port associated with the specified *path*, or the port connected to the specified *device*. **all** specifies all storage-system ports to all devices eligible for write throttling. The default is **all**.  
  
<port\_dev> can be used instead of <dev>. <port\_dev> is included only for compatibility with prior PowerPath releases. Use <dev> instead.

## powermt unmanage

**Purpose** Exclude devices or storage system classes from PowerPath management.

**Syntax** `powermt unmanage dev=<path>|<device> | class=<class>`

**Description** **powermt unmanage** excludes a specified logical device or storage system class from PowerPath management.

When applied to a storage system class, **powermt unmanage** succeeds in completely removing the class from PowerPath control only after you stop applications with open devices of the specified class. If you cannot stop applications, you may need to reboot the host.

When you remove a storage system class from PowerPath control, custom configurations for that class (such as policy or mode) are lost with the next reboot or execution of **powermt save**.

A **powermt unmanage** class operation is legal only if the class's status is managed or unmanage\_incomplete. If the class status is unmanage\_incomplete, the manage operation must be completed—or the host rebooted—before **powermt unmanage** succeeds. The `powermt display options` command displays the status of a class.

When applied to a specific device, **powermt unmanage** differs from **powermt remove**:

- ◆ **powermt unmanage** applies to the entire device, not just the specified path.
- ◆ The effect of **powermt unmanage** is persistent: **powermt config** does not restore the device to PowerPath control.

Mappings of pseudo names to native names do not persist when you unmanage and then remanage devices.

### Arguments

`dev=<path>|<device>`  
Specifies the pathname of a logical device to be excluded from PowerPath management.

`class={symm|vnx|clariion|vplex|invista|netapp|ess|xtremio|hitachi|hp  
|generic}`  
Specifies the storage system class to be excluded from PowerPath management. For descriptions of all the class types, see [“Command syntax: class argument specification” on page 24](#).

If **class** is specified along with **dev**, the command is applied to HBAs or devices that meet all specified constraints.

## powermt update lun\_names

**Purpose** Update user-assignable LUN names.

**Syntax** `powermt update lun_names`

**Description** **powermt update lun\_names** causes PowerPath to retrieve the latest VNX, CLARiiON, XtremIO, and IBM XIV user-assignable LUN names. A user-assignable LUN name is a character string that a user or system manager associates with a LUN and assigns through Navisphere.

## powermt version

**Purpose** Get the PowerPath version number.

**Syntax** `powermt version`

**Description** `powermt version` prints the version of PowerPath that is installed on the host. Versions later than PowerPath 5.1 (for example, PowerPath 5.1 SP1 or PowerPath 5.1 with hot fix 1), include the following **powermt version** changes:

- ◆ Installed hot fixes are displayed by default.
- ◆ PowerPath version names use the term service pack (SP) for 5.1.x level releases (for example, 5.1 SP1 instead of 5.1.1) and patch (P) for 5.1.x.x level releases (for example, 5.1 SP1 P01).
- ◆ The `PP_SHOW_VERSION_STYLE` environment variable allow you to view **powermt version** output in the legacy format.

**Examples** The following example shows the output for PowerPath 5.7:

```
powermt version
EMC powermt for PowerPath (c) Version 5.7 build (421)
```



# emcpreg: Manage PowerPath license registration

**Platforms** This command is supported on UNIX and Linux platforms, and PowerPath 5.2 and later for Windows.

---

**Note:** PowerPath 5.2 and 5.3 for Windows supports only **emcpreg -add** and **emcpreg -remove**, while PowerPath 5.3 SP1 and later for Windows supports all **emcpreg** arguments.

---

**Syntax** On UNIX and Linux hosts:

```
emcpreg [-f pathname] -add key [key ...]
emcpreg [-f pathname] -check key [key ...]
emcpreg [-f pathname] -edit
emcpreg [-f pathname] -install
emcpreg [-f pathname] -list
emcpreg [-f pathname] -remove key [key ...]
```

On PowerPath 5.3 SP1 and later for Windows hosts:

```
emcpreg -add key [key ...]
emcpreg -check key [key ...]
emcpreg -edit
emcpreg -install
emcpreg -list
emcpreg -remove key [key ...]
```

**Description** **emcpreg** manages PowerPath license registration. On some platforms, **emcpreg** is invoked automatically by the PowerPath installation processes. On other platforms, you must invoke the command manually. Once installation is complete, **emcpreg** can be invoked manually on any supported platform.

To run **emcpreg**, you must be logged in as the root user.

The **emcpreg** utility stores license keys in the registration database, by default in `/etc/emcp_registration` (on UNIX and Linux) and in the Windows registry (on Windows).

---

**Note:** **emcpreg** only adds or removes licenses from the *database*. The licenses are not activated until one of the following commands are executed: **powermt config**, **powermt check\_registration**, **powermt load**, or **reboot**.

---

**Arguments**

`-f <pathname>`  
Provides the optional pathname of the registration database for UNIX and Linux hosts (this argument does not apply to Windows hosts). The default is `/etc/emcp_registration`. This file should not be manually edited. It also should not be moved, as that will prevent installed products from finding valid license keys.

`-add <key> [<key> ...]`  
Adds one or more keys to the registration database.

---

**Note:** To activate the license, perform this additional step: run the **powermt config**, **powermt check\_registration**, **powermt load**, or the **reboot** command.

---

- check <key> [<key> ...]  
Checks the validity of one or more keys and display the product information associated with each key.
- edit  
Edits the product information database. Each key in the database is displayed, and the user is prompted to remove or retain it.
- install  
Installs license key for PowerPath.
- list  
Lists the product information for all keys in the registration database.
- remove <key> [<key> ...]  
Removes one or more keys from the registration database. The keys may be added back later if needed.

---

**Note:** To change the driver state, you need to do an additional step: run the **powermt config**, **powermt check\_registration**, **powermt load**, or the **reboot** command.

---

# emcphostid: Set the Host ID

**Platforms** This command is supported on HP-UX (11i v2.0 and later) and Windows.

**Syntax** On HP-UX hosts:

```
emcphostid set [id=<id>] [-no_prompt]
emcphostid get
emcphostid help
```

On PowerPath 5.5 and later for Windows:

```
emcphostid.exe set [-id <id>] [-host <HostName>] [-no_prompt]
emcphostid.exe get [-host <HostName>]
emcphostid.exe check -host <HostName1> <HostName2> ...<HostNameN>
emcphostid.exe help
```

**Description** PowerPath requires a unique host ID to reserve shared disks in clustered systems. On HP-UX hosts, the PowerPath installation sets the host ID to the largest numerical MAC address or the hash of the Machine ID found at the time of installation. PowerPath installation sets the host ID to the largest numerical MAC address or the hash of the Machine ID found at the time of installation. On Windows hosts, the host ID is created by hashing the NetBIOS name of the computer on every boot. The host ID is saved in `/etc/emc/emcphostid` on HP-UX hosts.

On Windows hosts, run the Microsoft sysprep utility if the same operating system image is used to set up cluster nodes to ensure that each node receives a unique host ID.

## For Windows Hosts

After it is set by the installation, the host ID rarely needs to be changed. It is also not recommended to change this ID unless required. You should, however, use **emcphostid** to change the host ID in the following circumstances:

- ◆ When there is a conflict. The host ID can be overwritten by using a user specified value. This entry is stored in the registry. Setting the host ID to match the default host ID replaces the user specified value with the default ID.
- ◆ The same image of Windows operating system is run on several cluster nodes and you do not run the Microsoft sysprep utility.
- ◆ The machine security identifiers on two or more cluster nodes produced the same hash values.

## For HP-UX hosts

After set by the installation, the host ID rarely needs to be changed. You should, however, use **emcphostid** to change the host ID in the following circumstances:

- ◆ An Ethernet card is removed and may be installed elsewhere.

---

**Note:** If the Ethernet card will not be reused, for example if it is damaged, you need not change the host ID.

---

- ◆ A SAN-booted host is swapped; for example, a two-processor host is replaced by a four-processor host.

- ◆ A host has no network cards with MAC addresses at the time of PowerPath installation.

In this case, you must use the `id=<id>` argument to the **emcphostid set** command to set the host ID.

- ◆ A host had no valid MAC address at the time PowerPath was installed but now has one.

## Commands

**set** [-id <id>] [-host <HostName>] [-no\_prompt]  
 Sets or resets the host ID. By default, if a host ID already exists, prompts for confirmation before resetting the host ID. Use **-no\_prompt** to suppress this prompt. **-host <HostName>** does not apply to HP-UX hosts.

Without the *id* argument, **emcphostid.exe set** sets the host ID to the highest MAC address found on the host or the hash of the Machine ID at the time the command is run.

If the MAC address or Machine SID cannot be retrieved, a warning is displayed, and the host ID is set to 0xFFFFFFFFFFFFFF (HP-UX) or not set or changed (Windows). On HP-UX, you can rerun the command using the `id=<id>` argument.

Depending on the platform, you can specify *id* in one of several ways:

**Table 23** Platform specific IDs

On platform	Specify host ID as
Windows	An 8-digit hexadecimal number with or without the 0x prefix
HP-UX	A 16-character (8 byte) hexadecimal string without the 0x prefix (for example 1a2b3c4d5e6f7890)

**get** [-host <HostName>]  
 Prints the current host ID. **emcphostid.exe get** retrieves and prints the host ID that is currently stored in `/etc/emc/emcphostid` (HP-UX) or the Windows registry.

**-host <HostName>** does not apply to HP-UX hosts.

On Windows, **emcphostid.exe get** also checks to see whether the current host ID differs from the default host ID, and if so, prints a warning.

**check** -host <HostName1> <HostName2> ... <HostNameN>  
 Checks the uniqueness of host IDs for the specified list of hosts. This argument does not apply to HP-UX.

On Windows, run **emcphostid.exe check** to check for conflicts in the following scenarios:

- After installing PowerPath in cluster environment.
- After changing the host ID manually.

**help**  
 Provides help for the **emcphostid** commands.

## Changing the host ID

To change the host ID on Windows:

1. Reboot the host if necessary (for example, if you are replacing host hardware or a network card that requires reboot).

2. Run the **emcphostid.exe set** command.
3. Reboot the host.

## emcpadm: List or rename PowerPath pseudo devices

**Platforms** This command is supported on AIX, Linux, and Solaris.

**Description** The **emcpadm** command lists available PowerPath pseudo device names and used PowerPath pseudo device names or renames a PowerPath pseudo device from a used pseudo device name to an available pseudo device name. This command requires root privileges.

---

**Note:** The **emcpadm** command cannot rename pseudo devices that are in use. For example, devices that are part of a VxVM disk group are considered in use, and therefore cannot be renamed.

---



**Using the `rendev` command to rename a pseudo device fails. To rename a pseudo device, use the `emcpadm renamepseudo` command only.**

---

**New Features** The PowerPath command line utility, `emcpadm`, has been enhanced to export and import pseudo device name and LU (logical unit) data in XML format.

A command line option, `-x`, has been added to `emcpadm export_mappings`, `import_mappings`, and `check_mappings`. This directs `emcpadm` to export or import data using an XML format. The XML data is compatible with that viewed via tools such as PowerPath **powermt** and Symmetrix **syminq**.

### Supported arrays

The enhanced `emcpadm` with XML-formatted mapping data supports importing and exporting of Symmetrix, VNX, and CLARiiON device arrays as well as that of third-party arrays. The importing and exporting of mapping data in non-XML format is limited to the arrays supported by PowerPath.

## Syntax

**Note:** The options **[-f {i|d|m}]** in the **emcpadm getfreepseudo** and **emcpadm getusedpseudo** commands apply only on Solaris.

Pseudo device/LU mappings:

```
-----
emcpadm check_mappings [-v] [-x] -f <pseudo device/LU mappings file>
emcpadm export_mappings [-x] -f <pseudo device/LU mappings file>
emcpadm import_mappings [-v] [-x] -f <pseudo device/LU mappings file>
emcpadm print_mappings [-p] | -f <pseudo device/LU mappings file>
```

Pseudo device/LU renaming:

```
-----
emcpadm getfreepseudos [-f {i|d|m}] [-n <# of pseudo instances>] [-b <start pseudo device>]
emcpadm getusedpseudos
emcpadm renamepseudo [-f {i|d|m}] -s <src pseudo device> -t <tgt pseudo device>
```

Other:

```
-----
emcpadm help
emcpadm version
```

## Commands

### Pseudo device/LU mappings

#### check\_mappings

The `check_mappings` output displays a comparison of the current mappings and the device remappings in *mappings\_file*. Use this command to preview device remappings before importing a mappings file:

```
check_mappings [-v] [-x] -f <pseudo device/LU mappings file>
```

where **-x** directs `emcpadm` to read from or write to the PowerPath pseudo device name to LU mappings in XML format to the file named by the **-f** option. This command functions similarly to the `import_mappings` command but only checks which currently mapped devices will be affected by an import of data contained a given mappings file.

#### Arguments

**-f** <pseudo device/LU mappings file>

Specifies the file name and location for *mappings\_file*.

**-v**

Specifies verbose mode.

**-x**

Specifies XML format for the mappings file data. You must use the XML format to manually edit the mappings file.

This will be displayed on your screen:

```
# ./emcpadm check_mappings -x -f map1.xml
```

```

Imported Device   Remapping Status
Mapping
-----
emcpower0a       no change
emcpower1000a    remaps: emcpower1a
emcpower26a      no change
```

emcpower24a	no change
emcpower21a	no change
emcpower18a	no change
emcpower17a	no change
emcpower16a	no change
emcpower11a	no change
emcpower40a	no change

## export\_mappings

export\_mappings writes and exports the current mappings to the named file.

**Note:** Using a mapping file that is not updated with new device IDs would show error messages about mapping failures. Therefore, export and import the mapping file only after the migration is complete.

```
rename [pseudo] [-f{i|d|m}] -s <src pseudo device> -t <tgt pseudo device>
```

```
export_mappings [-x] -f <pseudo device/LU mappings file>
```

### Arguments

-f <pseudo device/LU mappings file>

Specifies the file name and location for *mappings\_file*.

-v

Specifies verbose mode.

-x

Specifies XML format for the mappings file data. You must use the XML format to manually edit the mappings file.

## import\_mappings

import\_mappings replaces the current mappings with the mappings in *mappings file*. If differences exist among the current mappings and the file mappings, the mappings in *mappings file* take precedence. When you import the file mappings, current host devices are remapped according to the file mappings, where differences exist.



**If the pseudo device was renamed earlier with the `rendev` command, the `emcpadm import_mappings` command fails to import the mapping file.**

```
import_mappings [-v] [-x] -f <pseudo device/LU mappings file>
```

### Arguments

-f <pseudo device/LU mappings file>

Specifies the file you want to import.

-v

Specifies verbose mode, which displays remapping information as it occurs.

-x

Specifies XML format for the mappings file data. You must use the XML format to manually edit the mappings file.

Before importing new mappings on a node or server:



- Preview changes with **emcpadm check\_mappings**.
- Shut down all applications and database systems.
- Unmount file systems.
- Deport VxVM disk groups.

## print\_mappings

```
print_mappings [-p | -f <pseudo device/LU mappings file>]
```

**print\_mappings** (without arguments) displays the current mappings on the node or server.

**print\_mappings -f** <pseudo device/LU mappings file> displays the exported mappings in <pseudo device/LU mappings file>.

**print\_mappings -p** <pseudo device/LU mappings file> displays the persistent mappings maintained in the PowerPath device-mapping database. These mappings should be the same as the mappings on the current node or server.

Print mappings displays output for each device, as follows:

```
Device: emcpower1a
Device Id:
  Array LU identifier:
    4942 4d20 2020 2020 2020 3231 3035 2020 2020
    2020 2020 2020 2020 3745 3131 3537 3438
(Output truncated)
```

Where:

**Array LU identifier** matches the device ID displayed in the **powermt display dev=device** output.

## Arguments

**-f** <pseudo device/LU mappings file>

Specifies the file name that contains the device mappings you want to display.

**-p**

Displays the persistent device mappings in the PowerPath database.

You can edit an exported XML-formatted mappings file to establish custom mappings. However, care must be taken to ensure that the LUN XML nodes of the file contain `pseudo_dev`, `product_id` and `logical_devid` elements. In the case of Symmetrix LUs, a `frame_id` element is also required. The imported mappings file must also be constructed using valid XML syntax.

## Pseudo device/LU renaming

```
getfree [pseudos] [-f {i|d|m}] [-n <# of pseudo instances>] [-b <start pseudo device>]
```

Retrieves a list of available PowerPath pseudo devices. For each device, lists the pseudo device name, instance number, major device number, and minor device number. By default, retrieves the first free pseudo device name.

## Arguments

**-f** {i|d|m}

Use the **-f {i|d|m}** argument to indicate the format of the pseudo device:

- **-f i** — Indicates an instance number, for example, *1*.
- **-f d** — Indicates a pseudo device name, for example, *emcpower1*.
- **-f m** — Indicates a minor number, for example, *8*.

---

**Note:** If you do not include the **-f** argument, the **emcpadm** command treats a numeric entry (for example, *2*) as an instance number. The command recognizes an **emcpower*N*** name without the **-f** argument.

---

`-n <# of pseudo instances>`

Retrieves information for the indicated number of free pseudo devices, starting with the first free device.

`-b <start pseudo device>`

Retrieves information for the first available pseudo devices, starting from the indicated device.

If both `-n <# of pseudo instances>` and `-b <start pseudo device>` are specified, retrieves information for up to the indicated number of free pseudo devices, starting at the indicated device.

`getused [pseudos]`

Lists the pseudo device name, instance number, major device number, and minor device number for each used device.

`rename [pseudo] [-f {i|d|m}] -s <src pseudo device> -t <tgt pseudo device>`

Changes the name of the pseudo device from the pseudo device name specified by source to the pseudo device name specified by target. The target pseudo device name must be available (that is, not in use).

## Arguments

`-f {i|d|m}`

Use the **-f {i|d|m}** argument to indicate the format of the pseudo device:

- **-f i** — indicates an instance number, for example, *1*.
- **-f d** — indicates a pseudo device name, for example, *emcpower1*.
- **-f m** — indicates a minor number, for example, *8*.

---

**Note:** If you do not include the **-f** argument, the **emcpadm** command treats a numeric entry (for example, *2*) as an instance number. The command recognizes an **emcpower*N*** name without the **-f** argument.

---

`-s <src pseudo instance> -t <tgt pseudo instance>`

Changes the name of the pseudo device from the pseudo device name specified by **-s** (source) to the pseudo device name specified by **-t** (target). The target pseudo device name must be available (that is, not in use).

You can specify for the device name either the instance number (for example, *12*) or the complete pseudo device name (for example, *emcpower12*).

## Other commands

`help`

Provides help for the **emcpadm** command.

`version`

Identifies the version of the **emcpadm** command and the PowerPath base driver.

## Environment variables

The following environment variables are available for the **emcpadm** command:

`EMCPADM_DBG=<value>`

Setting this environment variable to any value enables debugging. By default this environment variable is not set.

`EMCPADM_ERR=<value>`

Setting this environment variable to any value enables error logging. By default this environment variable is not set.

## Examples

This section includes the following examples:

- ◆ [“Example 1: Manually edit the mappings file” on page 123](#)
- ◆ [“Example 2: Rename devices; export and import device mappings” on page 128](#)
- ◆ [“Example 3: Rename a pseudo device” on page 131](#)
- ◆ [“Example 4: Enable debugging” on page 133](#)

### Example 1: Manually edit the mappings file

When you manually edit the mappings file, you can change multiple pseudo device names simultaneously. This section describes the procedure for exporting, manually editing, and importing device mappings.

---

**Note:** Use the XML format to manually edit the mappings file.

---



**While editing the mappings file, if a conflict occurs between a host’s existing mappings and those in the mappings file, the mappings file content takes precedence.**

---

#### Step 1: Display the current device mappings

To display the PowerPath pseudo device name and logical-device mappings on the node where you run this command, type:

```
emcpadm print_mappings
```

Output similar to the following appears:

```
Device: emcpower0a
 6005 08b4 0001 2ab5 0000 c000 0bc7 0000
 0000 0000 0000 0000 0000 0000 0000 0000
Device: emcpower1a
 6005 08b4 0001 2ab5 0000 c000 0bc1 0000
 0000 0000 0000 0000 0000 0000 0000 0000
Device: emcpower2a
 6005 08b4 0001 2ab5 0000 c000 0bbb 0000
 0000 0000 0000 0000 0000 0000 0000 0000
Device: emcpower3a
```

```
6005 08b4 0001 2ab5 0000 c000 0bb5 0000
0000 0000 0000 0000 0000 0000 0000 0000
```

(Output truncated)

### Compare device mappings with powermt display output

Optionally, view the **powermt display** output for the devices whose mappings you want to change. The logical device ID in the **powermt display** output correlates to the device IDs displayed when you run **emcpadm print\_mappings**.

For example, in the following device mappings excerpt, the device IDs for **emcpwer3a** appear in bold text.

```
Device: emcpower3a
6005 08b4 0001 2ab5 0000 c000 0bb5 0000
0000 0000 0000 0000 0000 0000 0000 0000
```

The characters in bold from the mappings correlate to the logical device ID in the **powermt display dev=3** output.

To view the **powermt display** output, type:

```
powermt display dev=3
```

The logical device ID appears in bold text.

Output similar to the following appears:

```
Pseudo name=emcpower3a
HP HSx ID=50001FE15002F000
Logical device ID=600508B400012AB50000C0000BB50000
state=alive; policy=BasicFailover; queued-IOs=0
Controller: HSV100          Controller A: 66C5APF02H          Controller B: 66C5APF01U
Default owner=Unknown      Current owner=A
=====
----- Host ----- - Stor -          -- I/O Path - -- Stats ---
### HW Path          I/O Paths   Interf.      Mode  State  Q-IOs Errors
=====
3072 pci@1f,4000/lpfc@4/fp@0,0 c2t50001FE15002F008d2s0 A 0      active  alive    0    0
3072 pci@1f,4000/lpfc@4/fp@0,0 c2t50001FE15002F009d2s0 A 1      unlic   alive    0    0
3072 pci@1f,4000/lpfc@4/fp@0,0 c2t50001FE15002F00Cd2s0 B 0      unlic   alive    0    0
3072 pci@1f,4000/lpfc@4/fp@0,0 c2t50001FE15002F00Dd2s0 B 1      active  alive    0    0
3073 pci@1f,2000/lpfc@1/fp@0,0 c3t50001FE15002F008d2s0 A 0      unlic   alive    0    0
3073 pci@1f,2000/lpfc@1/fp@0,0 c3t50001FE15002F009d2s0 A 1      unlic   alive    0    0
3073 pci@1f,2000/lpfc@1/fp@0,0 c3t50001FE15current mapping
```

When you export the currently configured mappings, the device mappings are saved to the file you specify.

To export the mappings shown in [“Step 1: Display the current device mappings” on page 123](#) as XML data, type:

```
emcpadm export_mappings -x -f initial_mappings.xml
```

### Step 3: Manually edit the mappings file

This section describes the contents of the device mappings file in XML format and shows how to edit it.

#### Mappings file format

A device mapping in the mappings file consists of the **pseudo\_dev**, **product\_id**, **frame\_id**, and **logical\_devid** elements.

Where:

- ◆ **pseudo\_dev** = pseudo device name
- ◆ **product\_id** = type pf array
- ◆ **frame\_id** = array identifier
- ◆ **logical\_devid** = LU identifier

This set of elements is used by emcpadm to establish device name to LU mappings during an import mappings process.

---

**Note:** A VNX and CLARiiON logical device ID is a WWN and uniquely identifies a given LU. A Symmetrix device is uniquely identified by an array ID (frame\_id) and a logical device ID.

---

For example, the device mapping information for emcpower0 appears as follows:

```
<?xml version="1.0" encoding="ISO-8859-1"?>
<powerpath_dev_mappings>
<powerpath_version>4.5.0 (build 301)</powerpath_version>
<src_hostname>lclcl169</src_hostname>
<lun>
<pseudo_dev>emcpower0a</pseudo_dev>
<udev>0x1020000000</udev>
<product_id>clariion</product_id>
<logical_devid_type>wwn</logical_devid_type>
<logical_devid>600508b400012ab50000c0000bc70000</logical_devid>
</lun>
...
</powerpath_dev_mappings>
```

### File editing

To swap the device mappings for two existing devices, swap the device **devid** information for the two devices. To swap the mappings for these two devices, swap their devid content.

After editing the file, save it with a different file name, for example, **my\_mappings.xml**. Optionally, you can delete from the file any content that is not changing.

There may be occasions in which you might like to identify a set of array LUs by specific pseudo device names. For example, it might be desirable to associate LU x, LU y, and LU z with emcpower1000, emcpower1001, and emcpower1002, respectively. This can be accomplished by simply editing the pseudo\_dev elements of these devices in an exported mappings file. Alternatively, you can create the XML file directly. For example:

### Before editing:

```
<?xml version="1.0" encoding="ISO-8859-1"?>
<powerpath_dev_mappings>
  <lun>
    <pseudo_dev>emcpower0a</pseudo_dev>
    <product_id>clariion</product_id>
    <logical_devid>600508b400012ab50000c0000bc70000</logical_devid>
  </lun>
  <lun>
    <pseudo_dev>emcpower1a</pseudo_dev>
    <product_id>clariion</product_id>
    <logical_devid>600508b400012ab50000c0000bc10000</logical_devid>
  </lun>
```

```
</powerpath_dev_mappings>
```

#### After editing:

```
<?xml version="1.0" encoding="ISO-8859-1"?>
<powerpath_dev_mappings>
  <lun>
    <pseudo_dev>emcpower1000a</pseudo_dev>
    <product_id>clariion</product_id>
    <logical_devid>600601760B090000A4F4356E2B17D711</logical_devid>
  </lun>
  <lun>
    <pseudo_dev>emcpower1001a</pseudo_dev>
    <product_id>clariion</product_id>
    <logical_devid>600601780B090000874DEDF62A17D711</logical_devid>
  </lun>
</powerpath_dev_mappings>
```

Similarly, you could sort target logical device IDs and assign each an PowerPath pseudo device name by increasing value (for example, emcpower0a, emcpower1a, emcpower2a, etc.).

These mappings will become active following the execution of `emcpadm import_mappings`.

#### Step 4: Compare mappings

This step compares the mappings in the mappings file with those currently configured on the host where you run the command. A list of device mappings that will change when you import the new mappings is displayed.

To compare the mappings in the mappings file with those on the host where you want to import the file mappings, type `emcpadm check_mappings -x -f map1.xml`

Output similar to the following appears:

Imported Device Mapping	Remapping Status
emcpower0a	no change
emcpower1000a	remaps: emcpower1a
emcpower26a	no change
emcpower24a	no change
emcpower21a	no change
emcpower18a	no change
emcpower17a	no change
emcpower16a	no change
emcpower11a	no change
emcpower40a	no change

Where:

**Imported Device Mappings** = The mappings that exist in the specified mappings file.

**Remapping Status** = The devices whose mappings will change when you import the file.

#### Step 5: Import the modified mappings file

You can now import the modified mappings file on the same node or on another node in the cluster.

```
emcpadm import_mappings -x -f my_mappings.xml
```

**Display the new mappings with powermt display**

Optionally, view the effected mappings with the **powermt display** command. To display the output for emcpower0, type:

```
powermt display dev=0
```

After importing the new mappings, the **powermt display** output displays the new logical device ID for emcpower0:

```
Pseudo name=emcpower1000a
HP HSx ID=50001FE15002F000
Logical device ID=600508B400012AB50000C0000BC70000
state=alive; policy=Adaptive; queued-IOs=0
Controller: HSV100          Controller A: 66C5APF02H          Controller B: 66C5APF01U
Default owner=A           Current owner=A
=====
----- Host ----- - Stor - -- I/O Path - -- Stats -----
### HW Path          I/O Paths   Interf.    Mode   State  Q-IOs Errors
=====
3072 pci@1f,4000/lpfc@4/fp@0,0 c2t50001FE15002F008d3s0 A 0 active alive 0 0
3072 pci@1f,4000/lpfc@4/fp@0,0 c2t50001FE15002F009d3s0 A 1 active alive 0 0
3072 pci@1f,4000/lpfc@4/fp@0,0 c2t50001FE15002F00Cd3s0 B 0 active alive 0 0
3072 pci@1f,4000/lpfc@4/fp@0,0 c2t50001FE15002F00Dd3s0 B 1 active alive 0 0
3073 pci@1f,2000/lpfc@1/fp@0,0 c3t50001FE15002F008d3s0 A 0 active alive 0 0
3073 pci@1f,2000/lpfc@1/fp@0,0 c3t50001FE15002F009d3s0 A 1 active alive 0 0
3073 pci@1f,2000/lpfc@1/fp@0,0 c3t50001FE15002F00Cd3s0 B 0 active alive 0 0
3073 pci@1f,2000/lpfc@1/fp@0,0 c3t50001FE15002F00Dd3s0 B 1 active alive 0 0
```

To display the output for **emcpower2**, type:

```
powermt display dev=2
```

Output similar to the following appears:

```
Pseudo name=emcpower1001a
HP HSx ID=50001FE15002F000
Logical device ID=600508B400012AB50000C0000BC10000
state=alive; policy=Adaptive; queued-IOs=0
Controller: HSV100          Controller A: 66C5APF02H          Controller B: 66C5APF01U
Default owner=A           Current owner=A
=====
----- Host ----- - Stor - -- I/O Path - -- Stats ---
### HW Path          I/O Paths   Interf.    Mode   State  Q-IOs Errors
=====
3072 pci@1f,4000/lpfc@4/fp@0,0 c2t50001FE15002F008d5s0 A 0 active alive 0 0
3072 pci@1f,4000/lpfc@4/fp@0,0 c2t50001FE15002F009d5s0 A 1 active alive 0 0
3072 pci@1f,4000/lpfc@4/fp@0,0 c2t50001FE15002F00Cd5s0 B 0 active alive 0 0
3072 pci@1f,4000/lpfc@4/fp@0,0 c2t50001FE15002F00Dd5s0 B 1 active alive 0 0
3073 pci@1f,2000/lpfc@1/fp@0,0 c3t50001FE15002F008d5s0 A 0 active alive 0 0
3073 pci@1f,2000/lpfc@1/fp@0,0 c3t50001FE15002F009d5s0 A 1 active alive 0 0
3073 pci@1f,2000/lpfc@1/fp@0,0 c3t50001FE15002F00Cd5s0 B 0 unlic alive 0 0
3073 pci@1f,2000/lpfc@1/fp@0,0 c3t50001FE15002F00Dd5s0 B 1 unlic alive 0 0
```

## Example 2: Rename devices; export and import device mappings

This example describes how to:

- ◆ Rename pseudo devices with **emcpadm renamepseudo**
- ◆ Export device mappings to a file
- ◆ Import the file mappings on a cluster node

The import feature allows you to remap multiple devices at once and to more easily maintain consistent mappings among cluster nodes or servers.

### Step 1: Display pseudo devices in use

To display the emcpower devices in use, type:

```
emcpadm getused
```

Output similar to the following appears:

```
Pseudo device names in use :
Pseudo Device Name      Instance#   Major#   Minor# *
      emcpower0           0         262      0
      emcpower1           1         262      8
      emcpower2           2         262     16
      emcpower3           3         262     24
      [...]
      emcpower62          62         262    496
      emcpower63          63         262    504
```

\* Minor # is for slice 'a'

### Step 2: Display current mappings

View the current device mappings on the server or node.

**Note:** In the **emcpadm print\_mappings** output, the **Array LU identifier** is the same as the logical device ID in the **powermt display dev=device** output.

Type:

```
emcpadm print_mappings
```

Output similar to the following appears:

```
Device: emcpower63a
Device Id:
  Array LU identifier:
    0000 0000 0554 40e0 0000 0000 0000 0000

[...]

Device: emcpower3a
Device Id:
  Array LU identifier:
    4942 4d20 2020 2020 3231 3035 2020 2020
    2020 2020 2020 2020 3745 3031 3537 3438

Device: emcpower2a
Device Id:
  Array LU identifier:
    4942 4d20 2020 2020 3231 3035 2020 2020
    2020 2020 2020 2020 3744 4331 3537 3438
```



```
Device: emcpower1a
Device Id:
  Array LU identifier:
    4942 4d20 2020 2020 3231 3035 2020 2020
    2020 2020 2020 2020 3745 3131 3537 3438
```

```
Device: emcpower0a
Device Id:
  Array LU identifier:
    4942 4d20 2020 2020 3231 3035 2020 2020
    2020 2020 2020 2020 3744 4431 3537 3438
```

### Step 3: Rename emcpower devices

Prerequisite: The specified target device name must be available (that is, not in use).

To rename the emcpower0 device to emcpower100, type:

```
emcpadm renamepseudo -s emcpower0 -t emcpower100
```

To rename the emcpower1 device to emcpower101, type:

```
emcpadm renamepseudo -s emcpower1 -t emcpower101
```

To rename the emcpower2 device to emcpower0, type:

```
emcpadm renamepseudo -s emcpower2 -t emcpower0
```

### Step 4: Export the new mappings

Export the new device mappings. To export the device mappings to a file called **mappings.exp**, and save the file in /tmp, type:

```
emcpadm export_mappings -f /tmp/mappings.exp
```

### Step 5: Display new device mappings on node A

After renaming the pseudo devices and exporting the mappings, verify the updated mappings. Type:

```
emcpadm print_mappings
```

Output similar to the following appears:

```
PowerPath device name, logical device id mappings
-----
```

```
Device: emcpower63a
Device Id:
  Array LU identifier:
    0000 0000 0554 40e0 0000 0000 0000 0000
```

[...]

```
Device: emcpower3a
Device Id:
  Array LU identifier:
    4942 4d20 2020 2020 3231 3035 2020 2020
    2020 2020 2020 2020 3745 3031 3537 3438
```

```
Device: emcpower101a
Device Id:
  Array LU identifier:
    4942 4d20 2020 2020 3231 3035 2020 2020
    2020 2020 2020 2020 3745 3131 3537 3438
```

```
Device: emcpower100a
Device Id:
  Array LU identifier:
    4942 4d20 2020 2020 3231 3035 2020 2020
    2020 2020 2020 2020 3744 4431 3537 3438
```

## Step 6: Preview mapping changes on cluster node B

The mappings in **mappings.exp** can be transferred to another node or server in the cluster. Before importing **mappings.exp**, preview the device remappings on node B.

**Note:** In this example, the current mappings on node B are the same as the original mappings on node A (shown in [“Step 2: Display current mappings” on page 128](#)).

To preview device remapping on node B, type:

```
emcpadm check_mappings -f /tmp/mappings.exp
```

Output similar to the following appears:

Imported Device Mapping	Remapping Status
emcpower63a	no change
[...]	
emcpower3a	no change
emcpower0a	remaps: emcpower2a emcpower0a
emcpower101a	remaps: emcpower1a
emcpower100a	remaps: emcpower0a

Where:

**Imported Device Mappings** = The device mappings in the specified file (**mappings.exp** in this example).

**Remapping Status** = The device mapping changes that occur when you import the mappings file. For example, **emcpower1a** would be remapped to **emcpower100a**.

## Step 7: Import the new mappings on node B

After previewing the device remappings, import the **mappings.exp** file on node B.

To import the mappings in **mappings.exp**, type:

```
emcpadm import_mappings -f /tmp/mappings.exp
```

## Step 8: Verify the new mappings on node B

To verify the devices were successfully remapped, compare the node B mappings with those in **mappings.exp**. Type:

```
emcpadm check_mappings -f /tmp/mappings.exp
```

Output similar to the following appears:

Imported Device Mapping	Remapping Status
emcpower63a	no change
emcpower53a	no change

```
[...]
```

```
emcpower0a      no change
emcpower101a   no change
emcpower100a   no change
```

The output verifies that the node B mappings are the same as the mappings in **mappings.exp**.

### Example 3: Rename a pseudo device

In the following example, PowerPath pseudo devices on a Linux host are renamed using **emcpadm renamepseudo**.

1. List the PowerPath pseudo device names in use:

```
emcpadm getused
```

```
Pseudo device names in use:
```

Pseudo Device Name	Major#	Minor#
emcpowera	232	0
emcpowerb	232	16
emcpowerc	232	32
emcpowerd	232	48
emcpowere	232	64
emcpowerf	232	80
emcpowerg	232	96
emcpowerh	232	112
emcpoweri	232	128
emcpowerj	232	144
emcpowerk	232	160
emcpowerl	232	176
emcpowerm	232	192
emcpowern	232	208
emcpowero	232	224
emcpowerp	232	240

2. List the next 10 free PowerPath pseudo device names, beginning at the device named *emcpowerk*:

```
emcpadm getfree -n 10 -b emcpowerk
```

```
PowerPath pseudo device names not in use:
```

Pseudo Device Name	Major#	Minor#
emcpowerq	233	0
emcpowerr	233	16
emcpowers	233	32
emcpowert	233	48
emcpoweru	233	64
emcpowerv	233	80
emcpowerw	233	96
emcpowerx	233	112
emcpowery	233	128
emcpowerz	233	144

3. Verify that *emcpowera* exists.

```
powermt display dev=emcpowera
```

```
Pseudo name=emcpowera
CLARiION ID=WRE00100100415
Logical device ID=6006016018C80800B5069570B8A9D911
state=alive; policy=CLAROpt; queued-IOS=0
```

```

Owner: default=SP A, current=SP A
=====
----- Host ----- - Stor - -- I/O Path - -- Stats ---
### HW Path          I/O Paths  Interf.  Mode   State  Q-IOs Errors
=====
  3 lpfc              sdaa     SP A0    active alive   0      0
  2 lpfc              sdj      SP A1    active alive   0      0

```

- Rename the *emcpowera* device to *emcpowerq*:

```
emcpadm renamepseudo -s emcpowera -t emcpowerq
```

- List the used pseudo devices names. Note that *emcpowera* is no longer included in the output, and that *emcpowerq* is included.

```
emcpadm getused
```

PowerPath pseudo device names in use:

Pseudo Device Name	Major#	Minor#
emcpowereb	232	16
emcpowerc	232	32
emcpowerd	232	48
emcpowere	232	64
emcpowerf	232	80
emcpowerg	232	96
emcpowerh	232	112
emcpoweri	232	128
emcpowerj	232	144
emcpowerk	232	160
emcpowerl	232	176
emcpowerm	232	192
emcpowern	232	208
emcpowero	232	224
emcpowerp	232	240
emcpowerq	233	0

- Verify that *emcpowera* does not exist:

```
powermt display dev=emcpowera
```

Bad dev value emcpowera, or not under PowerPath control.

- Verify that *emcpowerq* exists:

```
powermt display dev=emcpowerq
```

```

Pseudo name=emcpowerq
CLARiiON ID=WRE00100100415
Logical device ID=6006016018C80800B5069570B8A9D911
state=alive; policy=CLAROpt; queued-IOs=0
Owner: default=SP A, current=SP A
=====
----- Host ----- - Stor - -- I/O Path - -- Stats ---
### HW Path          I/O Paths  Interf.  Mode   State  Q-IOs Errors
=====
  3 lpfc              sdaa     SP A0    active alive   0      0
  2 lpfc              sdj      SP A1    active alive   0      0

```

- Update the **powermt.custom** file with the new pseudo device mappings:

```
powermt save
```

## Example 4: Enable debugging

With debugging enabled, you can view additional pseudo device information. To enable debugging and display the current mappings, type:

```
EMCPADM_DBG=1 emcpadm print_mappings
```

Output similar to the following appears:

```
Device: emcpower0a
6005 08b4 0001 2ab5 0000 c000 0bc7 0000
0000 0000 0000 0000 0000 0000 0000 0000
----- DEBUG -----
userDev=0x10200000000
LAM class=Mpap
instance=14
configured=1
hidden=0
installed=1
registered=1
----- DEBUG -----
Device: emcpower1a
6005 08b4 0001 2ab5 0000 c000 0bc1 0000
0000 0000 0000 0000 0000 0000 0000 0000
----- DEBUG -----
userDev=0x10200000008
LAM class=Mpap
instance=16
configured=1
hidden=0
installed=1
registered=1
----- DEBUG -----
```

(Output truncated)

---

**Note:** In general a reboot is not necessary if the remapping does not include a boot device. In this case the remappings can occur on a running, quiesced system. Clearly one does not want to attempt such remappings with running applications.

---

- Caveats**
- ◆ Tools that can modify the **emcp.conf** file (for example, `powercfg`) cannot be run until the remapping process has completed.
  - ◆ The user-supplied data is assumed to be correct. Only the source device IDs can be verified.

**References**    PowerPath `emcpadm` (1) man page

## emcpcfg script: Save and restore a PowerPath configuration

**Platforms** This command is supported on Solaris platforms.

**Storage systems** This command is supported on all storage systems.

**Syntax**

```
emcpcfg backup
emcpcfg list_cfgs
emcpcfg restore [-reboot] <saved_PowerPath_configuration>
emcpcfg set_cfgs_limit <number>
emcpcfg cleanup
```

**Description** The **emcpcfg** script saves and restores a PowerPath configuration if any of the configuration files have been corrupted or removed. The files that represent the current configuration are:

- ◆ /kernel/drv/emcp.conf
- ◆ powermt\_custom.xml
- ◆ /etc/powermt.custom
- ◆ /etc/emcp\_devicesDB.dat
- ◆ /etc/emcp\_devicesDB.idx

### Arguments

**backup**  
Saves the current configuration files for PowerPath as TAR files in the /etc/emc/cfgs directory. **powermt save** is automatically run when **emcpcfg backup** is run.

**list\_cfgs**  
Lists the available saved configurations.

**restore [-reboot] <saved\_PowerPath\_configuration>**  
Restores the specified last good configuration that was saved. After running **emcpcfg restore**, reboot the host to return to the last saved configuration.

---

**Note:** The **restore [-reboot]** option performs an automatic reboot; otherwise, a manual reboot is required.

---

**set\_cfgs\_limit <number>**  
Sets the limits on number of configurations that can be backed up. For example, if the new limit is set to three configurations, as opposed to a previous limit of five configurations, then the two older configurations are removed to maintain only three backups. You do not need to delete the configurations manually.

**cleanup**  
Removes/deletes all of the backed up configuration files (which were backed up using PowerPath **emcpcfg** script only).

# CHAPTER 2

## PowerPath Migration Enabler CLI

This chapter describes the PowerPath Migration Enabler command line utilities.

- ◆ Migration Enabler new and changed commands..... 136
- ◆ Migration Enabler technology type support ..... 138
- ◆ Migration Enabler syntax summary..... 139
- ◆ powermigcl: Configure Migration Enabler cluster resources ..... 168

## Migration Enabler new and changed commands

Table 24 on page 136 lists the new and changed PowerPath Migration Enabler commands added in versions 5.0 and later, with the earliest supported version listed by platform. Support in later versions (including service packs) can be assumed if no exception in the table notes otherwise. Version 5.0 is the first Migration Enabler release.

**Table 24 Migration Enabler new and changed commands (page 1 of 2)**

Command	Earliest supported version				
	AIX	HP-UX	Linux	Solaris	Windows
powermig abort	5.0	5.0	5.0	5.0	5.0
Support for -all	5.7	5.2	5.7	5.5	5.7
powermig cleanup	5.0	5.0	5.0	5.0	5.0
Support for -all	5.7	5.2	5.7	5.5	5.7
powermig commit	5.0	5.0	5.0	5.0	5.0
Support for -all	5.7	5.2	5.7	5.5	5.7
powermig getHandle	5.0	5.0	5.0	5.0	5.0
powermig help	5.0	5.0	5.0	5.0	5.0
powermig info	5.0	5.0	5.0	5.0	5.0
powermig options	5.5	5.2	5.5	5.5	5.7
Support for hostcopy_ceiling	5.5	5.2	5.5	5.5	5.7
powermig pause	5.3	5.1	5.3	5.3	5.3
Support for -all	5.7	5.2	5.7	5.5	5.7
powermig query	5.0	5.0	5.0	5.0	5.0
Support for -all	5.7	5.2	5.7	5.5	5.7
powermig recover	5.0	5.0	5.0	5.0	5.0
Support for -all	5.7	5.2	5.7	5.5	5.7
powermig resume	5.3	5.1	5.3	5.3	5.3
Support for -all	5.7	5.2	5.7	5.5	5.7
powermig selectSource	5.0	5.0	5.0	5.0	5.0
Support for -all	5.7	5.2	5.7	5.5	5.7
powermig selectTarget	5.0	5.0	5.0	5.0	5.0
Support for -all	5.7	5.2	5.7	5.5	5.7
powermig setup	5.0	5.0	5.0	5.0	5.0
Support for -cluster	N/A <sup>a</sup>	N/A	N/A	N/A	5.7 SP1 <sup>b</sup>
Removal of -suspendTime	5.5	5.2	5.5	5.3	5.5
Support for -techHost and -techPort	5.3	5.1 SP2	5.3	5.3	5.3



**Table 24 Migration Enabler new and changed commands (page 2 of 2)**

Command	Earliest supported version				
	AIX	HP-UX	Linux	Solaris	Windows
Support for -file	5.7	5.2	5.7	5.5	5.7
powermig sync	5.0	5.0	5.0	5.0	5.0
Support for -all	5.7	5.2	5.7	5.5	5.7
powermig throttle	5.3	5.0	5.3	5.0	5.0
Removal of -suspendTime	5.5	5.2	5.5	5.3	5.5
Support for -all	5.7	5.2	5.7	5.5	5.7
powermig undoRedirect	5.0	5.0	5.0	5.0	5.0
Support for -all	5.7	5.2	5.7	5.5	5.7
powermig version	5.0	5.0	5.0	5.0	5.0
powermigcl	N/A	N/A	N/A	N/A	5.7 SP1 <sup>c</sup>

- a. N/A means the command or utility is not supported on the platform.
- b. The -cluster argument is not supported on Windows Server 2003.
- c. This command is not supported when migrating devices on Windows Server 2003.

## Migration Enabler technology type support

Table 25 on page 138 lists for all versions of PowerPath Migration Enabler the commands supported with each platform and technology type.

**Table 25** Migration Enabler technology type support

Commands	Supported technology types				
	Version 5.2	Version 5.3	Version 5.5	Version 5.6	Version 5.7
	HP-UX, Solaris, Windows	AIX, Linux, Solaris, Windows	AIX, Linux, Solaris, Windows	Linux	AIX, Linux, Windows
powermig abort	HOSTCOPY, INVE, OR, SYMCLONE <sup>a</sup>	HOSTCOPY, INVE <sup>b</sup> , OR, SYMCLONE	HOSTCOPY, INVE <sup>b</sup> , OR, SYMCLONE	HOSTCOPY, INVE <sup>b</sup> , OR, SYMCLONE	HOSTCOPY, INVE <sup>b</sup> , OR, SYMCLONE
powermig cleanup					
powermig commit					
powermig getHandle					
powermig help					
powermig info					
powermig options	HOSTCOPY (HP-UX only)	N/A <sup>c</sup>	HOSTCOPY (AIX, Linux, Solaris only)	HOSTCOPY	HOSTCOPY
powermig pause	HOSTCOPY, OR, SYMCLONE <sup>a,d</sup>	HOSTCOPY, OR, SYMCLONE <sup>d</sup>	HOSTCOPY, OR, SYMCLONE <sup>d</sup>	HOSTCOPY, OR, SYMCLONE <sup>d</sup>	HOSTCOPY, OR, SYMCLONE <sup>d</sup>
powermig query	HOSTCOPY, INVE, OR, SYMCLONE <sup>a</sup>	HOSTCOPY, INVE, OR, SYMCLONE	HOSTCOPY, INVE, OR, SYMCLONE	HOSTCOPY, INVE, OR, SYMCLONE	HOSTCOPY, INVE, OR, SYMCLONE
powermig recover					
powermig resume	HOSTCOPY, OR, SYMCLONE <sup>a,d</sup>	HOSTCOPY, OR, SYMCLONE <sup>d</sup>	HOSTCOPY, OR, SYMCLONE <sup>d</sup>	HOSTCOPY, OR, SYMCLONE <sup>d</sup>	HOSTCOPY, OR, SYMCLONE <sup>d</sup>
powermig selectSource	HOSTCOPY, INVE, OR, SYMCLONE <sup>a</sup>	HOSTCOPY, INVE, OR, SYMCLONE	HOSTCOPY, INVE, OR, SYMCLONE	HOSTCOPY, INVE, OR, SYMCLONE	HOSTCOPY, INVE, OR, SYMCLONE
powermig selectTarget					
powermig setup					
powermig sync					
powermig throttle	HOSTCOPY, OR, SYMCLONE <sup>a</sup>	HOSTCOPY, OR, SYMCLONE	HOSTCOPY, OR, SYMCLONE	HOSTCOPY, OR, SYMCLONE	HOSTCOPY, OR, SYMCLONE
powermig version	HOSTCOPY, INVE, OR, SYMCLONE <sup>a</sup>	HOSTCOPY, INVE, OR, SYMCLONE	HOSTCOPY, INVE, OR, SYMCLONE	HOSTCOPY, INVE, OR, SYMCLONE	HOSTCOPY, INVE, OR, SYMCLONE

a. Migration Enabler 5.2 for Solaris and Windows does not support SYMCLONE. Supported only on Migration Enabler 5.2 HP-UX.

b. Encapsulation (INVE) is supported only with Invista storage systems. It is not supported for EMC VPLEX storage systems.

c. N/A means the command or utility is not supported on the platform.

d. TimeFinder/Clone supports pause and resume only when the source and target are the same size.

## Migration Enabler syntax summary

A summary of the syntax of all **powermig** commands follows:

```
powermig abort -handle <migrationHandle>|-all [-noPrompt|-no]
powermig cleanup -handle <migrationHandle>|-all
    [-format] [-force] [-noPrompt|-no]
powermig commit -handle <migrationHandle>|-all [-noPrompt|-no]
powermig getHandle -dev <device>|-src <source> -tgt <target>
powermig help <operation> <options>
powermig info -handle <migrationHandle> [-query]|-all [-query]
powermig options [-<option> <value>]
powermig pause -handle <migrationHandle>|-all [-noPrompt|-no]
powermig query -handle <migrationHandle>|-all [-noPrompt|-no]
powermig recover -handle <migrationHandle>|-all [-noPrompt|-no]
powermig resume -handle <migrationHandle>|-all [-noPrompt|-no]
powermig selectSource -handle <migrationHandle>|-all [-noPrompt|-no]
powermig selectTarget -handle <migrationHandle>|-all [-noPrompt|-no]
powermig setup -techType <techType> -src <sourceName> -tgt
    <targetName>|-file <fileName> [-throttleValue <throttleValue>]
    [-cluster] [-techHost <hostname> [-techPort <port#>]] [-nothin]
    [-force] [-noprompt|-no]
powermig sync -handle <migrationHandle>|-all [-noPrompt|-no]
powermig throttle -handle <migrationHandle>|-all -throttleValue
    <throttleValue> [-noPrompt|-no]
powermig undoRedirect -handle <migrationHandle>|-all [-force]
    [-noPrompt|-no]
powermig version
```

## Command syntax

The **powermig** command keywords such as **selectTarget**, and arguments such as **-handle**, are not case-sensitive. The capitalization used in this guide is for clarity only. You can type command keywords and arguments using all lowercase letters or mixed lower and uppercase letters. You can also run the arguments in any order. This guide lists arguments in alphabetical order.

For example, the following command syntax is valid:

```
powermig setup -techType or -src <devicename> -tgt <devicename>
```

If you type an unsupported command option, Migration Enabler does not display an error; it ignores the unsupported option and executes the valid command and options.

## Argument abbreviations

The **powermig** command supports the abbreviations shown in [Table 26 on page 141](#).

**Table 26** Argument abbreviations

Argument	Abbreviation
-device	-dev
-handle	-hd
-force	-forc
-format	-form
-help	-h
-noPrompt	-no
-noThin	-noth
-query	-q
-source	-src
-target	-tgt
-techHost	-th
-techPort	-tp
-techType	-tt
-throttleValue	-tv

## powermig abort

**Purpose** Aborts the migration.

**Syntax** `powermig abort -handle <migrationHandle>|-all [-noPrompt|-no]`

**Description** **powermig abort** aborts the migration and returns the migration session to the **setup** state. From **Setup**, you can either restart the synchronization with the **powermig sync** command, or clean up the migration with the **powermig cleanup** command. When you clean up the migration, you remove selected data copied from the source to the target.

---

**Note:** You can abort a migration session at any point after running **powermig sync** and before running **powermig commit**. Once a migration has been committed, it cannot be aborted.

---

After aborting a migration, you can restart synchronization of the source and target or you can clean up the target to remove selected data that was copied from the source.

**powermig abort** fails if the handle is invalid or migration is not in a state that permits an abort.

**Arguments**

`-handle <migrationHandle>`  
Number that identifies the migration session. It is assigned when you run the **powermig setup** command. You must enter the handle when running commands that reference the migration session.

`-all`  
Performs abort for all migration sessions in the **syncing**, **sourceSelected**, or **targetSelected** state. Migrations in other states are ignored. Migrations in the **needsRecovery** state that cannot be recovered with the **powermig recover** command must be aborted individually by handle.

`-noPrompt|-no`  
Suppresses the default confirmation that appears when you run a command.

**Migration state** Required migration state to run the command: **syncing**, **sourceSelected**, or **targetSelected**  
Migration state after the command completes: **setup**

**Example** `powermig abort -handle 1powermig cleanup`

**Purpose** Cleans up data on the source or target logical unit.

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

**Description** **powermig cleanup** cleans up data on the source or target logical unit, depending on the state of the migration when you run the command. This cleanup prevents the application or operating system from potentially seeing two logical units with identical data, which could cause confusion.

- ◆ Run this command from the **committed** state to clean up selected data on the source logical unit.

- ◆ Run this command from the **setup** state to clean up selected data copied from the source to the target logical unit. For example, you may want to destroy an aborted migration or a migration that has been setup, but is no longer desired.

Once this command completes successfully, the migration handle is no longer valid, and no record of the migration exists.

## Arguments

**-handle** <migrationHandle>

Number that identifies the migration session. It is assigned when you run the **powermig setup** command. You must type the handle when running commands that reference the migration session.

**-all**

Performs cleanup for all migration sessions in the **committed** or **setup** state. Migrations in other states are ignored.

---

**Note:** Migrations in the **needsRecovery** state requiring the **-force** argument must be cleaned up individually by handle.

---

**-force**

Removes the migration session from PowerPath Migration Enabler control. Use this argument when the following are true:

- **powermig cleanup** generates an error because the logical unit requiring cleanup is permanently inaccessible.
- you are sure you want to erase all knowledge of the migration from PowerPath Migration Enabler.

**-format**



This command option destroys all data on the source logical unit leaving it labeled and ready for use.

---



---

**Note:** This option is supported only on Linux and Solaris hosts and is not supported with Virtual Encapsulation (INVE) or with encrypted devices. This option will cause **powermig cleanup** to take substantially longer because it performs a full format of the disk.

---

Used for migrations in the **committed** state when a full disk format of the source logical unit is desired. This option has no effect on migrations in the **setup** state.

To shorten this argument, use **-form**.

**-noPrompt** | **-no**

Suppresses the default confirmation that appears when you run a command.

## Migration state

Required migration state to run the command: **committed** (cleans up source) or **setup** (cleans up target).

Migration state after the command completes: Not applicable. No record of the migration session remains.

**Example** `powermig cleanup -handle 1`

Output:

Cleanup migration? [Yes]/No:



## powermig commit

**Purpose** Commits the migration.

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

**Description** **powermig commit** commits the migration by permanently designating the target as the recipient of all I/O requests. After this command is run, Migration Enabler no longer keeps the source and target synchronized and I/O is not sent to the source. The migration session must be in the **targetSelected** state for this command to succeed.

When you commit a migration, the migration enters one of two states, depending on whether the source logical unit is a pseudo- or native-named device. If the source is a pseudo-named device, the migration enters the **committed** state. If the source is a native-named device, the migration enters the **committedAndRedirected** state.

The **powermig commit** command fails if the handle is not valid or the migration is not in the **targetSelected** state.

**Arguments**

- `-handle <migrationHandle>`  
Number that identifies the migration session. It is assigned when you run the **powermig setup** command. You must type the handle when running commands that reference the migration session.
- `-all`  
Performs commit for all migration sessions in the **targetSelected** state. Migrations in other states are ignored.
- `-noPrompt|-no`  
Suppresses the default confirmation that appears when you run a command.

**Migration state** Required migration state to run the command: **targetSelected**

Migration state after the command completes: **committed** (for a pseudo source device) or **committedAndRedirected** (for a native source device).

**Example** `powermig commit -handle 1`

No output appears when this command completes successfully.

## powermig getHandle

**Purpose** Displays the handle for a migration session in progress.

**Syntax** `powermig getHandle -dev <device>|-src <source> -tgt <target>`

**Description** **powermig getHandle** displays the handle for a migration session in progress. The handle is assigned when you run the **powermig setup** command, and it is needed for other **powermig** commands that reference the migration session.

### Arguments

`-dev <device>`

The source or target device name for the migration whose handle you want to retrieve. The device can be a pseudo- or native-named device. The device name must match the name entered when you set up the migration session.

`-src <source>`

The source-device name involved in the migration whose handle you want to retrieve. The **-source** argument requires that you also enter the target-device name as an argument in the **-target** flag. Alternatively, use the **-dev** argument with either the source- or target-device name as an easier way to retrieve the handle.

`-tgt <target>`

The target-device name involved in the migration whose handle you want to retrieve. The **-target** argument requires that you also enter the source-device name as an argument in the **-source** flag. Alternatively, use the **-dev** argument with either the source- or target-device name as an easier way to retrieve the handle.

**Migration state** Required migration state to run the command: Any state.

Migration state after the command completes: This command does not change the migration state.

If a migration with the specified device name or names does not exist, you see this message: **Handle not found**

### Example

```
powermig getHandle -dev emcpower72a
```

Output:

```
Migration Handle = 1
```

# powermig help

**Purpose** Displays a brief description and usage summary for the specified **powermig** command.

**Syntax** `powermig help <operation> <options>`

**Description** **powermig help** displays a brief description and usage summary for the specified **powermig** command. By default, supported commands are listed. Specify a specific command (for example, **powermig help setup**) to display help for that command.

**Arguments** `<operation>`  
Name of the **powermig** operation about which you want to view information. The possible values are:

abort	resume
cleanup	selectSource
commit	selectTarget
getHandle	setup
info	sync
pause	throttle
query	undoRedirect
recover	version

`<options>`  
Name of the argument you want to view help. The possible values are:

-handle	-source
-format	-target
-force	-techType
-noPrompt	-throttleValue
-query	

**Migration state** Required migration state to run the command: Any state.

Migration state after the command completes: This command does not change the migration state.

## powermig info

**Purpose** Displays information about a specific migration session or about all active migrations.

**Syntax** `powermig info -handle <migrationHandle> [-query] | -all [-query]`

**Description** **powermig info** displays information about a specific migration session or about all active migrations. Output includes:

- ◆ source and target device names for the migration session
- ◆ underlying technology used in the migration
- ◆ state of the migration
- ◆ logical-unit fault information, where relevant

The synchronization status appears if you use the **-query** option.

### Arguments

`-handle <migrationHandle>`

Number that identifies the migration session. It is assigned when you enter the **powermig setup** command. You must enter the handle when running commands that reference the migration session.

`-query`

Queries all migrations for their synchronization status. This is a resource-intensive operation if a large number of migrations are in progress and in the **syncing** state. The abbreviated form of this argument is **-q**.

`-all`

Specifies that information be displayed for all PowerPath Migration Enabler migrations in progress.

### Migration state

Required migration state to run the command: Any state.

Migration state after the command completes: This command does not change the migration state.

A **no migrations found** message displays when no migrations are present on the host.

### Example

`powermig info -handle 14`

The output shows a Host Copy migration in the **syncing** state:

```
=====
Hnd  Source      Target      Tech      State
===  =====
14   emcpower48a  emcpower45a  HostCopy  syncing (12%) [PAUSED]
```

This output is applicable only for migrations in an MSCS cluster environment. The **powermig info** command here displays cluster and non-cluster migrations initiated from the migration node. It does not display cluster migrations initiated from other nodes. You will be able to differentiate between cluster and non-cluster migration by **(cl)** flag near the techType value.

```
=====
Hnd  Source      Target      Tech      State
===  =====
```

```
1 harddisk18 harddisk19 HostCopy(cl) setup
2 harddisk16 harddisk17 HostCopy      setup
```

## powermig options

**Purpose** Set PowerPath Migration Enabler options which apply to all migrations, or all migrations within a specific technology type.

**Syntax** `powermig options [-<option> <value>]`

**Description** The **powermig options** command lets you set PowerPath Migration Enabler options which apply to all migrations, or all migrations within a specific technology type. Running **powermig options** with no option merely displays all available options and their values.

---

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

---

### Arguments

`<option>`

Name of the option for which you want to view. The possible value is: `hostcopy_ceiling`

`<value>`

Name of the value for which you want to view. The possible values are: *CeilingValue*, `none`.

`-hostcopy_ceiling <ceilingValue>|none`

Used to control the aggregate copy rate for all host copy migrations in the **syncing** state. Once set, all per migration throttle values are ignored.

`<ceilingValue>`

Specifies Host Copy ceiling value in megabytes (MB) per second. This integer value can range from 1 to 1000000 (1,000,000) MB per second.

`none`

Disables Host Copy ceiling. When set to **none**, per migration throttling and the active sessions limit are enabled.

### Migration state

Required migration state to run the command: Any state.

Migration state after the command completes: This command does not change the migration state.

---

**Note:** This command changes the Host Copy ceiling value, and also switches between the traditional throttle based behavior and the new Host Copy ceiling behavior. This command affects all the current and new Host Copy migrations.

---

### Example

```
powermig options -hostcopy_ceiling 10
```

The output shows that the Host Copy ceiling value set is 10 MB per second.

```
hostcopy_ceiling: 10 Megabytes per second
```

## powermig pause

**Purpose** Pauses a migration session.

**Syntax** `powermig pause -handle <migrationHandle>|-all [-noPrompt|-no]`

**Description** **Note:** This command is supported with all technologies except Virtual Encapsulation (INVE).

**powermig pause** suspends a migration session in the **syncing** state and pauses the synchronization of the source and target logical units. In the case of a TimeFinder/Clone migration, **powermig pause** is supported only if the source and target are the same size.

Pausing a migration allows you to free host resources for other operations. When you resume the migration, the synchronization process picks up where it left off.

**Arguments**

- `-handle <migrationHandle>`  
Number that identifies the migration session. It is assigned when you enter the **powermig setup** command. You must enter the handle when running commands that reference the migration session.
- `-all`  
Performs pause for all migration sessions in the **syncing** state. Migrations in other states are ignored.
- `-noPrompt|-no`  
Suppresses the default confirmation that appears when you run a command.

**Migration state** Required migration state to run the command: **syncing**  
Migration state after the command completes: **syncing [PAUSED]**

**Example** `powermig pause -handle 1`  
No output appears when this command completes successfully.

## powermig query

**Purpose** Displays the state of a migration.

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

**Description** **powermig query** displays the state of a migration.

**Arguments**

- `-handle <migrationHandle>`  
Number that identifies the migration session. It is assigned when you enter the **powermig setup** command. You must enter the handle when running commands that reference the migration session.
- `-all`  
Performs query for all migration sessions.
- `-noPrompt|-no`  
Suppresses the default confirmation that appears when you run a command.

**Migration state** Required migration state to run the command: Any state.

Migration state after the command completes: This command does not change the migration state.

**Output field descriptions**

---

**Note:** The output fields that display vary based on the migration state.

---

**Handle**  
Number that identifies the migration session.

**Source**  
Source logical unit. Includes the LUN size for technologies that sync and **thin** if the source is virtually provisioned. LUN size is not displayed for Virtual Encapsulation (INVE) migrations. Once committed, if the devices have been swapped, the size information will also be swapped.

**Target**  
Target logical unit. Includes the LUN size for technologies that sync and **thin** if the target is virtually provisioned. LUN size is not displayed for Virtual Encapsulation (INVE) migrations. Once committed, if the devices have been swapped, the size information will also be swapped.

**Technology**  
Underlying technology used in the migration.

**Migration state**  
State of the migration.

**Percent InSync**  
Percentage the source and target logical units are in-sync and displays the amount of data copied.

**Throttle Value**  
Value of the throttle, which slows down or speeds up the synchronization of the source and target logical units.

**Sync Start Time**  
Date and time that the migration begins syncing. This begins after running **powermig sync**.



**Sync End Time**

Date and time that the migration completes syncing. This occurs after the last **powermig query** runs and moves the migration to the **sourceSelected** state.

---

**Note:** For Host Copy only, the last **powermig query** occurs automatically at the final checkpoint when the copy is complete.

---

**Total Sync Time**

Elapsed time from the start of syncing until the last **powermig query** during the **syncing** state runs. This includes time that the host was down, but not time the migration was paused.

**Recent Throughput**

Throughput and the amount of data copied since **powermig query** was last run.

**Overall Average Throughput**

Determined by dividing the source logical unit size by the total sync time. This field displays after the migration has finished syncing.

**Estimated Time to Completion**

Estimated sync time remaining using the most recent throughput calculation.

**Estimated Completion Time**

Estimated date and time when sync should complete.

**Example**

```
powermig query -hd 92
```

The following output shows a Host Copy migration in the **syncing** state:

```
Handle: 92
Source: emcpower1a (512.01 MB)
Target: emcpower2a (4.21 GB, thin)
Technology: HostCopy
Migration state: syncing
Throttle Value: 2
Percent InSync: 46% (237.50 MB copied)
Sync Start Time: Mon Aug 1 11:09:21 2011
Recent Throughput: 14.84 MB/s (118.75 MB in 8 seconds)
Estimated Time to Completion (using recent throughput): 18 seconds
Estimated Completion Time: Mon Aug 1 11:09:53 2011
```

## powermig recover

**Purpose** Recover the migration command that was in progress.

**Syntax** `powermig recover -handle <migrationHandle>|-all [-noPrompt|-no]`

**Description** **powermig recover** tries to recover the migration command that was in progress when an interruption occurred due to a migration error or process crash. Run this command whenever the migration is in the **needsRecovery** state.

If a migration error occurs, the recovery may fail until the cause of the error is identified and resolved.

### Arguments

`-handle <migrationHandle>`

Number that identifies the migration session. It is assigned when you enter the **powermig setup** command. You must enter the handle when running commands that reference the migration session.

`-all`

Performs recover for all migration sessions in the **needsRecovery** state. Migrations in other states are ignored.

`-noPrompt|-no`

Suppresses the default confirmation that appears when you run a command.

### Migration State

Required migration state to run the command: **needsRecovery**

Migration state after the command completes: The next state in the migration process. For example, if the transition from **syncing** to **sourceSelected** is interrupted, after running **powermig recover** the migration session enters the **sourceSelected** state.

### Example

`powermig recover -handle 3`

No output appears when this command completes successfully.

## powermig resume

**Purpose** Restarts a migration that was paused.

**Syntax** `powermig resume -handle <migrationHandle>|-all [-noPrompt|-no]`

**Description** **Note:** This command is supported with all technology types except Virtual Encapsulation (INVE).

**powermig resume** restarts a migration that was paused by the **powermig pause** command. In the case of a TimeFinder/Clone migration, the **powermig resume** command works only if the source and target are the same size.

Use the **powermig info -all -query** command to determine which migrations are paused. A migration state of **PAUSED** indicates a paused migration session.

**Arguments**

- `-handle <migrationHandle>`  
Number that identifies the migration session. It is assigned when you enter the **powermig setup** command. You must enter the handle when running commands that reference the migration session.
- `-all`  
Performs resume for all migration sessions in the **paused** state. Migrations in other states are ignored.
- `-noPrompt|-no`  
Suppresses the default confirmation that appears when you run a command.

**Migration state** Required migration state to run the command: **paused**  
Migration state after the command completes: **syncing**

**Example** `powermig resume -handle 1`  
No output appears when this command completes successfully.

## powermig selectSource

**Purpose** Designates the source logical unit as the recipient of all I/O requests.

**Syntax** `powermig selectSource -handle <migrationHandle>|-all [-noPrompt|-no]`

**Description** **powermig selectSource** designates the source logical unit as the recipient of all I/O requests. When this command completes successfully, the migration is in the source and target logical units continue to be synchronized.

In the **sourceSelected** state, the This command allows you to back out of the **targetSelected** state if necessary. Synchronization of the source and target continues after running this command.

The **powermig selectSource** command fails if the handle is invalid or if the migration is not in the **targetSelected** state or the **sourceSelected** state.

**Arguments**

- `-handle <migrationHandle>`  
Number that identifies the migration session. It is assigned when you enter the **powermig setup** command. You must enter the handle when running commands that reference the migration session.
- `-all`  
Performs selectSource for all migration sessions in the **targetSelected** state. Migrations in other states are ignored.
- `-noPrompt|-no`  
Suppresses the default confirmation that appears when you run a command.

**Migration state** Required migration state to run the command: **targetSelected**

Migration state after the command completes: **sourceSelected**

You cannot commit a migration from this step. To commit a migration, select the target first and then commit the migration.

**Example** `powermig selectSource -handle 1`

Output:

```
Transition to sourceSelected state? [Yes]/No:
```

## powermig selectTarget

**Purpose** Designates the target logical unit as the recipient of all I/O requests.

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

**Description** **powermig selectTarget** designates the target logical unit as the recipient of all I/O requests. When this command completes successfully, the migration transitions to the **TargetSelected** state. In the **targetSelected** state the source and target logical units continue to be synchronized.

The **powermig selectTarget** command fails if the handle is invalid or the device is not in the **sourceSelected** state or the **targetSelected** state.

---

**Note:** When the target logical unit is larger than the source, additional space on the target is unusable until the migration is committed.

---



---

**Note:** This command is not supported on MSCS cluster migrations.

---

### Arguments

**-handle <migrationHandle>**  
Number that identifies the migration session. It is assigned when you enter the **powermig setup** command. You must enter the handle when running commands that reference the migration session.

**-all**  
Performs **selectTarget** for all migration sessions in the **sourceSelected** state. Migrations in other states are ignored.

**-noPrompt|-no**  
Suppresses the default confirmation that appears when you run a command.

**Migration state** Required migration state to run the command: **sourceSelected**  
Migration state after the command completes: **targetSelected**

### Example

```
powermig selectTarget -handle 1
```

Output:

```
Transition to targetSelected state? [Yes]/No:
```

## powermig setup

**Purpose** Defines the source and target device name involved in the migration and the underlying technology used in conjunction with PowerPath Migration Enabler.

**Syntax**

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

**Description** **powermig setup** command is the first step in the migration process. It defines the source and target device name involved in the migration and the underlying technology used in conjunction with PowerPath Migration Enabler.

The **powermig setup** command output displays a migration handle, which identifies the migration session. You need the migration handle in subsequent **powermig** commands that reference the migration session.

The migration handle is a number between 1 and 30,000. The first migration handle is 1, and each subsequent migration handle increases by 1, until the maximum of 30,000 is reached. Once the maximum is reached, the handle number restarts at 1, or at the lowest number not in use by another migration session. The abbreviation for this argument is **-hd**.

After the **powermig setup** command completes, the migration transitions to the **setup** state. You can then continue the migration by synchronizing the source and target logical units. Or you can abandon the migration by running the **powermig cleanup** command to remove selected data from the source or target logical unit.

This command fails if:

- ◆ The source or target is inaccessible.
- ◆ The source or target is part of another migration session.
- ◆ The source and target devices have mismatched labels or the target is unlabeled.

After this command completes, Migration Enabler protects the source and target from removal or remapping. Once the migration enters the **setup** state, the target-device name and any aliases become inaccessible, and attempts to access the target fail.

---

**Note:** If environment variables are present, they will be used to replace the values stored in the PowerPath Migration Enabler database (UMD) at setup, causing powermig setup to fail.

---

### Arguments

**-techType <techType>**

The technology type used in conjunction with Migration Enabler.

Where *techType* is OR (Open Replicator), INVE (Virtual encapsulation), HOSTCOPY (Host Copy), and SYMCLONE (TimeFinder/Clone).

Migrations to an encrypted thin target, or from an encrypted thin source, do not preserve the source's allocation on the target. The abbreviation for this argument is **-tt**.

`-src <sourceName>`  
 Source device from which you want to migrate data. Specify the device name as a base name (for example, emcpowerc), or as a fully qualified path name (for example, /dev/rdisk/emcpowerc). For the **-src** argument, you must specify the source Name used by the application.

---

**Note:** The *<sourceName>* must be the same as the name applications are configured to use.

---

`-tgt <targetName>`  
 Target device to which you want to migrate data. Specify the device name as a base name (for example, emcpowerc), or as a fully qualified path name (for example, /dev/rdisk/emcpowerc). On Windows, use the name as displayed by **powermt display dev=all**.

`-file <fileName>`  
 File name that contains migration pairs. For example, `-file /tmp/pairs`. The migration pair consists of a source and target device. Create the text file with one migration pair per line. Migration pairs with handles that exist from previous migrations will be ignored.

If an error displays for a migration pair, setup will continue with the next pair. After you resolve the cause of the error, run the **powermig setup** command with the same file again to complete.

`-throttleValue <throttleValue>`  
 Value of the throttle, which slows down or speeds up the synchronization of the source and target logical units. Values range from 0 to 9, where 0 is the fastest speed and 9 is the slowest. The default value is 5 for Open Replicator, 0 for SYMCLONE, and 2 for Host Copy migrations. The throttle allows you to change (while a synchronization is in progress) the speed at which data is synchronized. The default values are: 5 for Open Replicator; 2 for Host Copy; and 0 for TimeFinder/Clone.

This option is valid only with Open Replicator, SYMCLONE, or Host Copy migrations. With Open Replicator, the throttle value determines resource consumption on the target array. With Host Copy, throttle value affects resource consumption on the host. The abbreviation for this argument is **-tv**.

If you need a faster copy, lower it, keeping in mind this will utilize more host resources. If you need to free up your host's resources increase the throttle value to slow down Host Copy. If you need to postpone utilizing host resources use **powermig pause**, then **powermig resume** when host resources are more readily available for copying.

`-techHost <hostname>`  
 Host name or IP address of the remote Solutions Enabler server used to optionally specify remote server name (or IP address) where the Solutions Enabler server daemon is running. This is only applicable when using Solutions Enabler remotely with Migration Enabler technology types Open Replicator (OR) and SYMCLONE (TimeFinder/Clone).

`-techPort <port#>`  
 Used with **-techHost** to optionally specify the port number used by the remote Solutions Enabler server daemon. The default is 2707. Use the default value.

`-nothin`

Applies only to PowerPath Encryption with RSA migrations, which use the Host Copy technology type. Use this option when the target is virtually provisioned, and the source and or target is encrypted, so as to avoid unnecessary I/O operations.

---

**Note:**

- The **-nothin** option is not supported on Linux hosts.
  - Migrations to an encrypted thin target, or from an encrypted thin source, do not preserve the source's allocation on the target.
- 

**-force**

Overrides error when the target is in use. On Solaris, **-force** also overrides label validation error after source label is copied to target.

**-cluster**

Used to setup a cluster migration.

This argument is not supported when migrating devices on Windows Server 2003.

**-noPrompt | -no**

Suppresses the default confirmation that appears when you run a command.

**Migration state**

Required migration state to run the command: Not applicable as the migration session is not set up until after running the command.

Migration state after the command completes: **setup**

**Platform-specific notes**

On AIX, when the **max\_transfer** attribute of the target is smaller than that of the source, **powermig setup** may fail. Modify the target's **max\_transfer** size to be at least equal to that of the source.

On Solaris, you may see the following warning message in a heterogeneous migration when the disks are close in size. The migration succeeds.

```
Possible alternate cylinder overlap detected using native source geometry.
```

On Windows, when the **max\_transfer** attribute of the target is smaller than that of the source, **powermig setup** may fail. Modify the target's **max\_transfer** size to be at least equal to that of the source.

**Example**

```
powermig setup -src emcpowerc -tgt emcpowerh -tt hostcopy
```

Output similar to the following appears:

```
Setup migration? [Yes]/No: yes Migration Handle = 1
```

The following example uses the **-file** option to move all migration pairs listed in the **/tmp/pairs** file to the **setup** state.

```
powermig setup -techType HOSTCOPY -file /tmp/pairs -noprompt
```

Output similar to the following appears:

```
Source = emcpower21c, Target = emcpower4c
Migration Handle = 244
Source = emcpower20c, Target = emcpower19c
Migration Handle = 245
Source = emcpower7c, Target = emcpower5c
```



```
Migration Handle = 246  
Source = emcpower6c, Target = emcpower26c  
Migration Handle = 247  
...
```

## powermig sync

**Purpose** Starts the synchronization of the target logical unit with the source logical unit.

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

**Description** **powermig sync** command starts the synchronization of the target logical unit with the source logical unit. Depending on the size of the source, the synchronization can take some time to complete.

Once the source and target logical units are synchronized the migration transitions to the **sourceSelected** state. You can check the status of a synchronization using the **powermig query** command. For Open Replicator, HOSTCOPY, or SYMCLONE migrations, you can change the speed at which the source-target data is synchronized using the **powermig throttle** command.

The migration is in the **syncing** state while data is being copied. When the synchronization completes, the migration transitions to the **sourceSelected** state.

The **powermig sync** command fails if the handle is invalid or if the migration is not in the **setup** state.

**Arguments**

- `-handle <migrationHandle>`  
Number that identifies the migration session. It is assigned when you enter the **powermig setup** command. You must enter the handle when running commands that reference the migration session.
- `-all`  
Performs sync for all migration sessions in the **setup** state. Migrations in other states are ignored.
- `-noPrompt|no`  
Suppresses the default confirmation that appears when you run a command.

**Migration state** Required migration state to run the command: **setup**  
Migration state after the command completes: **sourceSelected**

**Example** `powermig sync -handle 1`

Output similar to the following appears:

```
Start sync? [Yes]/No: yes
```

# powermig throttle

**Purpose** Sets the speed at which data is synchronized between a source-and-target pair.

**Syntax** `powermig throttle -handle <migrationHandle>|-all -throttleValue <throttleValue> [-noPrompt|-no]`

**Description** **Note:** This command is supported with all technology types except INVE (Virtual encapsulation). If you run this command when using INVE (Virtual encapsulation) as the underlying technology, while the migration transitions is in the **sourceSelected** state, a message alerts you that the migration is in the wrong state to run this command.

**powermig throttle** sets the speed at which data is synchronized between a source-and-target pair. Run this command while a synchronization is in progress. This command affects resource consumption differently depending on the underlying technology used.

**Note:** **powermig throttle** can only be changed in an OR (Open Replicator) migration while state is **syncing**. It can also be set when running **powermig setup**.

**Arguments**

`-handle <migrationHandle>`  
Number that identifies the migration session. It is assigned when you enter the **powermig setup** command. You must enter the handle when running commands that reference the migration session.

`-all`  
Performs throttle for all migration sessions in the **syncing** state. Migrations in other states are ignored.

`-throttleValue <throttleValue>`  
Value of the throttle, which slows down or speeds up the synchronization of the source and target logical units. Values range from 0 to 9, where 0 is the fastest speed and 9 is the slowest. The default value is 5 for Open Replicator, 0 for SYMCLONE, and 2 for Host Copy migrations. The throttle allows you to change (while a synchronization is in progress) the speed at which data is synchronized. The default values are: 5 for Open Replicator; 2 for Host Copy; and 0 for TimeFinder/Clone.

This option is valid only with Open Replicator, SYMCLONE, or Host Copy migrations. The abbreviation for this argument is **-tv**.

If you need a faster copy, lower it, keeping in mind this will utilize more host resources. If you need to free up your host's resources increase the throttle value to slow down Host Copy. If you need to postpone utilizing host resources use **powermig pause**, then **powermig resume** when host resources are more readily available for copying.

`-noPrompt|-no`  
Suppresses the default confirmation that appears when you run a command.

**Migration state** Required migration state to run the command: **syncing, setup**

Migration state after the command completes: This command does not change the migration state.

**Example** `powermig throttle -handle 1 -throttleValue 0`

No output appears when this command completes successfully.

## Technology-specific notes

This section describes the throttle command and arguments in more detail for each technology type.

### Host Copy

The Host Copy technology type copies data from source to target in cycles that last five seconds by default. Application I/O resumes at the end of each cycle. The **throttleValue** defines the overall percentage of time the host spends copying data as opposed to performing non-migration activities. The throttle value affects resource consumption on the host.

When the suspend time expires, the host does one of the following, depending on the throttle value:

- ◆ Restarts a new copy cycle.
- ◆ Remains idle to free host resources for other operations, such as completing application I/O or allowing other migrations to enter the copy cycle.

[Table 27 on page 164](#) shows the percentage of the host's time spent copying data as defined by the throttle value.

**Table 27** 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%

### Open Replicator

For Open Replicator migrations, the throttle value adjusts the speed at which data is transferred between the source and target logical units. For Open Replicator, the throttle works similarly to the Solutions Enabler **symrcopy** command with the **set pace** argument. A lower throttle value can significantly increase resource consumption on the target logical unit's array. Throttle values range from 0 to 9, where 0 is the fastest.

### TimeFinder/Clone

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

Symmetrix QoS values range from 0 to 16 and Migration Enabler throttle values range from 0 to 9. Migration Enabler maps the throttle values to QoS values as shown in [Table 28 on page 165](#).

**Table 28** TimeFinder/Clone throttle and QoS values

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

## powermig undoRedirect

**Purpose** Stops the redirection of I/O from the source to the target.

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

**Description** **powermig undoRedirect** stops the redirection of I/O from the source to the target. Use this command only when the source logical unit is a native-named device.

**⚠ CAUTION**

**Before using this command, bring down any application sending I/O to the target. While the application is down, run `powermig undoRedirect`, and then reconfigure the application to use the target-device name. This prevents application I/O errors.**

**Arguments**

`-handle <migrationHandle>`  
Number that identifies the migration session. It is assigned when you enter the **powermig setup** command. You must enter the handle when running commands that reference the migration session.

`-all`  
Performs `undoRedirect` for all migration sessions in the **committedAndRedirected** state. Migrations in other states are ignored.

`-force`  
Forces the redirection of I/O from the source to the target to stop even if the source device is in use.

**Note:** If you use **-force** while I/O is running to the source device, then I/O errors or disk corruption may occur.

`-noPrompt|-no`  
Suppresses the default confirmation that appears when you run a command.

**State transition** Required migration state to run the command: **committedAndRedirected**

Migration state after the command completes: **committed**

**Example**

`powermig undoRedirect -handle 1`

Output:

Undo IO redirection? [Yes]/No:

## powermig version

**Purpose** Displays the version of PowerPath Migration Enabler running on the host.

**Syntax** `powermig version`

**Description** **powermig version** displays the version of PowerPath Migration Enabler running on the host. PowerPath Migration Enabler has the same version as the PowerPath software package it is distributed with.

**Arguments** There are no arguments for this command.

**Transition state** Required migration state to run the command: Any state.

Migration state after the command completes: This command does not change the migration state.

## powermigcl: Configure Migration Enabler cluster resources

A command line utility for configuring PowerPath Migration Enabler cluster resources.

**Description** **powermigcl** configures PowerPath Migration Enabler cluster resources for physical disks in a cluster. During the configuration, a new PowerPath Migration Enabler cluster resource is created for each cluster physical disk resource and is added to its dependencies.

**Platforms** This command is supported only on PowerPath 5.7 SP1 and later for Windows. On Windows Server 2008 and later, the **powermigcl** command is not disruptive.

### NOTICE

This command is not supported when migrating devices on Microsoft Windows Server 2003.

**Syntax summary** A summary of the syntax for all **powermigcl** command is as follows:

```
powermigcl display (-all | -group <group>{,<group>}) [-verbose]
powermigcl config (-all | -group <group>{,<group>}) [-force] [-offline] [-noprompt]
powermigcl unconfig (-all | -group <group>{,<group>}) [-force] [-noprompt]
powermigcl version
powermigcl help
```

**Argument abbreviations** The **powermigcl** command supports the abbreviations listed in [Table 29 on page 168](#)

**Table 29** Argument abbreviations

Argument	Abbreviation
-all	-a
-force	-f
-group	-g
-noprompt	-no
-verbose	-v
-offline	-o



## powermigcl display

**Purpose** Displays the information on all available cluster resource groups.

**Syntax** `powermigcl display (-all | -group <group>{,<group>}) [-verbose]`

**Description** **powermigcl display** displays the information on all available cluster resource groups.

**Argument** `-all|-a`  
Information on all available cluster resource groups is displayed.

`-group|-g <group_name>`  
Information on selected cluster groups is displayed.

### NOTICE

You cannot use **-all** and **-group** arguments together.

`-verbose|-v`  
Detailed information such as group name, owner of the group, and a list of all physical disk resources with their configuration status is displayed.

### Examples

```
# powermigcl display -all
Group name | Current owner | Status |
=====|=====|=====|
Group 2 | DELL2850-01 | Configured |
Group 1 | DELL2850-01 | Configured |
Group 0 | DELL2850-01 | Configured |
Cluster Group | DELL2850-01 | Unconfigured |
```

```
# powermigcl display -group <GroupName2>..
Group name: Group 2
Owner: DELL2850-03
```

```
Resources | State |
=====|=====|
Disk 2 | Not configured |
```

```
# powermigcl display -all -verbose
Group name: Group 2
Owner: DELL2850-03
```

```
Resources | State |
=====|=====|
Disk 2 | Not configured |
```

```
Group name: Group 1
Owner:DELL2850-03
```

```
Resources | State |
=====|=====|
Disk 1 | Not configured |
```

```
Group name: Group 0
Owner:DELL2850-03
```

```
Resources | State |
=====|=====|
Disk 0 | Not configured |
```

```
Group name: Cluster Group
Owner:DELL2850-03
```

```
Resources | State |
=====|=====|
Disk F: | Quorum resource |
```

Without the **-verbose** option, short information such as group name, current owner, and status is displayed.

Example output:

Group name	Current owner	Status
Group 2	DELL2850-01	Configured
Group 1	DELL2850-01	Configured
Group 0	DELL2850-01	Configured
Cluster Group	DELL2850-01	Unconfigured

Where:

**Configured** - All physical disk resources in the cluster group have PowerPath Migration Enabler resource dependency.

**Partly configured** - At least one physical disk resource in the cluster group has PowerPath Migration Enabler resource dependency.

**Unconfigured** - No physical disk resource has PowerPath Migration Enabler resource dependency or the cluster group does not have any physical disk.

**Quorum resource** - Cluster group containing Quorum resource will not be configured by PowerPath Migration Enabler and will be skipped.

## powermigcl config

**Purpose** Configures PowerPath Migration Enabler cluster resources for physical disks in a cluster.

**Syntax** `powermigcl config (-all | -group <group>{,<group>}) [-force] [-offline] [-noprompt]`

**Description** Configures PowerPath Migration Enabler cluster resources for physical disks in a cluster. Once successfully configured, the PowerPath Migration Enabler cluster resource dependency is added to each cluster disk. By default, PowerPath Migration Enabler will not bring the cluster group offline while configuring cluster resources.

**Argument** `-all|-a`  
Configures all available cluster resources for physical disks in a cluster. However, cluster groups containing Quorum resource will not be configured and will be skipped.

---

**Note:** Without the **-force** argument, `powermigcl` attempts to configure only those groups owned by the current node.

---

`-group|-g`  
Configures the selected cluster groups.

### NOTICE

You cannot use **-all** and **-group** arguments together.

---

`-force|-f`  
Configures all selected cluster groups. Without this option, `powermigcl` attempts to configure only those groups owned by the current node.

`-offline|-o`  
Brings the cluster group offline for configuration. By default, PowerPath Migration Enabler will not bring the cluster group offline while configuring cluster resources.

`-noprompt|-no`  
Suppresses the default confirmation that is displayed when you run a command.

**Examples**

```
powermigcl config -group <GroupName1>..
Resource name: Disk 1
Skipped: Resource is already configured

# powermigcl config -group <GroupName0>..
Group name:Group 0
Resource name: Disk 0
Skipped: Resource is already configured

# powermigcl config -all -force -no
Group name:Cluster Group
Skipped: Group contains Quorum resource
```

## powermigcl unconfig

**Purpose** Unconfigures PowerPath Migration Enabler cluster resources for physical disks in a cluster.

**Syntax** `powermigcl unconfig (-all | -group <group>{,<group>}) [-force] [-noprompt]`

**Description** Unconfigures PowerPath Migration Enabler cluster resources for physical disks in a cluster. The PowerPath Migration Enabler cluster resource dependency which was added to each cluster disk at the time of configuration will be removed. By default, PowerPath Migration Enabler will not bring the cluster group offline while unconfiguring cluster resources.

**Argument** `-all | -a`  
Unconfigures all available cluster resources for physical disks in a cluster.

---

**Note:** Without the **-force** argument, this command affects only those groups that are owned by the current node.

---

`-group | -g`  
Unconfigures the selected cluster group.

---

**Note:** Without the **-force** argument, this command affects only those groups that are owned by the current node.

---

`-force | -f`  
Unconfigures all selected cluster groups. Without this option, `powermigcl` attempts to unconfigure only those groups owned by the current node.

`-offline | -o`  
Brings the cluster group offline for unconfiguration.

`-noprompt | -no`  
Suppresses the default confirmation that is displayed when you run a command.

**Examples** `# powermigcl unconfig -all`

```
Group name: Cluster Group
Skipped: Group is quorum.
```

---

**Note:** If the command or argument successfully executes, all default confirmation messages will be suppressed. The `powermigcl` command line utility will only display information when there is an occurrence of failed or skipped resources.

---

```
# powermigcl unconfig -group <GroupName1>,<GroupName2>..
```

```
# powermigcl unconfig -all -force -no
```

## powermigcl version

**Purpose** Displays the current installed version of powermigcl.

**Syntax** `powermigcl version`

**Description** Displays the current installed version of powermigcl.

**Example** `powermigcl version`

Output:

```
EMC powermigcl for PowerPath (c) Version 6.0 (build xx).
```

## powermigcl help

**Purpose** Provides help for the powermigcl command.

**Syntax** `powermigcl help`

**Description** Provides help for the powermigcl command.

**Example** `powermigcl help config`

Output:

Provides help for the config command. **powermigcl help** command can also be used in conjunction with other commands or arguments to get help.

# CHAPTER 3

## PowerPath Encryption CLI

This chapter describes the PowerPath Encryption command line utilities.

- ◆ PowerPath Encryption new and changed commands ..... 166
- ◆ ckmadm: configure and manage the key manager client..... 167
- ◆ cstadmin: manages the lockbox ..... 175
- ◆ powervt: manages virtual logical units ..... 179
- ◆ xcrypt\_config script, RKM\_Config.bat batch file, and RKM\_Config.exe file ..... 188

## PowerPath Encryption new and changed commands

Table 30 on page 166 lists the new and changed PowerPath Encryption commands added in versions 5.2 and later, with the earliest supported version listed by platform. Support in later versions (including service packs) can be assumed if no exception in the table notes otherwise. Version 5.2 is the first PowerPath Encryption with RSA release. The following commands require administrator (root) privileges.

### IMPORTANT

PowerPath 5.7 is the last version of PowerPath to support encryption.

**Table 30** PowerPath Encryption new and changed commands

Command	Earliest supported version				
	AIX	HP-UX	Linux	Solaris	Windows
ckmadm help	N/A <sup>a</sup>	N/A	5.3	5.2	5.2
ckmadm getconfig	N/A	N/A	5.3	5.2	5.2
ckmadm getkeyclass	N/A	N/A	5.3	5.2	5.2
ckmadm setconfig	N/A	N/A	5.3	5.2	5.2
ckmadm setup	N/A	N/A	5.3	5.2	5.2
ckmadm setkeyclass	N/A	N/A	5.3	5.2	5.2
ckmadm upgrade	N/A	N/A	5.3	5.2	5.2
cstadmin initialize	N/A	N/A	5.3	5.2	5.2
cstadmin new-passphrase	N/A	N/A	5.3	5.2	5.2
cstadmin sign-file	N/A	N/A	5.3	5.2	5.2
powervt help	N/A	N/A	5.3	5.2	5.2
powervt resize	N/A	N/A	5.5	5.2	5.2
powervt update	N/A	N/A	5.3	5.2	5.2
powervt version	N/A	N/A	5.3	5.2	5.2
powervt xcrypt -on	N/A	N/A	5.3	5.2	5.2
powervt xcrypt -off	N/A	N/A	5.3	5.2	5.2
powervt xcrypt -info	N/A	N/A	5.3	5.2	5.2
xcrypt_config	N/A	N/A	5.3	5.2	N/A
RKM_Config.bat, RKM_Config.exe	N/A	N/A	N/A	N/A	5.2
xcrypt_upgrade	N/A	N/A	5.3	5.2	N/A

a. N/A means the command or utility is not supported on the platform.



## ckmadm: configure and manage the key manager client

Use the **ckmadm** commands to configure and manage the key manager client on a PowerPath Encryption host.

**Syntax**

```
ckmadm help|-h [<operation>]
ckmadm getconfig
ckmadm getkeyclass
ckmadm setconfig -file rkm_svc.conf|-transport <transport_name>|
  -cache <cache_name>|-log <log_name>
ckmadm setup -file rkm_init.conf
ckmadm setkeyclass -file rkm_keyclass.conf
ckmadm upgrade [-cache_file <filename>] [-init_file rkm_init.conf]
  [-config_file rkm_svc.conf] [-keyclass_file rkm_keyclass.conf]
```

**Prerequisites** Complete the following prerequisites before running the **ckmadm** command:

- ◆ Install the PowerPath Encryption software on the host following the instructions in the platform specific PowerPath Installation and Administration Guide.
- ◆ Configure the PowerPath Encryption software.
- ◆ Prepare the RSA Data Protection Manager (DPM) appliance to support PowerPath Encryption.

**Status codes** The **ckmadm** command returns a zero for success and a non-zero number if it encounters a problem.

## ckmadm help

**Purpose** Provides a description and usage for the **ckmadm** commands and options.

**Syntax** `ckmadm help|-h [<operation>]`

**Description** The **ckmadm help** command displays a brief description and usage summary for the specified **ckmadm** command. By default, the supported commands are listed. Specify a specific command (for example, **ckmadm help setconfig**) to display help for that command.

**Arguments** `<operation>`  
 Name of the **ckmadm** operation about which you want to view information. The possible values are:

setup	setconfig	setkeyclass
getconfig	getkeyclass	upgrade

**Examples**

```
ckmadm help setconfig
ckmadm -h setup
```

## ckmadm getconfig

**Purpose** Returns the names of the active transport service configuration, cache service configuration, and log service configuration on a PowerPath host.

**Syntax** `ckmadm getconfig`

**Description** The **ckmadm getconfig** command returns the names of the active transport service configuration, cache service configuration, and log service configuration on a PowerPath Encryption host.

**Arguments** None

**Examples** `ckmadm getconfig`

## ckmadm getkeyclass

**Purpose** Returns the default key class for a PowerPath Encryption host.

**Syntax** `ckmadm getkeyclass`

**Description** The **ckmadm getkeyclass** command returns the default key class for a PowerPath Encryption host. The default key class is set by the **ckmadm setkeyclass** command.

**Arguments** None

**Examples** `ckmadm getkeyclass`

## ckmadm setconfig

**Purpose** Activates a transport service, cache service, or log service configurations on a PowerPath Encryption host.

**Syntax** `ckmadm setconfig -file rkm_svc.conf|-transport <transport_name>|  
-cache <cache_name>|-log <log_name>`

**Description** The **ckmadm setconfig** command activates a specific transport configuration, cache service configuration, or log service configuration on a PowerPath Encryption host.

During PowerPath Encryption configuration, the configuration script calls the **ckmadm setup** command to define a set of transport, cache, and log service configurations. The configuration script then calls the **ckmadm setconfig** command to activate the initial transport service, cache service, and log service configuration.

You can run the **ckmadm setconfig** command to activate a different transport service, cache service, or log service after initial PowerPath Encryption configuration. The new configuration takes effect after the **emcp\_xcryptd** daemon is stopped and restarted, or the next time that the **powervt** command is run.

**Arguments**

- file *rkm\_svc.conf*  
ASCII file that identifies the transport service, cache service, and log service configurations to activate.
- transport <transport\_name>  
Identifies the transport service configuration to activate. The *transport\_name* argument corresponds to a **configName** set in the *rkm\_init.conf* file.
- cache <cache\_name>  
Identifies the cache service configuration to activate. The *cache\_name* argument corresponds to a **configName** set in the *rkm\_init.conf* file.
- log <log\_name>  
Identifies the log service configuration to activate. The *log\_name* argument corresponds to a **configName** set in the *rkm\_init.conf* file.

**Examples**

```
ckmadm setconfig -file /etc/emc/rsa/rkm_client/config/rkm_svc.conf
ckmadm setconfig -file <install-dir>
files\emc\rsa\rkm_client\config\rkm_svc.conf
ckmadm setconfig -transport transport_cfg1
```

## ckmadm setup

**Purpose** Defines the transport service, cache service, and log service configurations on a PowerPath Encryption host.

**Syntax** `ckmadm setup -file rkm_init.conf`

**Description** The **ckmadm setup** command defines the transport service, cache service, and log service configurations on a PowerPath Encryption host.

You do not call the **ckmadm setup** command directly. The **xcrypt\_config** (Linux and Solaris), **RKM\_Config.bat** (Windows) and **RKM\_Config.exe** (Windows) scripts call the **ckmadm setup** command during initial PowerPath Encryption configuration. The **ckmadm setup** command:

- ◆ Creates the transport, cache, and log service configurations specified in the **rkm\_init.conf** file.
- ◆ Prompts for the password used to protect the PowerPath Encryption PKCS#12 credentials file and stores it in the lockbox. This password was set when the PKCS#12 credentials file was created by the certification authority.

After the **ckmadm setup** command completes successfully, the set of configurations specified in the **rkm\_init.conf** file are defined on the PowerPath Encryption host. The **ckmadm setconfig** command activates specific configurations. The **ckmadm setconfig** command is called by the **xcrypt\_config** (Linux and Solaris), **RKM\_Config.bat** (Windows) and **RKM\_Config.exe** (Windows) scripts during initial PowerPath Encryption configuration.

**Arguments** `-file rkm_init.conf`  
ASCII file that sets values for the transport service, cache service, and log service configuration attributes.

**Examples**

```
ckmadm setup -file /etc/emc/rsa/rkm_client/config/rkm_init.conf
ckmadm setup -file <install-dir>
files\emc\rsa\rkm_client\config\rkm_init.conf
```

## ckmadm setkeyclass

**Purpose** Sets the default key class to be used by the RSA Data Protection Manager (DPM) server when provisioning keys for a PowerPath Encryption host.

**Syntax** `ckmadm setkeyclass -file rkm_keyclass.conf`

**Description** The **ckmadm setkeyclass** command sets the default key class used by the DPM server to provision keys for a PowerPath Encryption host. The **ckmadm setkeyclass** command is called by the **xcrypt\_config** script during PowerPath Encryption configuration to set the initial default key class.

You can run this command at any time after initial PowerPath Encryption configuration to change the default key class. The new default key class only affects logical units where encryption is turned on after the completion of the **ckmadm setkeyclass** command.

**Arguments** `-file rkm_keyclass.conf`  
ASCII file that contains the key class name. The key class name must correspond to the name of a key class defined on the DPM server.

**Examples**

```
ckmadm setkeyclass -file
/etc/emc/rsa/rkm_client/config/rkm_keyclass.conf
ckmadm setkeyclass -file <install-dir>
files\emc\rsa\rkm_client\config\rkm_keyclass.conf
```

## ckmadm upgrade

**Purpose** Loads the new RKM configuration files into CST.

**Syntax** `ckmadm upgrade [-cache_file <filename>] [-init_file rkm_init.conf] [-config_file rkm_svc.conf] [-keyclass_file rkm_keyclass.conf]`

**Description** The **ckmadm upgrade** command combines the functionality of the **ckmadm setup**, **ckmadm setconfig**, and **ckmadm setkeyclass** commands so that you can update the configuration files in CST simultaneously. Running the options individually produces the same results as the **ckmadm setup**, **ckmadm setconfig**, and **ckmadm setkeyclass** commands. Depending on the options you type, this command:

- ◆ Identifies the cache file specified in the **rkm\_init.conf** file (-cache\_file <filename>).
- ◆ Creates the transport, cache, and log service configurations specified in the **rkm\_init.conf** file (-init\_file rkm\_init.conf).
- ◆ Activates the transport, cache, and log service configurations on a PowerPath Encryption host (-config\_file rkm\_svc.conf).
- ◆ Sets the default key class used by the DPM server to provision keys for a PowerPath Encryption host (-keyclass\_file rkm\_keyclass.conf).

Run **ckmadm upgrade** when using a new DPM server or after upgrading PowerPath Encryption.

**Arguments** `-cache_file <filename>`  
Identifies the cache file in the **rkm\_init.conf** file. For example, where <filename> is **km27apps.cache**. When you use this option, `-init_file rkm_init.conf` and `-config_file rkm_svc.conf` are required.

`-init_file rkm_init.conf`  
ASCII file that sets values for the transport service, cache service, and log service configuration attributes. Prompts for the password used to protect the PowerPath Encryption PKCS#12 credentials file and stores it in the lockbox.

`-config_file rkm_svc.conf`  
ASCII file that identifies the transport service, cache service, and log service configurations to activate.

`-keyclass_file rkm_keyclass.conf`  
ASCII file that contains the key class name. The key class name must correspond to the name of a key class defined on the DPM server.

**Examples** The following updates all configuration files:

```
ckmadm upgrade -cache_file km27apps.cache -init_file rkm_init.conf
               -config_file rkm_svc.conf -keyclass_file rkm_keyclass.conf
```



## cstadmin: manages the lockbox

PowerPath Encryption securely stores configurations and passwords in a lockbox (an encrypted file).

**Syntax**

```
cstadmin initialize "<directory>" -unattended -overwrite  
cstadmin new-passphrase [-cstdir="<directory>"]  
cstadmin sign-file -[cstdir="<directory>"] <fileName>
```

## cstadmin initialize

**Purpose** Creates the lockbox.

**Syntax** `cstadmin initialize "<directory>" -unattended -overwrite`

**Description** The **cstadmin initialize** command creates the lockbox. The **cstadmin initialize** command is run by the **xcrypt\_config** (Linux and Solaris), **RKM\_Config.bat** (Windows) and **RKM\_Config.exe** (Windows) scripts during PowerPath Encryption configuration. You do not run this command directly.

The **cstadmin initialize** command:

- ◆ Prompts for the lockbox passphrase and then creates the lockbox using the provided passphrase. On Linux and Solaris hosts, the default location for the encrypted lockbox file is `/etc/emc/rsa/cst/config`. On a Windows host, the default location for the encrypted lockbox file is `<install-dir>\RSA\CST\Lib\`.
- ◆ Generates the keys to sign and encrypt the configuration files and stores them in the lockbox.
- ◆ Digitally signs the configuration files.
- ◆ Retrieves the system fingerprint, comprised of stable system values (SSVs) from the application host and stores this fingerprint in the lockbox.

**Arguments** `"<directory>"`  
The path to the directory containing the configuration files, installation files, and the lockbox.

`-unattended`  
Creates a lockbox that can be used in an unattended (automatic) mode. Manual entry of a password is not required to open the lockbox. Rather, a set of SSVs from the host platform are checked prior to allowing the lockbox to be opened.

`-overwrite`  
Enables a new lockbox to overwrite an existing lockbox.

**Examples**

```
cstadmin initialize /etc/emc/rsa/cst/config -unattended -overwrite
cstadmin initialize "<install-dir>\EMC\RSA\CST\Lib" -unattended
-overwrite
```

## cstadmin new-passphrase

**Purpose** Changes the passphrase for the lockbox.

**Syntax** `cstadmin new-passphrase [-cstdir="<directory>"]`

**Description** The **cstadmin new-passphrase** command changes the passphrase for the lockbox. The initial lockbox passphrase is set when the **cstadmin initialize** command is run by the **xcrypt\_config** (Linux and Solaris), **RKM\_Config.bat** (Windows) and **RKM\_Config.exe** (Windows) scripts during PowerPath Encryption configuration.

The new lockbox passphrase must be a minimum of eight characters in length and contain at least one each of the following:

- ◆ Numeric character
- ◆ Uppercase character
- ◆ Lowercase character
- ◆ Non-alphanumeric character, such as # or !

The **cstadmin new-passphrase** command fails if one or more of the following conditions exist:

- ◆ The lockbox cannot be found.
- ◆ The current passphrase entered is invalid.
- ◆ The new passphrase does not meet the password strength criteria.

**Arguments** `-cstdir="<directory>"`  
Indicates the directory containing the configuration files. On Linux and Solaris hosts, this directory is `/etc/emc/rsa/cst/config`. On a Windows host, this directory is `<install-dir>\EMC\RSA\CST\Lib`.

**Example** `cstadmin new-passphrase -cstdir="<install-dir>\EMC\RSA\CST\Lib"`

## cstadmin sign-file

**Purpose** Re-signs the CST lockbox.

**Syntax** `cstadmin sign-file -[cstdir="<directory>"] <fileName>`

**Description** The **cstadmin sign-file** command re-signs the CST lockbox after a host name or IP address change.

**Arguments**

- cstdir="<directory>"  
Indicates the directory containing the configuration files. On Linux and Solaris hosts, this directory is /etc/emc/rsa/cst/config. On a Windows host, this directory is <install-dir>\EMC\RSA\CST\Lib.
- <fileName>  
Full file name of the configuration file to be re-signed. For example, Config.xml or "\*.xml".

**Example** `cstadmin sign-file -cstdir="<install-dir>\EMC\RSA\CST\Lib" "*.xml"`

## powervt: manages virtual logical units

The **powervt** command manage virtual logical units.

**Syntax**

```
powervt help|-h [<operation>] [<option>]
powervt resize -dev <devname> [<devname> ...] |-file <devfile>
powervt update -dev <devname> [<devname> ...] |-file <devfile>
powervt version
powervt xcrypt -on -dev <devname> [<devname> ...] |-file <devfile>
[-noprompt|-no] [-force]
powervt xcrypt -off -dev <devname> [<devname> ...] |-file <devfile>
[-noprompt|-no] [-force]
powervt xcrypt -info [<encrypted>] -dev <devname> [<devname>
...]|all|-file <devfile> [-refresh] [-detail]
```

**Prerequisites** PowerPath Encryption must be installed and configured on the host.

**Status codes** [Table 31 on page 179](#) lists the **powervt xcrypt** status codes.

**Table 31** powervt xcrypt status codes

Status code	Description
0	The command succeeded with no warnings or errors.
2	The command failed.
3	The command did not supply a required argument.
4	The command specified an invalid argument.
5	The command specified an invalid operation, or no operation was specified.

**Audit information** **powervt xcrypt -on** and **powervt xcrypt -off** command actions are recorded in the audit log.  
**powervt xcrypt -info** command actions are not recorded in the audit log.

## powervt help

**Purpose** Provides a description and usage for the **powervt** commands and options.

**Syntax** `powervt help|-h [<operation>] [<option>]`

**Description** The **powervt help** command displays a brief description and usage summary for the specified **powervt** command. By default, the supported commands are listed. Specify a specific command (for example, **powervt help xcrypt**) to display help for that command.

**Arguments**

`<operation>`  
Name of the **powervt** operation about which you want to view information. The possible values are:

<code>xcrypt</code>	<code>version</code>	<code>resize</code>	<code>update</code>
---------------------	----------------------	---------------------	---------------------

`<option>`  
Name of the option for which you want to view help. Possible values are:

<code>-on</code>	<code>-off</code>	<code>-info</code>
<code>-device</code>	<code>-file</code>	<code>-noprompt</code>
<code>-force</code>	<code>-refresh</code>	

**Examples**

```
powervt help resize
powervt -h xcrypt -on
```

## powervt resize

**Purpose** Adjusts the usable space on an encrypted virtual logical unit to reflect its actual capacity.

**Syntax** `powervt resize -dev <devname> [<devname> ...] |-file <devfile>`

**Description** The **powervt resize** command adjusts the usable space on an encrypted virtual logical unit to reflect its actual capacity.

Migrating plaintext data to an encrypted virtual logical unit requires that the target virtual logical unit be larger than the source plaintext logical unit. After the migration completes, the usable space on the encrypted virtual logical unit is equivalent to the size of the source plaintext logical unit. The **powervt resize** command grows the usable space on the virtual logical unit to its actual size.

**Arguments** `-dev <devname> [<devname> ...] |-file <devfile>`  
 Specifies the encrypted virtual logical units that will be resized. The specified device name must be a PowerPath pseudo device.

At the command prompt, use a space to separate the logical units. In a file, place logical units on separate lines.

**Examples**

```
powervt resize -dev emcpowera
powervt resize -dev emcpowera emcpowerb
powervt resize -file dev_list.txt
```

## powervt update

**Purpose** Forces the PowerPath kernel to verify metadata access and to look up the key lookup for the target logical unit.

**Syntax** `powervt update -dev <devname> [<devname> ...] |-file <devfile>`

**Description** The **powervt update** command forces the PowerPath kernel to verify metadata access and to look up the key for the target logical unit.

On Linux, Solaris, and Windows hosts, PowerPath Encryption cannot detect state changes (established to split, split to established) for replicas. Run the **powervt update** command:

- ◆ after a replica changes state to ensure that PowerPath recognizes the change
- ◆ after an encrypted device is added as a resource to ensure that each node recognizes the encrypted device
- ◆ after a LUN is grown from the array to move the metadata to claim the new space

**Arguments** `-dev <devname> [<devname> ...] |-file <devfile>`  
Specifies the encrypted virtual logical units that will be updated. The specified device name must be a PowerPath pseudo device.

At the command prompt, use a space to separate the logical units. In a file, place logical units on separate lines.

**Examples**

```
powervt update -dev emcpowera
powervt update -dev emcpowera emcpowerb
powervt update -file dev_list.txt
```



## powervt version

**Purpose** Displays the version of the **powervt** command running on the host.

**Syntax** `powervt version`

**Description** The **powervt version** command displays the version of the **powervt** command running on the host. The **powervt** command has the same version as the PowerPath Encryption software with which it is distributed.

**Arguments** None

**Example** `powervt version`

## powervt xcrypt -on

**Purpose** Enables encryption on one or more plaintext logical units.

**Syntax** `powervt xcrypt -on -dev <devname> [<devname> ...] | -file <devfile> [-noprompt | -no] [-force]`

**Description** The **powervt xcrypt -on** command enables encryption on one or more plaintext logical units. After the **powervt xcrypt -on** command completes successfully:

- ◆ Any previously existing user data on the logical unit is inaccessible through PowerPath Encryption. Specifically, any plaintext data before enabling encryption is returned as ciphertext on subsequent reads.
- ◆ User data is accessible only through PowerPath Encryption.
- ◆ All writes to the logical unit are encrypted; all reads are decrypted.
- ◆ The logical unit can be put into service as a new encrypted virtual logical unit or it can be used as the target of a Migration Enabler migration where the source is a plaintext logical unit with existing user data.

### Arguments

**-on**  
Enables encryption for the plaintext logical units specified by the **-dev** or **-file** argument.

**-dev <devname> [<devname> ...] | -file <devfile>**  
Specifies the plaintext logical units for which encryption will be enabled. The specified device name must be a PowerPath pseudo device.

At the command prompt, use a space to separate the logical units. In a file, place logical units on separate lines.

**-noprompt | -no**  
Suppresses the default confirmation that appears when you run a command.

**-force**  
By default, the **powervt xcrypt -on** command fails if you attempt to enable encryption on an open logical unit. Specifying the **-force** option enables encryption even if the logical unit is open.

### Examples

```
powervt xcrypt -on -dev emcpowera
powervt xcrypt -on -dev emcpowera emcpowerb
powervt xcrypt -on -file dev_list
```

## powervt xcrypt -off

**Purpose** Disables encryption for one or more encrypted virtual logical units.

**Syntax** `powervt xcrypt -off -dev <devname> [<devname> ...] | -file <devfile> [-noprompt | -no] [-force]`

**Description** The **powervt xcrypt -off** command disables encryption on one or more encrypted virtual logical units. After the **powervt xcrypt -off** command completes successfully:

- ◆ All previously existing user data on the logical unit becomes inaccessible. Specifically, any encrypted data on the logical unit is not decrypted and appears as ciphertext on subsequent reads.
- ◆ The logical unit is available to store plaintext data.
- ◆ Writes to the logical unit are not encrypted; reads are not decrypted.

### Arguments

**-off**  
Disables encryption for the virtual logical units specified by the **-dev** or **-file** argument.

**-dev <devname> [<devname> ...] | -file <devfile>**  
Specifies the virtual logical units for which encryption will be disabled. The specified device name must be a PowerPath pseudo device.

At the command prompt, use a space to separate the virtual logical units. In a file, place virtual logical units on separate lines.

**-noprompt | -no**  
Suppresses the default confirmation that appears when you run a command.

**-force**  
By default, the **powervt xcrypt -off** command fails if you attempt to disable encryption on an open virtual logical device. Specifying the **-force** option disables encryption even if the logical unit is open.

### Examples

```
powervt xcrypt -off -dev emcpowera
powervt xcrypt -off -dev emcpowera emcpowerb
powervt xcrypt -off -file dev_list.txt
```

## powervt xcrypt -info

**Purpose** Indicates whether one or more logical units are encrypted.

**Syntax** `powervt xcrypt -info [<encrypted>] -dev <devname> [<devname> ...]|all|-file <devfile> [-refresh] [-detail]`

**Description** The **powervt xcrypt -info** command indicates the encryption status of a logical unit. A logical unit can be in one of the states described in [Table 32 on page 186](#).

**Table 32** Encryption status

Status	Description
Encrypted	Logical unit is encrypted.
Encrypted, but no key available	PowerPath Encryption can read the metadata on an encrypted virtual logical unit, but cannot access the key. To resolve this problem: <ul style="list-style-type: none"> <li>• Ensure that the <b>emcp_xcryptd</b> (Linux and Solaris) and the <b>emcp_xcryptd/EMC PowerPath RSA Encryption Service</b> (Windows) service is running on the PowerPath host.</li> <li>• Ensure the DPM appliance is running and accessible from the PowerPath host.</li> <li>• Ensure the key manager client configuration is correct.</li> <li>• Look for messages from the key manager client in the system log.</li> </ul> After attempting to resolve the problem, run <b>powervt xcrypt -info -refresh</b> to retry the key lookup.
Encryption status is unknown	PowerPath Encryption cannot read the metadata on the logical unit. This state may indicate that the logical unit is offline or has some other media error. Ensure that the PowerPath host can access the device and then run <b>powervt xcrypt -info -refresh</b> to retry the key lookup.
Not encrypted	Logical unit is not encrypted.

### Arguments

**-info**

Indicates whether the logical units specified by the **-dev** or **-file** argument are encrypted.

**-dev <devname> [<devname> ...]|all|-file <devfile>**

Specifies the logical units for which encryption information will be displayed. The specified device name must be a PowerPath pseudo device.

At the command prompt, use a space to separate the logical units. Use the **-dev all** argument to display information for all logical units on the host. In a file, place logical units on separate lines.

**-refresh**

Forces the kernel to refresh its internal data structures before performing a key lookup. Use this option to get the keys from the RKM again and when an instance of the **powervt xcrypt -info** command indicates that the state of a logical unit is either **encrypted but no key available** or **unknown**.

**-detail**

Displays additional information about the logical unit, including its pseudo name, storage array ID, logical device ID, and state.

### Examples

```
powervt xcrypt -info -dev emcpowera
powervt xcrypt -info -dev emcpowera emcpowerb
```

```
powervt xcrypt -info -dev all  
powervt xcrypt -info -file dev_list.txt
```

## xcrypt\_config script, RKM\_Config.bat batch file, and RKM\_Config.exe file

**Purpose** Configures the PowerPath Encryption software on a host.

**Syntax** `xcrypt_config` (Linux and Solaris)  
`RKM_Config.bat` (Windows)  
`RKM_Config.exe` (Windows)

**Description** The **xcrypt\_config** command (Linux and Solaris), the `RKM_Config.bat` batch file (Windows), and the `RKM_Config.exe` file (Windows) configure the PowerPath Encryption software on a PowerPath Encryption host. The configuration script simplifies the configuration process by providing a single point that marshals all configuration steps needed by PowerPath Encryption. Specifically, the configuration script and GUI.

- ◆ Validate that a PowerPath Encryption license is registered on the host. The script returns an error if there is no valid PowerPath Encryption license on the host (for **xcrypt\_config** and `RKM_Config.bat` only).
- ◆ Provide the interface to fill all the configuration details into the following configuration files:
  - `rkm_init.conf`
  - `rkm_keyclass.conf`
  - `rkm_registration.conf`
  - `rkm_svc.conf`
- ◆ Run the **cstadmin initialize** command to configure the lockbox. The **cstadmin initialize** command:
  - Prompts for the passphrase for the lockbox.
  - Configures a fixed location for the lockbox. On Linux and Solaris hosts, the default location for the encrypted lockbox file is `/etc/emc/rsa/cst/lib`. On a Windows host, the default location for the encrypted lockbox file is `<install-dir>\EMC\RSA\CST\Lib`.
- ◆ Run the following **ckmadm** commands to configure and initialize the key manager client:
  - The **ckmadm setup** command:
    - Prompts for the password used to protect the PowerPath host's PKCS#12 credentials file and stores it in the CST lockbox. This password was set when the credentials file was created by the certification authority.
    - Generates the cache file password and stores it in the lockbox.
    - Creates the transport service, cache service, and log service configurations defined in the `rkm_init.conf` file.
  - The **ckmadm setconfig** command activates the transport service, cache service, and log service configuration specified in the `rkm_svc.conf` file.
  - The **ckmadm setkeyclass** command sets the default key class specified in the `rkm_keyclass.conf` file.

- ◆ On a Linux host, create the `/etc/emc/xcrypt_cfg_done` file, which indicates that PowerPath Encryption configuration has completed successfully on the host.
- ◆ Start the PowerPath Encryption **emcp\_xcryptd** daemon on a Linux host or the PowerPath Encryption with RSA Encryption service on a Windows host
- ◆ Enable the startup code for reboots.
- ◆ Create the `<install-dir>\EMC\RSA\CST\Lib\xcrypt_cfg_done` file, which indicates that PowerPath Encryption configuration has completed successfully on the host (for **RKM\_Config.exe** only).

After the configuration script or GUI completes, PowerPath Encryption is configured on the host and ready to encrypt data written to virtual logical units.

The configuration script only needs to be run once to perform the initial PowerPath Encryption configuration. Use the **ckmadm** command to update configuration attributes after initial configuration. If **RKM\_config.bat** or **cstadmin initialize** command is run again, there is a high probability that the cache file may be corrupted. To rectify this problem, delete the cache file (`<install-dir>\EMC\RSA\Rkm_client\config\km27apps.cache`) and reboot the PowerPath Encryption host with the DPM appliance running.

## Reconfigure the PowerPath Encryption host

In case you cannot access the LUNs on a PowerPath Encryption host when the primary DPM appliance is not functional, use the following steps to reconfigure PowerPath Encryption host to use Secondary DPM appliance.

---

**Note:** Except Windows, this procedure is applicable for all platforms.

---

1. Edit the `rkm_init.conf` file to define a new global service, transport service, cache service, or log service configuration that contains the updated attributes.
2. Run **ckmadm setup -file rkm\_init.conf** to define the new service configuration. At the prompt type the client credential password.
3. Edit the `rkm_svc.conf` file to point to the new service.
4. Run **ckmadm setconfig -file rkm\_svc.conf** to activate the new service configuration.
5. Stop and restart **EMC PowerPath RSA Encryption Service** through the services console. For Windows, follow step 6.

---

**Note:** Do not run the **RKM\_config.bat** configuration script or **cstadmin initialize** command.

---

6. On Windows, click **All Programs > EMC > Configuration > RKM client configuration**.

This displays a dialog box that requires lockbox and **passphrase** password details. Enter the details to proceed with the configuration.

---

**Note:** Before you configure the RKM client, check if the **xcrypt\_cfg\_done** file is present. This file is hidden by default. Delete this file and proceed with reconfiguration.

---

**Arguments**      None.

## Prerequisites

Ensure that the following prerequisites have been completed before running the configuration script:

- ◆ The PowerPath Encryption software is installed on the host. The platform specific PowerPath Installation and Administration Guide and *EMC PowerPath Encryption with RSA User Guide* describes how to install the PowerPath Encryption software.
- ◆ The required PowerPath host's PKCS#12 credentials file and trusted root certificate are available on the PowerPath Encryption host.

## Audit information

Configuration script actions are recorded in the audit log (for **xcrypt\_config** script and **RKM\_Config.bat** batch file only).

## Errors

The configuration script fails if one or more of the following conditions exist (for **xcrypt\_config** script and **RKM\_Config.bat** batch file only):

- ◆ A valid PowerPath Encryption license is not registered on the host.
- ◆ The configuration script has already been run on the host.
- ◆ If any of the commands called by the configuration script encounters an error.
- ◆ On Linux hosts, PowerPath is not started.

[“Error messages” on page 242](#) describes the error messages returned by the configuration scripts.

Errors are similar to those which pertain to **ckmadm** and **cstadmin** commands (for **RKM\_Config.exe** only).



## PART 2

# PowerPath Family System Messages Reference

This section lists error messages for PowerPath, PowerPath Migration Enabler, and PowerPath Encryption with RSA and suggests corrective action to resolve the error. It also describes audit logging for each product.

### [Chapter 4, “PowerPath Multipathing Messages”](#)

Describes PowerPath multipathing error messages and suggests corrective action.

### [Chapter 5, “PowerPath Migration Enabler Messages”](#)

Describes Migration Enabler error messages and suggests corrective action.

### [Chapter 6, “PowerPath Encryption Messages”](#)

Describes PowerPath Encryption error messages and suggests corrective action.



# CHAPTER 4

## PowerPath Multipathing Messages

This chapter lists messages returned by the PowerPath driver, PowerPath installation process, **powermt** utility, and other PowerPath utilities, and suggests how to respond to them.

- ◆ [Audit and error messages](#) ..... 193
- ◆ [Introduction to PowerPath error messages](#) ..... 196
- ◆ [PowerPath error messages](#) ..... 197
- ◆ [Windows system event IDs](#) ..... 225
- ◆ [Windows application event IDs](#) ..... 227

## Audit and error messages

PowerPath reports any errors, diagnostic messages, and failover recovery messages through the syslog file or Event viewer (Windows) that is specified by the administrator.

**Table 33** Default message logging location

Platform	Default message logging location
AIX	/usr/safe.log
HP-UX	/var/adm/syslog/syslog.log
Linux	/var/log/messages
Solaris	/var/adm/messages
Windows	Event Viewer

## Audit log messages

Some commands generate audit messages that allow you to track the commands run on a particular host. The following commands generate audit messages.

**Table 34** Commands that generate audit messages (page 1 of 2)

Command	Version support was added				
	AIX	HP-UX	Linux	Solaris	Windows
emcpadm renamepseudo	5.3	N/A <sup>a</sup>	5.3	5.3	N/A
powermt check	5.3	5.1 SP2	5.3	5.3	5.3
powermt config	5.3	5.1 SP2	5.3	5.3	N/A
powermt disable hba	5.3	5.0	5.3	5.0	N/A
powermt enable hba	5.3	5.0	5.3	5.0	N/A
powermt load	5.3	5.1 SP2	5.3	5.3	5.3
powermt manage	5.3	5.1 SP2	5.3	5.3	N/A
powermt release	N/A	N/A	5.3	N/A	N/A
powermt remove	5.3	5.1 SP2	5.3	5.3	5.3
powermt restore	5.3	5.1 SP2	5.3	5.3	5.3
powermt save	5.3	5.1 SP2	5.3	5.3	5.3
powermt set mode	5.3	5.1 SP2	5.3	5.3	5.3
powermt set path_latency_monitor	5.3	5.0	5.0	5.0	5.3
powermt set path_latency_threshold	5.3	5.0	5.0	5.0	5.3
powermt set periodic_autorestore	5.3	5.1 SP2	5.3	5.3	5.3
powermt set policy	5.3	5.0	5.0	5.0	5.3
powermt set port_disable	5.3	5.0	5.3	5.0	N/A
powermt set streamio_threshold	N/A	5.1 SP2	5.7	5.3	5.5

**Table 34** Commands that generate audit messages (page 2 of 2)

Command	Version support was added				
	AIX	HP-UX	Linux	Solaris	Windows
powermt set write_throttle	5.3	5.1 SP2	N/A	5.3	N/A
powermt set write_throttle_queue	5.3	N/A	N/A	5.3	N/A
powermt unmanage	5.3	5.1 SP2	5.3	5.3	N/A
powermt update lun_names	5.3	5.1 SP2	5.3	5.3	5.3

a. *N/A* means the command, utility or feature is not supported on the platform.

## Message format and description

Audit log messages have the following format:

### On AIX and Linux:

```
<timestamp> <hostname> EMCPP: emcpAudit: Info: cmd=<message> (<user_ID>)
```

For example:

```
Oct 28 13:06:57 Host1a local0:info EMCPP: emcpAudit: Info: cmd=powermt:manage
class=clariion(user ID real=0 effective=0)
```

### On HP-UX and Solaris:

```
<timestamp> <hostname> emcpAudit: [ID <pid> local0.info] CMD: <message>
```

For example:

```
Oct 28 13:06:57 Host1a emcpAudit: [ID 702911 local0.info] CMD: Info Mode of the Array
APM00042002027 port SP B0 is changed to Disabled
```

### On Windows

```
EMC PowerPath Information: emcpAudit: Info: cmd=<message>
```

For example:

```
EMC PowerPath Information: emcpAudit: Info: cmd=powermt: save
```

## Arguments

where:

- ◆ *timestamp* is the date and time the event occurred.
- ◆ *hostname* is the name of the host where PowerPath resides.
- ◆ EMCPP: emcpAudit: is the program performing the auditing function. This is the tag used for audit messages for PowerPath.
- ◆ *pid* is the operating system's process ID and the syslog.conf (4) priority of the message. The message priority includes the facility (kern) and the level (notice).
- ◆ *cmd* is the utility that was used.
- ◆ *message* is the audit message.

- ◆ *user\_ID* is the identification of the user.

## Error log messages

Error log messages appear along with audit messages in the when you configure PowerPath for common logging. The error-log messages capture unexpected events that occur. Some error-log messages convey information that appears on screen when a command fails.

### Message format and description

Error log messages have the following format:

#### On AIX and Linux:

```
<timestamp> <hostname> kernel: <level>:<module>:<message>
```

For example:

```
Sep  7 10:20:45 Host1a kernel: Error:Mpx:Path Bus 5 Tgt 0 Lun 0 to FCNBD054405535 is dead.
```

#### On HP-UX and Solaris:

```
<timestamp> <hostname> <module>: [ID <pid> local0.<module>: <message>
```

For example:

```
Sep  7 10:20:45 Host1a emcp: [ID 801593 kern.notice] Error: Killing bus 3072 to Symmetrix 000190100195 port 13bA.
```

#### On Windows:

```
EMC PowerPath Error: Module: Error: <message>
```

For example:

```
EMC PowerPath Error: emcp_xcryptd: Error: Failed to start the daemon with err 0
```

## Arguments

where:

- ◆ *timestamp* is the date and time the event occurred.
- ◆ *hostname* is the name of the host where PowerPath resides.
- ◆ *level* is the message type (**Info**, **Error**, or **Warning**).
- ◆ *pid* is the operating system's process ID and the syslog.conf (4) priority of the message. The message priority includes the facility (kern) and the level (notice).
- ◆ *module* is the program or module that is reporting the message.
- ◆ *message* is the error message.

# Introduction to PowerPath error messages

PowerPath writes to the system error log a variety of error messages.

## Message prefixes

Some platforms add one or more of the following prefixes to error messages generated by the PowerPath driver:

- ◆ A platform-specific reference to PowerPath:

**Table 35**

Platform	PowerPath error message prefix
Windows	emc
UNIX, Linux	emcp (stands for EMC PowerPath)

- ◆ A reference to the module that logged the message: Mp (multipathing), Pn (pseudo name), or power2 (PowerPath base driver).
- ◆ The message type: Panic, Warning, Error, Info, or Debug.

The above items can be separated by various forms of punctuation, depending on the platform.

Thus, for example, a message on different platforms could have any of the following prefixes:

```
PowerPath:Error
emcp: mod=Mpx
emcp:
emc
```

## Messages for different storage systems

To specify the type of storage system to which an error applies, some error messages include the variable *storage\_system*. This can be, for example, **Symm**.

## Licensing errors

If there is any problem with the license (for example, it is expired, invalid, or cannot be located), or if you have a PowerPath Base license:

- ◆ Certain **powermt set** commands cannot be executed, as explained [“powermt check\\_registration” on page 28](#).
- ◆ **powermt save** cannot be executed on AIX or HP-UX. However, on Linux, Solaris, NetWare, and Windows hosts, you can run **powermt save** without a license.
- ◆ At boot time, the load-balancing and failover policy is set to basic failover or no redirect, as described in [“powermt set policy” on page 99](#).

## Message sources

Error messages can be generated by several commands and utilities. The **Source** column of [Table 36 on page 198](#) lists the following message sources:

- ◆ PowerPath driver — Usually there are driver-related error messages associated with PowerPath messages. For instance, PowerPath marks a path dead if the path test fails, and a path test is done as a result of an I/O failure. These failed I/Os almost always are logged by the HBA driver. These errors can be associated by time and device with PowerPath errors.
- ◆ PowerPath installation — These messages can appear while PowerPath is being installed.
- ◆ powermt and other PowerPath utilities — If a message is generated by a specific powermt command, that command is listed. Otherwise, the message can be generated by any powermt command.
- ◆ emcprep — If a message is generated by a specific emcprep command, that command is listed. Otherwise, the message can be generated by any emcprep command.
- ◆ cfgmgr — As noted in [Chapter 1, “PowerPath Multipathing CLI,”](#) on AIX, for detailed PowerPath error information, run `cfgmgr -vl powerpath0` instead of `powermt config`. The `cfgmgr` command performs the same function as `powermt config` but generates more detailed error messages.

## PowerPath error messages

PowerPath error messages are listed in [Table 36](#).

Windows users can also refer to [Table 37 on page 226](#). That table lists all messages that appear on Windows platforms, along with their IDs.



Windows users who have only message IDs can locate a message's text in [Table 37 on page 226](#), and then use [Table 36 on page 198](#) to get an explanation of and suggested actions for the message.

**Table 36** PowerPath error messages (page 1 of 28)

Message	Platform	Source	Explanation	Suggested action
0301-154 bosboot: missing proto file: /usr/lib/boot/network/ch rp.hdiskpower.proto	AIX	bosboot	After a system boots from a PowerPath device, the <b>bosboot</b> tool cannot function correctly. This is because of the state of the configuration after booting from a PowerPath device and the fact that bosboot expects the boot device to be an hdisk, not an hdiskpower device.	Run <b>pprootdev fix</b> after the system boots using a PowerPath device. Refer to the <i>PowerPath for AIX Installation and Administration Guide</i> for more information.
All paths to 6006016016201A0066A8D722 9B07DC11 are dead	All UNIX, Linux	PowerPath driver	All paths to a volume are dead.	Wait for automatic restore. Otherwise, repair the failed paths and then run the <b>powermt restore</b> command.
All paths to <device_id> are dead	All	PowerPath driver	All paths to a logical device are dead. The logical device failed.	Repair the physical paths. Wait for automatic restoration of the logical device or, to avoid delay, run <b>powermt restore</b> to restore the device. On Solaris 9 hosts with Leadville drivers, add the following line to the /etc/system file: fcps:ssfcps_enable_auto_configuration=1
An unidentified program wants to access your computer. Don't run the program unless you know where it's from or you've used it before. Unidentified publisher	Windows	All MSI-based setup packages on Windows 2008	This message is displayed (at the UAC prompt) when a user attempts to uninstall PowerPath (although the user installed a digitally signed copy of PowerPath setup program).	No action is required.
Array array Port port is Disabled	All UNIX, Linux	PowerPath driver	The specified port has been disabled with the <b>powermt set port_disable</b> command.	To enable the port, use the command <b>powermt set port_disable</b> . Refer to <a href="#">“powermt set port_disable” on page 103</a> for more information.
Array array Port port is Enabled	All UNIX, Linux	PowerPath driver	The specified port is enabled.	To disable the port, use the <b>powermt set port_disable</b> command. Refer to <a href="#">“powermt set port_disable” on page 103</a> for more information.
Assigned volume <clariion_volume> to SP<clariion_ storage_processor>	All	PowerPath driver	Informational, VNX and CLARiiON only: The current SP is inaccessible. PowerPath issues a trespass to the peer SP and switches to the set of paths attached to the peer SP.	Repair physical paths or the SP. Wait for automatic restoration of the volume or, to avoid delay, run <b>powermt restore</b> to restore the volume.

**Table 36** PowerPath error messages (page 2 of 28)

Message	Platform	Source	Explanation	Suggested action
disable/enable is not allowed on this hba type	Linux, UNIX	powermt	The message is returned by 'powermt enable hba=<>' and 'powermt disable hba=<>' command in an unsupported configuration.	No action required. Refer to the section on powermt disable hba in <a href="#">“powermt disable hba” on page 32</a> or <a href="#">“powermt enable hba” on page 76</a> .
ERROR: Cannot open PowerPath. Initialization error	All	powermt	The PowerPath modules are not loaded. This situation can occur when the PowerPath service (on Linux) is stopped.	To view the modules that are installed, use lsmod (modinfo for Solaris). If the PowerPath modules are already installed, start the PowerPath service and then retry the command. Type /etc/init.d/PowerPath start. Refer to the PowerPath Installation and Administration Guide for steps to install the PowerPath modules.
"autostandby iopflimit" must be a number greater than 1 and less than or equal to 100000000	HP-UX, Solaris, Windows	powermt	This message is shown when incorrect value is set for the average number of I/Os between failures.	Ensure that the minimum value is more than 1 and the maximum value is less than or equal to 100000000. You can set the value by using the <b>powermt set autostandby iopflimit</b> command.
Attempted action is not allowed for volumes with BF or NR policy. Verify the policy for the volume <device_name>.	UNIX, Linux	powermt	This message appears when you attempt to disable an HBA or port and the load balancing and failover policy is set to <b>bf</b> or <b>nr</b> . The <b>powermt disable hba</b> and <b>powermt set_port_disable</b> commands fail when the <b>bf</b> or <b>nr</b> policy is set.	Change to another load balancing and failover policy supported with the storage array, and then retry the command to disable an hba or port.
Attempted action is not allowed for volumes with BF or NR policy.	HP-UX, Solaris	powermt	This message appears when you attempt to disable an HBA or port and the load balancing and failover policy is set to <b>bf</b> or <b>nr</b> . The <b>powermt disable hba</b> and <b>powermt set_port_disable</b> commands fail when the <b>bf</b> or <b>nr</b> policy is set.	Change to another load balancing and failover policy supported with the storage array, and then retry the command to disable an hba or port.
Bad dev value <device>, or not under Powerpath control.	Solaris	powermt	You may see this message on a Solaris 10 host when trying to enable with <b>powermt</b> commands a port or HBA. If the Fibre Channel cable for the port or HBA had been previously disconnected and then reconnected, this error may occur.	On a Solaris 10 host, ensure U4 with Solaris patch 122012-1 or later is installed. Run <b>cfgadm -c configure controller#</b> , Enable the port or HBA with the appropriate <b>powermt</b> command.

**Table 36** PowerPath error messages (page 3 of 28)

Message	Platform	Source	Explanation	Suggested action
Bad HBA value HBA#	AIX, Linux, Windows	Powermt	The <b>hba</b> argument is invalid.	Run the command again, specifying a valid HBA number. Use <b>powermt display</b> to list valid HBA numbers.
ERROR: Bad hba value <hba>, or not under PowerPath control.	HP-UX, Solaris	Powermt	The <b>hba</b> argument is invalid.	Run the command again, specifying a valid HBA number. Use <b>powermt display</b> to list valid HBA numbers.
Bus <bus_id> to <storage_system_name> port <port_name> is alive.	Linux, Windows	PowerPath driver	The HBA is alive after a path failure. Or, the HBA was enabled with the <b>powermt enable hba</b> command.	No action is required.
Bus <bus_id> to <storage_system_name> port <port_name> is dead.	Linux, Windows	PowerPath driver	PowerPath proactively killed a path due to a path test failure. Or, the HBA was disabled with the <b>powermt disable hba</b> command (Linux).	Repair the physical path, and run <b>powermt restore</b> to restore it.
Bus hba is Disabled.	All UNIX, Linux	PowerPath driver	The HBA has been disabled using the <b>powermt disable hba</b> command.	To enable the HBA, refer to <a href="#">“powermt enable hba” on page 76.</a>
Bus hba is Enabled.	All UNIX, Linux	PowerPath driver	The HBA is enabled.	To disable the HBA, refer to <a href="#">“powermt disable hba” on page 32.</a>
Cannot assemble drivers for root /pseudo/emcp@5:a,blk Cannot mount root on /pseudo/emcp@5:a,blk fstype ufs	Solaris	PowerPath driver	Booting from an emcpower device has failed.	<p>If booting from an R1 device or an array device:</p> <ol style="list-style-type: none"> <li>1. Boot from the native device.</li> <li>2. Remove the /kernel/drv/emcp.conf file.</li> <li>3. Rebuild the emcp.conf file. Type: powercf -q powermt config</li> <li>4. Make sure the /etc/vfstab and /etc/system files are pointing to the emcpower device used for the boot.</li> <li>5. Reboot from the emcpower device.</li> </ol> <p>If booting from an R2 device:</p> <ol style="list-style-type: none"> <li>1. Do failback.</li> <li>2. Boot from the R1 external device.</li> <li>3. Run the following commands: powercf -Z powermt save</li> </ol>

**Table 36** PowerPath error messages (page 4 of 28)

Message	Platform	Source	Explanation	Suggested action
<pre>&lt;date&gt; &lt;time&gt; &lt;host_name&gt; emcp: [ID 801593 kern.notice] Error: Latency Threshold exceeded by UserDev Bus &lt;bus_number&gt; Tgt &lt;target&gt; Lun &lt;LUN&gt; Vol &lt;volume&gt; with a new high latency of &lt;latency_value&gt; msec  &lt;date&gt; &lt;time&gt; &lt;host_name&gt; emcp: [ID 801593 kern.notice] Error: Latency Threshold exceeded by UserDev C#T#D# Vol &lt;volume&gt; with a new high latency of &lt;latency_value&gt; msec</pre>	HP-UX, Solaris, Windows	powermt	The I/O completion time has exceeded the threshold value.	<p>Estimate expected latencies for your environment and set the value suitable for your environment.</p> <p>To estimate, check the <b>powermt display latency</b> output.</p> <p>Change the threshold as below:</p> <ol style="list-style-type: none"> <li>1. Disable path latency monitoring so that high watermarks are zeroed.</li> <li>2. Set the new threshold.</li> <li>3. Re-enable path latency monitoring with the <b>powermt set path_latency_monitor</b> command.</li> <li>4. Check the I/O completion time by using the <b>powermt display latency</b> command.</li> </ol>
Failover mode PNR for CLARiION LUN <WWN#> is not recommended. Please consider PAR or ALUA	AIX	cfgpower	The Failover mode is set to PNR, which is not recommended.	Set the Failover mode to either PAR or ALUA.
Device(s) not found.	Linux	powermt	There is no device specified in the <b>powermt display perf dev=&lt;dev all&gt;</b> or <b>powermt display perf bus</b> commands.	Specify a device in the <b>powermt display perf dev=&lt;dev all&gt;</b> or <b>powermt display perf bus</b> commands.
Invalid option: class=<class> is not supported for this version	HP-UX, Linux, Solaris, Windows	powermt	The [class=<class>] argument is used in the <b>powermt display perf dev=all</b> command.	The [class=<class>] is not supported, because the default setting is [class=all]. Use the <b>powermt display perf dev=all</b> command without the [class] argument.
Invalid option: dev=all is required for this version	HP-UX, Linux, Solaris	powermt	A device ID value is specified for the [device] argument in the <b>powermt display perf bus</b> command.	The [device=<device ID>] is not supported, because the supported argument is [device=all]. Use the <b>powermt display perf bus</b> command with the [device=all] argument.
Invalid option: class=<class> is not supported for this version	HP-UX, Linux, Solaris, Windows	powermt	A class value is specified for the [class] argument in the <b>powermt display perf bus</b> command.	The [class=<class>] is not supported, because the default setting is [class=all]. Use the <b>powermt display perf bus</b> command without the [class] argument.
Killing bus 4 to Symmetrix 000192601699 port 7fB	All UNIX, Linux	powermt	A bus to a storage port is dead.	Wait for automatic restore. Otherwise, repair the failed paths and then run the <b>powermt restore</b> command.

**Table 36** PowerPath error messages (page 5 of 28)

Message	Platform	Source	Explanation	Suggested action
Not all paths to 6006016016201A00A32708D2 3C2EDC11 are dead	All UNIX, Linux	powermt	Some paths to volume are dead.	Wait for automatic restore. Otherwise, repair the failed paths and then run the <b>powermt restore</b> command.
Path Bus 4 Tgt 1 Lun 3 to APM00064303132 is alive	All UNIX, Linux	powermt	A path to a device is dead.	Wait for automatic restore. Otherwise, repair the failed paths and then run the <b>powermt restore</b> command.
Path Bus 4 Tgt 1 Lun 3 to APM00064303132 is alive	All UNIX, Linux	powermt	A path to a device has come alive.	No action is required.
Path is in Autostandby, force option required.	HP-UX, Linux, Solaris, Windows	powermt	Attempt to change the autostandby mode of a path failed, because the required <b>force</b> option is not used.	Override the autostandby mode with the <b>force</b> option.
<path> is in Autostandby, force option required.	HP-UX, Linux, Solaris, Windows	powermt	Attempt to change the mode of a path group, in which some paths are in autostandby mode, has failed. The absence of the <b>force</b> option has caused the failure.	
Performance monitoring already enabled. To change interval, set off first.	HP-UX, Linux, Solaris, Windows	powermt	The <b>powermt set</b> command is run on the system where the performance monitoring is already enabled.	<ol style="list-style-type: none"> <li>1. Run <b>powermt set =off</b> to disable performance monitoring.</li> <li>2. Run <b>powermt set =on interval=&lt;n&gt;</b> to change the performance monitoring interval.</li> </ol>
Performance measurement is not enabled" & their powermt display perf command exits	HP-UX, Linux, Solaris	powermt	The <b>powermt display performance</b> command is run when performance monitoring is disabled by another user.	<p>Before enabling and using the performance monitoring, check the required settings of your environment; and, contact the user who has disabled the performance monitoring regarding why monitoring is disabled.</p> <p>Enable the performance monitoring.</p>
PRO service action 6 failed, <Sense Key>/<ASC>/<ASCQ> for Bus <bus> Tgt <target> Lun <lun>	All	powermt	PRO command failed for the specific device shown in the message.	If Sense Key != 6 (unit attention, informational), report the values to support.
Reviving bus 4 to Symmetrix 000192601699 port 7fB	All UNIX, Linux	powermt	A bus to a storage port has come alive.	No action is required.

**Table 36** PowerPath error messages (page 6 of 28)

Message	Platform	Source	Explanation	Suggested action
Cannot open /dev/emcp, errno <number>	Solaris	powermt config	<b>errno 13:</b> By default, the PowerPath installation process sets file permissions such that only the superuser can run <b>powermt</b> . <b>errno other than 13:</b> A required file cannot be opened, for some reason other than access permission.	<b>errno 13:</b> Log in as superuser. Then run <b>powermt config</b> . <b>errno other than 13:</b> Report the error number to EMC Customer Support.
Cannot open dev dir for reading, errno <number>	Solaris	powermt config		
Cannot open <filename>, errno <number>	HP-UX, Solaris	powermt config		
/etc/powermt: Permission denied	HP-UX Solaris	powermt config	A required file could not be opened.	1. Log in as super user. 2. Run <b>powermt config</b> .
Cannot open string Initialization error	All	powermt	If string is a device file, PowerPath is not running or is running without permission to access the device.	Start PowerPath or run the command as root.
/etc/powermt: Permission denied	HP-UX Solaris	powermt	PowerPath attempts to open a device file without required permissions.  If <i>string</i> is a device file, PowerPath is not running or is running without permission to access the device.	Start PowerPath or run the command as root.
Cannot remove alive device device_id	HP-UX, Linux	powermt check, powermt remove	An application or mounted filesystem is using the device.	Stop the application or unmount the filesystem that is using the device, and run <b>powermt check</b> or <b>powermt remove</b> again.
Cannot remove device <device_id>, name in use by application.	Solaris			
Can't get Pirp_t!	AIX	powermt	Error: There was an error allocating a request buffer.	Contact EMC Customer Support.
command:illegal option -- option (or command: option requires an argument -- option)	All	emcp	An incomplete command was entered.	Ensure you are entering the correct command syntax.
Current/Active controller for volume has changed.	All	powermt	The controller that owns the volume has changed. PowerPath was unable to switch to the set of paths attached to the new owning controller.	Repair physical paths to the new owning SP. Then run <b>powermt restore</b> to restore the paths as soon as possible. Until paths are restored, performance may be degraded significantly.
Device is unavailable	All	PowerPath I/O	Device is inaccessible (system log may indicate whether message refers to source or target).	Bring problem device back online.
Device(s) not found.	All	All <b>powermt</b> commands with both the <b>class</b> and <b>dev</b> arguments	<b>class=all</b> was specified, and no storage-system devices were found.	No action is required.

**Table 36** PowerPath error messages (page 7 of 28)

Message	Platform	Source	Explanation	Suggested action
Device not found.	All	powermt	The dev argument is invalid.	Run the command again, specifying a valid device name. Use <code>powermt display dev=&lt;device&gt;</code> to list valid device names.
ERROR: Bad dev value <device>, or not under Powerpath control.	HP-UX Solaris	powermt	The dev argument is invalid.	Run the command again, specifying a valid device name. Use <code>powermt display dev=&lt;device&gt;</code> to list valid device names.
<device_id> is alive.	All	powermt	Warning: The specified device was killed at some point, due to a critical I/O error. The device was just re-enabled, so PowerPath no longer inhibits I/O to it. The prior I/O error may have caused data on the device to become inconsistent.	Run a data integrity checker to verify the integrity of the data. Fix any inconsistencies. This must be done before re-enabling the application using the device, to avoid worsening any integrity problem and potentially making the data on the device unrecoverable.
<device_id> is dead.	All	powermt	Error: The logical device failed.	Wait for automatic restoration of the logical device or, to avoid delay, run <b>powermt restore</b> to restore the device.
EMCP registration file does not exist. Use -add to create it.	All	emcprep -edit	There is no registration file.	Use <b>emcprep -add</b> to add a registration key, which will cause a registration file to be created.
EMCpower installation: Configuration files successfully converted.	Solaris	PowerPath installation	PowerPath 3.0.x (or earlier) configuration files have been converted to 4.x format.	No action is required.
EMCpower installation: Error!!! Configuration files conversion failed!!!	Solaris	PowerPath installation	The upgrade failed because the PowerPath 3.0.x (or earlier) configuration files could not be converted to the current format. Your custom configuration is no longer available.	Reconfigure any applications that are configured with PowerPath pseudo devices.
PowerPath cannot disable some of its devices because a page file is located in a LUN under PowerPath control or host is configured to boot from SAN. Under these scenarios, two host reboots are required to uninstall PowerPath. Click OK to proceed, or Cancel to exit.	Windows	PowerPath uninstallation	A page file located in a LUN is under PowerPath control or the host is configured to boot from SAN.	Two host reboots required.
EMCP registration file is corrupt.	All	emcprep	The license registration file was corrupted and cannot be used.	Contact EMC Customer Support.

**Table 36** PowerPath error messages (page 8 of 28)

Message	Platform	Source	Explanation	Suggested action
EMCP registration file is corrupt. Delete it?	All	emcprep -edit	The registration file was corrupted and cannot be used.	If desired, type <b>y</b> to delete the file, and then add back your registration keys. (Adding the first key will re-create the file.) Otherwise, type <b>n</b> and contact EMC Customer Support.
EMC registration data is missing or corrupt.	All	powermt check_registration	The registration file is missing, corrupt, or open in an editor.	Verify the existence of the registration file. If the file is open, close the editor and run the command again.
There are no license keys now registered.	HP-UX Solaris	powermt check_registration	The registration file is missing, corrupt, or open in an editor.	Verify the existence of the registration file. If the file is open, close the editor and run the command again.
Encountered unexpected error 0x<code> on Dek operation command for <device id> on handle <association_handle>.	Windows	powermt	There was a problem with the DEK Management protocol.	Contact EMC Customer Support.
WARNING!!! ODM PdDv database contains duplicate uniquetypes. Refer to EMC Knowledge Base Article Number: 000084491 or contact EMC Customer Support personnel.	AIX	powermt	Multiple entries are present for an uniquetype in PdDv	Contact EMC Customer Support.
Error number adding pncpath <hdisk_name> for <hdiskpower_name> to CuAt.	AIX	cfgmgr	The system cannot add to the <b>CuAt</b> ODM table.	Remove hdiskpower devices and run <b>powermt config</b> . Check the ODM state. Verify the root filesystem has sufficient free space.
Error number getting pncpath attribute for <hdiskpower_name>.	AIX	cfgmgr	The system cannot find the definition for a PowerPath ODM object.	Reinstall PowerPath.
Error: class is not a valid class.	All	powermt	The class value entered is invalid.	Check the syntax and value of the specified class. Valid values are <b>symm</b> , <b>vnx</b> , <b>clariion</b> , <b>vplex</b> , <b>invista</b> , <b>ess</b> , <b>generic</b> , <b>hitachi</b> , <b>hpxp</b> .
Error number updating CuDv attribute for <device_id>.	AIX	cfgmgr	The system cannot update the <b>CuDv</b> ODM table.	Remove hdiskpower devices and run <b>powermt config</b> . Check the ODM state. Verify the root filesystem has sufficient free space.
Error occurred loading saved driver state from file /etc/powermt_custom.xml Loading continues.	AIX, Linux	powermt load	If the custom file has information of the volumes which are not managed by PowerPath, e.g. using 'powermt unmanage dev=<>' or after removing luns from the host, the error message is displayed.	Contact EMC Customer Support.



**Table 36** PowerPath error messages (page 9 of 28)

Message	Platform	Source	Explanation	Suggested action
Error occurred loading saved driver state from file filename. Loading continues.	All	powermt load	The configuration file is corrupted, or there was an I/O error adding a device in the file.	Contact EMC Customer Support.
Error occurred loading saved driver state from file filename. Loading continues.	HP-UX, Solaris	powermt load	There is a problem reading the configuration file, because of insufficient permissions, corrupted configuration file, or there was an I/O error while adding a device in the file.	<ul style="list-style-type: none"> <li>• Check the file permissions.</li> <li>• Check validity of the configuration file.</li> </ul>
Error: All paths to vol volume are dead.	Solaris	powermt	In a Solaris Cluster with Hitachi Lightning, Hitachi TagmaStore, or HP XP arrays:	
			Unregistered cluster disk paths are dead after the node has booted into the cluster.	Run the <b>powermt display</b> command to check path status. If paths are dead, check the physical connection and then reboot the cluster node.
			In all other configurations:	
			All paths to a volume have failed.	Check the physical path state and run <b>powermt restore</b> . If <b>powermt restore</b> fails to restore the paths, perform a reconfiguration reboot.
Error: An older EMCpower package is already installed. Please exit the installation and remove the old EMCpower package first.	Solaris	PowerPath installation	An older version of PowerPath is already installed on the host. This aborts the installation.	Remove the earlier version of PowerPath and begin the installation again.
Error: Attempted action for <storage_array> storage system support is not allowed with the current license level.	HP-UX, Solaris, Windows	powermt	The PowerPath license is invalid or has expired. <b>powermt</b> requires a valid license.	Add a valid PowerPath license. Run <b>powermt check_registration</b> to determine the problem with the PowerPath license.
ERROR: License not installed.	HP-UX, Solaris	powermt	The PowerPath license is invalid or has expired. <b>powermt</b> requires a valid license.	Add a valid PowerPath license. Run <b>powermt check_registration</b> to determine the problem with the PowerPath license.
error - attempted upgrade from an unsupported version.	All UNIX	emcpupgrade	You cannot upgrade to PowerPath 5.2 from this version of PowerPath.	Upgrade to PowerPath 4.5 or later. Then, upgrade to PowerPath 5.2.
error - cannot generate new configuration	All UNIX	emcpupgrade	emcpupgrade could not generate new configuration files.	<p>Check whether the older version of the PowerPath driver is loaded. If the driver is loaded, run <b>rem_drv</b> to it.</p> <p>If this fails:</p> <ol style="list-style-type: none"> <li>1. Remove PowerPath 5.0.</li> <li>2. Reboot the host.</li> <li>3. Install PowerPath 5.0 again.</li> </ol>

**Table 36** PowerPath error messages (page 10 of 28)

Message	Platform	Source	Explanation	Suggested action
error - cannot install the new configuration files	All UNIX	emcpupgrade	<b>emcpupgrade</b> could not install new configuration files.	Ensure the /etc directory is not full.
Error: cannot manage. verify state of the class	All UNIX, Linux	powermt manage class	The status of the specified class is unmanage_incomplete.	Before issuing the <b>powermt manage class=&lt;class&gt;</b> command, do one of: <ul style="list-style-type: none"> <li>• Close all the applications using the devices belonging to the storage class &lt;class&gt; and issue the <b>powermt unmanage class=&lt;class&gt;</b> command.</li> <li>• Reboot the host.</li> </ul>
Error: cannot obtain driver state	All UNIX, Linux	PowerPath installation, powermt	<b>Installation:</b> You are making kernel modifications or installing kernel patches post-install. <b>powermt:</b> No storage-system logical devices are configured, one or more PowerPath drivers is not running, or the PowerPath control device file may not exist.	Ensure the driver's device file exists. If it does, make sure the user has permission to access the driver through that file.
Error: cannot open filename.	All	powermt save	The file is missing, the user does not have the correct access to read the file, or there was a filesystem error.	<ol style="list-style-type: none"> <li>1. Check the state of the configuration file and the filename's spelling.</li> <li>2. Check the files read and write permissions, and log in as superuser if needed.</li> <li>3. Run <b>powermt save</b>.</li> </ol>
Length of the filename exceeds the maximum limit of: <number> characters	HP-UX, Solaris	powermt	Filename length has exceeded the set limit.	Provide a filename within the expected length limitation.
Cannot write to <filename>	HP-UX, Solaris	powermt	The file is missing, the user does not have the correct access to read the file, or there was a filesystem error.	<ol style="list-style-type: none"> <li>1. Check the state of the configuration file and the filenames spelling.</li> <li>2. Check the files read and write permissions, and log in as superuser if needed.</li> <li>3. Run <b>powermt save</b>.</li> </ol>
Renaming of temporary file to file:<filename> failed	HP-UX, Solaris	powermt		
Cannot obtain lock on <filename>	HP-UX, Solaris	powermt	Multiple powermt sessions are occurring; unable to get lock.	Avoid multiple powermt sessions.

**Table 36** PowerPath error messages (page 11 of 28)

Message	Platform	Source	Explanation	Suggested action
error - cannot open configuration file filename.	All UNIX	emcpupgrade	<b>emcpupgrade</b> could not open the specified configuration file.	<ol style="list-style-type: none"> <li>1. Change (<b>su</b>) to root.</li> <li>2. Run <b>emcpupgrade</b>.</li> <li>3. Check that the specified configuration file <i>i</i> exists.</li> <li>4. If the file does not exist, contact EMC Customer Support.</li> </ol>
error - cannot parse PowerPath 4.x configuration.	All UNIX	emcpupgrade	<b>emcpupgrade</b> could not parse the existing files. Some of the configuration files might be corrupted.	Contact EMC Customer Support.
error - cannot produce new files.	All UNIX	emcpupgrade	<b>emcpupgrade</b> could not produce new configuration files.	Ensure that the /tmp directory is not full.
Error: cannot write to file.	All	powermt save	The user may not have permission to write to the file.	Check the file's write permissions, and log in as superuser if needed. Run <b>powermt save</b> again.
error - can't open /etc/path_to_inst. Cannot continue.	All UNIX	emcpupgrade	<b>emcpupgrade</b> could not open the system configuration file <b>path_to_inst</b> .	Repair the system and then run <b>emcpupgrade</b> again.
Error: class=<class> is not valid.	All	powermt	The class is invalid.	Run the command again, specifying a valid <class>. Alternately, run <b>powermt help</b> .
ERROR: Unsupported class name: <class>.  class={all symm vnx clarion vplex invista xtremio hitachi hpxp ess generic	All	powermt		
error - configuration database corrupted.	All UNIX	emcpupgrade	<b>emcpupgrade</b> could not read the configuration database.	Contact EMC Customer Support.
error - configuration file filename is corrupted.	All UNIX	emcpupgrade	The specified configuration file is corrupted. Upgrade is not possible.	Contact EMC Customer Support.
error - configuration files corrupted.	All UNIX	emcpupgrade	The configuration files are corrupted. Upgrade is not possible.	Contact EMC Customer Support.
error - configuration files corrupted/out of date (dev=device).	All UNIX	emcpupgrade	The configuration files are corrupted, or they do not match the present configuration on the system. To avoid any possibility of data corruption, the upgrade is not possible. Upgrade could cause data corruption and therefore will not take place.	Contact EMC Customer Support.

**Table 36** PowerPath error messages (page 12 of 28)

Message	Platform	Source	Explanation	Suggested action
error - currently installed powermt.custom has inconsistent version.	All UNIX	emcpupgrade	<b>emcpupgrade</b> does not recognize the version number in the /etc/powermt.custom configuration file.	Ensure the <b>emcpupgrade</b> and PowerPath version numbers match.
Error: error reading configuration file.	All	powermt load	An error occurred while reading the configuration file.	Run <b>powermt load</b> again. Also check the files read and write permissions, and log in as superuser if needed.
Error: custom file IO failure	All	powermt load		
ERROR: error retrieving minor number(s)	All UNIX	emcpminor	Either the base number you specified is too large or an entry exists in /devices/pseudo that corresponds to a free minor number.	Rerun <b>emcpminor</b> , specifying a base number in the allowable range. If an error still occurs, for each device in /devices/pseudo, run the <b>emcpminor check</b> command, supplying that device's corresponding minor number. If an error occurs, remove that device from /devices/pseudo. If the error persists, contact EMC Customer Support.
Error: failed to update CLARiiON user assignable LUN names.	All	powermt update lun_names	The command failed.	Contact EMC Customer Support.
error - incomplete installation detected.	All UNIX	emcpupgrade	<b>emcpupgrade</b> cannot find some files/parameters that should have been set when <b>pkgadd</b> was run.	Ensure <b>pkgadd</b> completed successfully before you run <b>emcpupgrade</b> .
error - internal DB read failed	All UNIX	emcpupgrade	The configuration database is corrupted.	Contact EMC Customer Support.
ERROR: invalid base minor number	All UNIX	emcpminor	The base minor number specified is not in a valid format (numeric) or is greater than the maximum possible minor number.	Rerun <b>emcpminor</b> , specifying as the base minor number a numeric value in the allowable range.
ERROR: invalid minor number	All UNIX	emcpminor	The minor number you specified is not in a valid format (numeric), is greater than the maximum possible minor number, or is thought to be free and has an entry in /devices/pseudo.	Rerun <b>emcpminor</b> , specifying as the minor number a numeric value in the allowable range of values. If an error still occurs, check /devices/pseudo for a device corresponding to that minor number, and remove that device. If the error persists, contact EMC Customer Support.
Error: license for <storage_system_name> storage system support is missing or expired.	All	powermt	The PowerPath license is invalid or has expired, and the <b>powermt</b> command that was run requires a valid license.	Run <b>powermt check_registration</b> to determine the problem with the PowerPath license.
ERROR: License not installed.	All	powermt		

**Table 36** PowerPath error messages (page 13 of 28)

Message	Platform	Source	Explanation	Suggested action
Error: Major number is already in use.	Solaris	PowerPath installation	When installing PowerPath for a clustered environment, you chose a major number that is already in use. This aborts the installation.	Reinstall PowerPath, specifying a unique major number.
Error: manage of device <device> failed.	All UNIX, Linux	powermt manage dev	<b>powermt manage</b> failed to place the specified device under PowerPath control.	Check the status of the class using <b>powermt display options</b> . If the status is <code>manage_incomplete</code> or <code>unmanage_incomplete</code> , the status will be updated after you close all applications using devices belonging to the class, or after you reboot the host. You can then re-issue the <b>powermt manage dev</b> command. For other statuses, the device might be unreachable.
error - mismatch in configuration (device device).	All UNIX	emcpupgrade	A mismatch was detected. SCSI inquiry for the specified device does not match the saved parameters. It might be that the configuration has changed since the last PowerPath installation. In this case, upgrade is impossible.	If PowerPath 3.0.x or 2.1 configuration files are corrupted, your custom configuration is no longer available after you install PowerPath 5.0. If any of your applications are configured with PowerPath pseudo devices, you must reconfigure those applications.
error - mismatch in configuration.	All UNIX	emcpupgrade	A mismatch was detected. The configuration files are not in sync.	Contact EMC Customer Support.
error - missing configuration files.	All UNIX	emcpupgrade	Some of the configuration files are not present. Cannot complete the action.	Contact EMC Customer Support.
error - mknod of device failed.	All UNIX	emcpupgrade	<b>emcpupgrade</b> was unable to create the emcpower device identified by <device>.	
ERROR: not a multiple of 8.	All UNIX	emcpminor	The minor number you specified is not a multiple of 8.	Rerun <b>emcpminor</b> , specifying a minor number that is a multiple of 8.
error - not enough memory.	All UNIX	emcpupgrade	Memory error.	Contact EMC Customer Support.
ERROR: number of minors out of range 1-1000.	All UNIX	emcpminor	The number of minor numbers specified for retrieval is out of the specified range.	Rerun <b>emcpminor</b> , specifying as the number of minor numbers to be retrieved a number in the range 1–1000.
Error: Path latency threshold value <number> is not in range (0 - 3600) or (0 - 3600000ms).	All	powermt set path_latency_threshold	The threshold value you specified is out of the specified range.	Run <b>powermt set path_latency_threshold</b> again, specifying a valid threshold value.

**Table 36** PowerPath error messages (page 14 of 28)

Message	Platform	Source	Explanation	Suggested action
ERROR: Path latency threshold value <value> is not in range (0 - 2147483647)	All	powermt set path_latency_threshold	The threshold value (in milliseconds) you specified is out of the specified range.	Run <b>powermt set path_latency_threshold</b> again, specifying a valid threshold value (in milliseconds).
Error: policy is not a valid policy for <storage_system_class> storage systems.	All	powermt set policy	The specified load-balancing and failover policy is invalid for the specified class of storage system.	Run <b>powermt set policy</b> again, specifying a valid policy/ storage-system class combination.
Error: powercf -i is obsolete.	Solaris	powercf	The <b>-i</b> option is no longer supported.	Refer to the <b>powercf</b> man page for supported options.
Error: powercf -p is obsolete.	Solaris	powercf	The <b>-p</b> option is no longer supported.	Refer to the <b>powercf</b> man page for supported options.
error - PowerPath configuration files present - exiting. For further information consult the release notes.	All UNIX	emcpupgrade	You either: <ul style="list-style-type: none"> <li>Upgraded once and then tried to upgrade again.</li> <li>Rebooted the host after you installed PowerPath but before you ran emcpupgrade.</li> </ul> Either of these actions would result in a valid PowerPath configuration on the host, and you cannot use <b>emcpupgrade</b> to convert files while there is valid PowerPath configuration on the host.	<ol style="list-style-type: none"> <li>Uninstall PowerPath.</li> <li>Copy the saved configuration files from the archive directory to the appropriate directory. (Refer to “If conversion fails” in the <i>PowerPath Installation and Administration Guide</i> for details.)</li> <li>Reinstall PowerPath.</li> </ol>
Error: PowerPath is not installed correctly.	All UNIX, Linux	powercf	PowerPath is missing a required file.	Linux: cannot find /etc/opt/emcpower/emcpmgr Linux: cannot find /etc/opt/emcpower/powercf Solaris: cannot find /etc/powercf Solaris: cannot find /etc/emc/bin/emcpmgr HP-UX, cannot find /sbin/powercf HP-UX, cannot find /sbin/emc/emcpmgr
error - read of file <file> failed.	All UNIX	emcpupgrade	File read failed.	Ensure permissions are set correctly on the specified file.

**Table 36** PowerPath error messages (page 15 of 28)

Message	Platform	Source	Explanation	Suggested action
error - SCSI error on device <device>.	All UNIX	emcpupgrade	Could not perform SCSI inquiry on the device. The device might be offline.	Check that all devices are online. Bring any offline devices online. If some of the paths have been removed from the system but not from the previously installed version of PowerPath: <ol style="list-style-type: none"> <li>1. Uninstall the new package.</li> <li>2. Reinstall the old package.</li> <li>3. Remove the dead paths from PowerPath using the command <b>powermt remove dev=&lt;dev&gt;</b>.</li> <li>4. Reinstall the new package.</li> </ol>
Error: seconds are not in the range (1-86400).	All	powermt display every=seconds	The refresh value is not in the valid range.	Run the command again, specifying a <i>seconds</i> value within the valid range.
ERROR: "every" value <value> is not in range (1 - 86400).	All	powermt display every=seconds	The refresh value is not in the valid range.	Run the command again, specifying a <i>seconds</i> value within the valid range.
ERROR: StreamIO threshold value <value> is not in range (64 - 2048).	All	powermt set streamio_threshold =threshold_count	The refresh value is not in the valid range.	Run the command again, specifying a <b>threshold_count</b> value within the valid range.
Error: The prerequisite patch(es) are not installed - please install them after installing the PowerPath package. Required patch(es): patch_number Continue installation? [y,n,q,?] (default: y):	Solaris,	PowerPath installation	The specified required patch is not installed on the host.	Either: <ol style="list-style-type: none"> <li>1. Continue the installation of PowerPath 5.0.</li> <li>2. Install the required patch immediately after the installation completes.</li> </ol> Or: <ol style="list-style-type: none"> <li>1. Abort the installation</li> <li>2. Install the required patch.</li> <li>3. Install PowerPath 5.0.</li> </ol>
Error: unmanage of device <device> failed.	All UNIX, Linux	powermt unmanage dev	A failure occurred while unmanaging the specified device.	Check the status of the class using <b>powermt display options</b> . If the status is <i>manage_incomplete</i> or <i>unmanage_incomplete</i> , the status will be updated after you close all applications using devices belonging to the class, or after you reboot the host. You can then re-issue the command. For other statuses, the devices might be unreachable.
error - unsupported configuration (R1/R2 setup detected)	All UNIX	emcpupgrade	Upgrade cannot take place on R1/R2-enabled systems.	Contact EMC Customer Support.
error - unsupported HBA.	All UNIX	emcpupgrade	The HBA on the system is not supported by PowerPath.	Refer to the E-Lab Interoperability Navigator for supported HBAs.

**Table 36** PowerPath error messages (page 16 of 28)

Message	Platform	Source	Explanation	Suggested action
error - upgrade from version <version> not supported.	All UNIX	emcpupgrade	Upgrade from the specified version is not supported.	Upgrade to PowerPath 3.x or later. Then upgrade to PowerPath 5.0.
error - validation failed. For further information consult the release notes.	All UNIX	emcpupgrade	Validation of the configuration files failed.	Look at the error messages that precede this one in the output, and locate them in this table.
Error: write throttle queue depth is not in range (1-2147483647).	All UNIX	powermt set write_throttle	The specified queue depth value is outside the valid range or invalid (for example, a non-integer value).	Run <b>powermt set write_throttle_queue</b> , specifying a valid <i>queue_depth</i> value. Refer to <a href="#">“powermt set write_throttle_queue” on page 109</a> .
ERROR: Write throttle queue depth value <value> is not in range (0 - 2048).	All UNIX	powermt set write_throttle	The specified queue depth value is outside the valid range or invalid (for example, a non-integer value).	Run <b>powermt set write_throttle_queue</b> , specifying a valid <i>queue_depth</i> value. Refer to <a href="#">“powermt set write_throttle_queue” on page 109</a> .
Failed opening adapter <device_number>.	AIX	powermt	Error: There was a failure opening a bus adapter device.	Check the adapter and bus.
Failed to load extension errno = <number>.	AIX	cfgmgr	The system cannot load the specified PowerPath extension.	Check the operating system level and PowerPath version. Verify the extension file is in <b>/usr/lib/drivers</b> . Check memory usage.
File deleted. Use -add to create a new file.	All	emcprep -edit	You chose to delete a corrupt registration file. This message appears after the <b>EMCP registration file is corrupt</b> message, if you type <b>y</b> to delete the file.	To use PowerPath again, add a registration key.
File not deleted. No licenses will be available until file is deleted and re-created.	All	emcprep -edit	You chose not to delete a corrupt registration file. This message appears after the <b>EMCP registration file is corrupt</b> message, if you type <b>n</b> to not delete the file.	No action is required. If you want to delete the corrupt file, run <b>emcprep -edit</b> again.
Get adapter <devno> failed number	AIX	cfgmgr	The system cannot determine the device number for the <b>hdisk</b> parent.	Remove hdiskpower devices and run <b>powermt config</b> . Check the ODM state
Get hdisk path data failed	AIX	cfgmgr	The system cannot get <b>hdisk</b> path data from the ODM.	
Get hdiskpower CuAt failed number	AIX	cfgmgr	The system cannot get CuAt for an hdiskpower device.	
Get hdiskpower CuDv failed number	AIX	cfgmgr	The system cannot get CuDv for an hdiskpower device.	



**Table 36** PowerPath error messages (page 17 of 28)

Message	Platform	Source	Explanation	Suggested action
Getting unmanaged devices failed with error.	All UNIX, Linux	powermt	The command failed to retrieve unmanaged devices.	There may be problems with some of the unmanaged devices. Check the configuration and connectivity of devices.
Information: An older version of the emcp driver is running. This installation will replace the older version.	Solaris	PowerPath installation	An earlier version of the PowerPath driver is running on the host.	No action is required. Rebooting the host after installation will replace the earlier version of the driver.
Info: Not all paths to vol volume are dead.	Solaris	powermt	In a Solaris Cluster with Hitachi Lightning, Hitachi TagmaStore, or HP xp arrays:	Run the <b>powermt display</b> command to check path status. If paths to the specified volume(s) are dead, try restoring them by running the <b>powermt restore</b> command. If the paths are still dead, check the physical connection and then reboot the host.
			Cluster disk paths that were temporarily dead after the node booted into a Solaris Cluster have been restored.	
			In all other configurations:	Some paths to a volume that was formerly in the AllPathsDead state have been restored.
Information: One or more expired PowerPath license key(s) detected. These are typically used for evaluation. You can use emcpreg utility to review your license keys.	Solaris	PowerPath installation	Your PowerPath evaluation license has expired.	Use <b>emcpreg</b> to add your new license key.
Info: Unable to unconfigure path Bus X Tgt Y Lun Z as firstpath policy is set.	Solaris	powermt remove	You tried to remove the first path an application opened to a multipathed volume. In certain situations (for example, when early versions of VxVM are running without DMP), removing this first path can cause data corruption and thus is prohibited.	If you are sure you want to remove the path, rerun the <b>powermt remove</b> command using the <b>firstpath</b> option: <b>powermt remove firstpath dev=device.</b>
Invalid instance number <number> in config file	Solaris	powermt config	Something is wrong with the configuration file.	Inspect /kernel/drv/ emcp.conf. If it appears corrupted, contact EMC Customer Support.
Iocmd mp.1 failed; wrong buffer size. Driver state not found.	Solaris	powermt	These two messages often occur together during an upgrade to PowerPath 5.0, when drivers from previous packages remain loaded after the new package is installed.	Reboot the system to complete the upgrade.

**Table 36** PowerPath error messages (page 18 of 28)

Message	Platform	Source	Explanation	Suggested action
Iocmd power2.# failed; reason.	All	powermt	The I/O command failed for the specified reason.	Recommended actions depend on the reason given in the error message.
Key <key_number> is already present, ignored.	All	emcprep -add, emcprep -install	The specified key is already registered.	Check the key you intended to register, and reenter it if it differs from what you specified.
Key <key_number> is invalid.	All	powermt check_registration  For all platforms: emcprep -add, emcprep -check, emcprep -install, emcprep -list	The specified key is invalid.	Run the command again, specifying a valid key.
Key <key_number> was not found.	All	emcprep -remove	The specified key is not in the registration file.	Enter a different key to be removed.
Key successfully installed.	All	emcprep -install	A new key was added to the registration file. This appears each time a key is added.	No action is required.
License file is full -- cannot add number keys.	All	emcprep -add, emcprep -install	The registration file cannot hold any more keys.	Identify and delete unneeded or unused keys, and then add the new key.
Missing option parameter.	All	emcprep -add, emcprep -check, emcprep -remove	No <i>key</i> parameter was specified, for a command option that requires one.	Reenter the command, specifying the <i>key</i> option.
MpEnable error number.	AIX	cfgmgr	The system cannot enable processing for PowerPath's multipath extension.	Verify that <code>/usr/lib/drivers/mpext</code> exists. Reboot.
MpSet error number.	AIX	cfgmgr	There is a configuration failure of PowerPath's multipath extension.	Verify that <code>/usr/lib/drivers/mpext</code> exists. Reboot.
No unmanaged devices found.	All UNIX, Linux	powermt	There are no unmanaged devices.	No action is required.
No unmanaged devices found for class class.	All UNIX, Linux	powermt	There are no unmanaged devices for the specified array class.	No action is required.
Not all paths to <device_id> are dead.	All	powermt	Informational: This is issued when a path comes alive, after all paths to the logical device were dead.	No action is required.
NOTE: There are no keys presently registered. Some functions and capabilities may be unavailable.	All	emcprep -install	No key is registered; that is, no key was added or upgraded, or the last key was removed.	Add a valid key.
number key(s) successfully added.	All	emcprep -install	The specified number of keys were registered. This appears at the end of the registration process.	No action is required.

**Table 36** PowerPath error messages (page 19 of 28)

Message	Platform	Source	Explanation	Suggested action
number key(s) successfully registered.	All	emcprep -install	The specified number of keys were registered.	No action is required.
Open powerpath0 failed number.	AIX	cfgmgr	The system cannot open the powerpath0 device.	Verify that /dev/powerpath0 exists and has correct permissions.
/opt/emcpower/ powercf: you have to be root to execute this program.	Linux	powermt config	By default, the PowerPath installation process sets file permissions such that only the superuser can run <b>powermt</b> . ( <b>powercf</b> is run automatically when the user runs <b>powermt config</b> .)	Log in as superuser. Then run <b>powermt config</b> .
/etc/powermt: Permission denied  /opt/emcpower/powercf does not exist	HP-UX, Solaris	powermt config	By default, the PowerPath installation process sets file permissions such that only the superuser can run <b>powermt</b> . ( <b>powercf</b> is run automatically when the user runs <b>powermt config</b> .)	Log in as superuser. Then run <b>powermt config</b> .
PartMgr N/A USWFP0008A Disk X will not be used because it is a redundant path for disk Y.	Windows		Disk Y and disk X is the VCMDDB. The VCMDDB has a disk signature written to it. PowerPath does not manage the VCMDDB. Windows 2003 sees that there are two drives with the same signature. It attempts to change the signature on one of the drives. It then reads the signature of the first drive again, and if the signature is the same, it disables the second path to the device. This message is recreated every time the device is remounted.	None. This is the expected behavior.
Path <path_name> to <device_id> is alive.	All	powermt	Informational: A path's state transitioned from dead to alive.	No action is required.
Path <path_name> to <device_id> is dead.	All	powermt	Error: A path's state transitioned from alive to dead.	Repair the physical path.
Pioc daemon: creating number pseudos.	Solaris	powermt	Informational: The specified number of logical devices are configured.	No action is required.
policy value is missing or invalid.	All	powermt	A policy that is already obsolete is specified or no value is specified for the policy.	Enter value of a policy that is still valid.

**Table 36** PowerPath error messages (page 20 of 28)

Message	Platform	Source	Explanation	Suggested action
<pre>ERROR: Unsupported policy:nr.  Usage: powermt set policy={ad co lb li re r r so si}  [dev=&lt;device&gt; all   class=&lt;class&gt; all]</pre>	All	powermt	A policy that is already obsolete is specified or no value is specified for the policy.	Enter value of a policy that is still valid.
<pre>Power Daemon &lt;process_id&gt; - exiting on signal &lt;signal_number&gt;.</pre>	AIX	powermt	Informational: The PowerPath daemon is exiting.	No action is required.
PowerPath driver loaded.	Windows	powermt	Informational.	No action is required.
PowerPath is not installed.	Linux	PowerPath installation	You are trying to uninstall PowerPath without a prior installation.	No action is required.
PowerPath modules are not loaded. Ensure that the modules are loaded and try again.	Linux, Solaris	powermt	The PowerPath modules are not loaded. This situation can occur when the PowerPath service (on Linux) is stopped.	To view the modules that are installed, use <b>lsmod (modinfo</b> for Solaris). If the PowerPath modules are already installed, start the PowerPath service and then retry the command. Type <b>/etc/init.d/PowerPath start</b> Refer to the appropriate PowerPath installation and administration guide for steps to install the PowerPath modules.
PowerPath pseudo bus <bus_name> cannot be configured. Uninstall and reinstall PowerPath.	Windows	powermt	An error is preventing creation of the software bus.	Uninstall and then reinstall PowerPath.
PowerPath pseudo(s) still in-use, busy or mounted.	Solaris	powermt	When trying to remove PowerPath from a Solaris 9 host configured to boot from a native device, the <b>pkgrm</b> command fails with this error, though no devices except the boot device are in use. This may occur when you switch from booting from a pseudo-named device to booting from a native-named device. On Solaris 9 hosts, PowerPath cannot be removed from a SAN boot device (when booting off a native device).	Follow the procedure documented in the troubleshooting section of the <i>PowerPath for Solaris Installation and Administration Guide</i> .

**Table 36** PowerPath error messages (page 21 of 28)

Message	Platform	Source	Explanation	Suggested action
PowerPath registration file has been upgraded.	All	emcpreg -install	Your existing PowerPath license was converted to a new PowerPath 5.0 license, and new key(s) were added to the registration file. This appears at the end of the registration process.	No action is required.
PRO service, action 46 failed, 0x5/0x80/0x20	Solaris	powermt	An initiator with no registered Symmetrix path group ID sends a SCSI-3 PERSISTENT RESERVE OUT registration to a LUN. A physical configuration change (such as reconnecting a dead path) causes a bus or device reset, which clears group registrations. This applies only in Symmetrix or VNX OE block and CLARiiON environments.	No action is required. An automatic retry register, including a group register, follows this message. In most cases, this restores all paths without intervention. If cluster problems persists, check the system and array logs for indications of persistent hardware problems.
Pseudo Device Name Major# Minor# <ERROR: unexpected internal error.>	AIX	cfgmgr	The <b>rendev</b> command failed to rename the minor number of the pseudo device.	Use the <b>emcpadm renamepseudo</b> command instead of the <b>rendev</b> command.
Recognize (loop) error number on device <device_id>.	AIX	cfgmgr	PowerPath failed to recognize the specified device.	Verify the system is a supported configuration and the EMC ODM package is loaded.
Registration expired month/year.	All	powermt check_registration	The PowerPath license was an evaluation license that has expired.	To get a permanent license, contact EMC Customer Support.
Expires: month, year=	All	powermt check_registration	The PowerPath license was an evaluation license that will expire on the shown date.	To get a permanent license, contact EMC Customer Support.
Expired: month, year	All	powermt check_registration	The PowerPath license has expired on the shown date.	To get a permanent license, contact EMC Customer Support.
Registration is valid.	All	powermt check_registration	The PowerPath license is valid.	No action is required.
Registration is valid until month/year.	All	powermt check_registration	The PowerPath license is a valid evaluation license that expires after the specified date.	No immediate action is required. When needed, get a permanent license by contacting EMC Customer Support.
Restored volume <clariion volume> to default: SP<clariion storage_processor>.	All	powermt	Informational, VNX and CLARiiON only: The specified volume was trespassed to back to the default storage processor.	No action is required.

**Table 36** PowerPath error messages (page 22 of 28)

Message	Platform	Source	Explanation	Suggested action
saved configuration verified.	All UNIX	emcpupgrade	PowerPath configuration on the system is current. The upgrade is possible.	No action is required.
Sleeping for memory.	AIX	powermt	Error: There is insufficient system memory.	Check your system configuration and memory use.
SPx Not Responding on array : Trespassing volumes.	Solaris	powermt	The current SP is inaccessible. PowerPath issues a trespass to the peer SP and switches to the set of paths attached to the peer SP.	No action is required.
Starting PowerPath daemon kernel process.	AIX	powermt	Informational: PowerPath is being configured.	No action is required.
Storage system <storage_system_name> PowerPath mode disabled.	All	powermt	Informational, VNX and CLARiiON only: PowerPath is not managing the specified VNX and CLARiiON system. PowerPath manages only those VNX OE and CLARiiON systems with PNR (Passive Not Ready) or, with some platforms, ALUA (Asymmetric Logical Unit Access) set.	If you want to manage this storage system with PowerPath, set the PNR or ALUA and LUNZ modes for it, and reboot. Otherwise, no action is required.
<storage_system_class> device(s) not found.	All	All <b>powermt</b> commands with both the <b>class</b> and <b>dev</b> arguments	The <b>powermt</b> command specified a specific <b>class</b> , and no devices were found for the specified type of storage system.	No action is required. If desired, run the command again with a different <b>class</b> value.
Syscall 8133 to device_id failed.	All UNIX and Linux	powermt	Error: An RDF control command to the Symmetrix system failed. This message is followed by a buffer dump of the failed call.	Check your configuration and path states. If the message persists, contact EMC Customer Support.
Temporary upgrade key has been superseded.	All	emcpreg -add, emcpreg -install	The temporary upgrade key was removed, because a site-specific key was added.	No action is required.
There are no license keys now registered.	All UNIX, Linux	powermt check_registration, emcpreg -list	No key is registered; that is, no key was added or upgraded, or the last key was removed.	Add a valid key.
This version of PowerPath is NOT binary compatible with your kernel.	Linux	PowerPath installation	You are installing PowerPath when running a kernel version other than the supported kernel. (Or Linux kernel symbol versioning has yielded false results.)	Make sure the supported kernel is running.

**Table 36** PowerPath error messages (page 23 of 28)

Message	Platform	Source	Explanation	Suggested action
Trespassed volume <clariion_volume> to SP<clariion_storage_processor>	Windows	powermt	Informational, VNX and CLARiiON only: The current SP is inaccessible. PowerPath issues a trespass to the peer SP and switches to the set of paths attached to the peer SP.	Repair physical paths or the SP. Wait for automatic restoration of the volume or, to avoid delay, run <b>powermt restore</b> to restore the volume.
Unable to create kernel process.	AIX	powermt	Panic: The system cannot create the PowerPath daemon.	Check file access settings for the daemon and system resource use.
Unable to follow volume volume to controller.	All	powermt	Informational, active-passive arrays only: The SP or controller that owns the volume has changed. PowerPath was unable to switch to the set of paths attached to the new owning SP or controller.	Repair physical paths to the new owning SP or controller. Then run <b>powermt restore</b> to restore the paths as soon as possible. Until paths are restored, performance may be degraded significantly.
Unable to open /etc/PowerPath Extensions.	AIX	cfgmgr	The system cannot open the configuration file.	Reinstall PowerPath.
Unable to restore volume volume to default controller.	All	powermt	Informational, active-passive arrays only: An attempt to trespass the volume back to the default SP or controller failed. PowerPath attempts this when <b>powermt restore</b> is run or a path to the default SP or controller transitions from dead to alive.	Examine the log file for message 100 or 112. Take the appropriate action for whatever message you find.
Unable to restore volume <clariion_volume> to default SP.	All	powermt	Informational, VNX and CLARiiON only: An attempt to trespass the volume back to the default storage processor (SP) failed. PowerPath attempts this when <b>powermt restore</b> is run or a path to the default SP transitions from dead to alive.	Examine the log file for message 100 or 112. Take the appropriate action for whatever message you find.
Unable to start kernel process.	AIX	powermt	Panic: The system cannot create the PowerPath daemon.	Check file access settings for the daemon and system resource use.
Unknown option -option_name.	All	emcprep	The specified option is invalid.	Re-issue the command, using a valid option.
Unrecoverable error occurred. Aborting.	All	powermt	An unexpected error occurred.	The suggested actions are context-dependent. Contact EMC Customer Support if needed.
User must have administrator privilege to run emcprep command on Windows.	Windows	emcprep	You are trying to run <b>emcprep</b> without administrator privileges.	Log in as administrator, and then run <b>emcprep</b> as desired.

**Table 36** PowerPath error messages (page 24 of 28)

Message	Platform	Source	Explanation	Suggested action
Volume <clariion_volume> followed to SP<clariion_storage_processor>.	All	powermt	Informational, VNX and CLARiiON only: The SP that owns the volume has changed. PowerPath is switching to the set of paths attached to the new owning SP.	No action is required.
Volume <clariion_volume> is unbound.	All	powermt	Error, VNX and CLARiiON only: All paths to the volume are marked dead, because the volume is unbound.	No action is required.
warning - configuration files contains a dead path (device).	All UNIX	emcpupgrade	The specified device is present in the configuration files, but cannot be validated using SCSI inquiry.	Reconfigure any application that was configured using this device with the previous installation of PowerPath.
Warning: cannot unmanage. Verify state of the class class.	All UNIX, Linux	powermt unmanage class	The class status is manage_incomplete.	<ol style="list-style-type: none"> <li>1. Close all applications using the devices of the class.</li> <li>2. Run <b>powermt manage class=class</b>.</li> <li>3. Run <b>powermt unmanage class=class</b>.</li> </ol>
Warning: cannot update persistent store.	All UNIX, Linux	powermt manage class, powermt unmanage class	The <b>powermt manage class</b> or <b>powermt unmanage class</b> command could not update persistent store.	Issue the <b>powermt save</b> command and then re-issue the <b>powermt manage class</b> or <b>powermt unmanage class</b> command.
Warning: class class already managed.	All UNIX, Linux	powermt manage class	You specified a class that is already managed. This command is valid only if the class status is unmanaged or manage_incomplete.	No action is required. The class is already managed.
Warning: class class is already unmanaged.	All UNIX, Linux	powermt manage class	You specified a class that is already unmanaged. This command is valid only if the class status is managed or unmanage_incomplete.	No action is required. The class is already unmanaged.
Warning: complete management of class class failed.	All UNIX, Linux	powermt manage class	The <b>powermt manage class</b> command failed.	<p>Either:</p> <ul style="list-style-type: none"> <li>• Stop all applications using the devices belonging to the class and re-issue the command.</li> <li>• Reboot the host.</li> </ul>
Warning: configuration file created by newer version of powermt.	All	powermt load	There is a mismatch between the version of powermt that is running and the version that created the configuration file. Some information in the file may not be valid or recognized.	No action is required. You can run powermt config (to ensure the file is current with respect to the host system's configuration) and powermt save (to update the file to the new powermt format).



**Table 36** PowerPath error messages (page 25 of 28)

Message	Platform	Source	Explanation	Suggested action
Warning: Device cannot be managed at this time. Please verify the state of the class class.	All UNIX, Linux	powermt manage dev	The class status is either unmanaged or unmanage_incomplete.	Manage the class using <b>powermt manage class</b> and then re-issue the <b>powermt manage dev</b> command.
Warning: device will be configured on closing all the applications to the device or upon reboot.	All UNIX, Linux	powermt manage class	The specified device could not be managed with the <b>powermt manage dev</b> command.	Either: <ul style="list-style-type: none"> <li>Stop all applications using the devices belonging to the class and re-issue the command.</li> <li>Reboot the host.</li> </ul>
ERROR: manage of device <device> failed	All UNIX, Linux	powermt manage class	The specified device could not be managed with the <b>powermt manage dev</b> command.	Either: <ul style="list-style-type: none"> <li>Stop all applications using the devices belonging to the class and re-issue the command.</li> <li>Reboot the host.</li> </ul>
Warning: fail generating emcp registration from pre-3.0.0 file.	Solaris	PowerPath installation	As of version 3.0.0, the format of PowerPath registration keys changed. The installation attempts to convert old keys to the new format. This message indicates that the conversion failed.	Contact EMC Customer Support.
WARNING: forceload of drv/sd failed.	Solaris	PowerPath installation	If the sd driver does not exist on the system, you see this message during boot.	You can safely ignore this message.
WARNING: forceload of drv/ssd failed.	Solaris	PowerPath installation	If the ssd driver does not exist on the system, you see this message during boot.	You can safely ignore this message.
Warning: illegal configuration for basic failover.	All	powermt set policy=bf	The <b>bf</b> (basic failover) policy is supported only in configurations with one HBA and two paths to a volume.	Run <b>powermt set policy</b> again, specifying a legal policy.
ERROR: Unsupported policy:bf.  Usage: powermt set policy={ad co lb li re r so si}  [dev=<device> all   class=<class> all]	All	powermt set policy=bf	The <b>bf</b> (basic failover) policy is supported only in configurations with one HBA and two paths to a volume.	Run <b>powermt set policy</b> again, specifying a legal policy.
Warning: Inconsistent powerpath device state detected. Please reboot first!!!	AIX	cfgmgr	The user did not correctly remove hdiskpower devices, including rebooting after an upgrade installation.	Reboot.

**Table 36** PowerPath error messages (page 26 of 28)

Message	Platform	Source	Explanation	Suggested action
Warning: Kernel symbol versioning is not enabled in your linux kernel.	Linux	PowerPath installation	You are installing on a kernel that does not have <b>CONFIG_MODVERSIONS</b> enabled. PowerPath will not check for binary compatibility with the kernel.	If checking for binary compatibility of kernel modules is necessary, reconfigure the kernel, enabling <b>CONFIG_MODVERSIONS</b> ; rebuild the kernel; and reboot. Otherwise, continue with the install.
WARNING: Key found in old registration file is invalid.	All	emcprep -install	The old key is corrupt or is an evaluation key. Evaluation keys are not accepted for automatic upgrade.	Add a valid registration key before installing PowerPath. If needed, get a new license key from EMC Customer Support.
Warning: key key number in license file is bad. Use -edit option to repair.	All	emcprep -add, emcprep -install	A key in the registration file is invalid.	Run <b>emcprep -edit</b> and remove the bad key. Alternately, remove the key by running <b>emcprep -remove &lt;key_number&gt;</b> .
Warning: license for <storage_system_name> storage system support is missing or expired.	All	powermt	The PowerPath license is invalid or has expired, but the powermt command that was run does not require a valid license.	Run <b>powermt check_registration</b> to determine the problem with the PowerPath license.
WARNING: License not installed.	All	powermt	The PowerPath license is invalid or has expired, but the powermt command that was run does not require a valid license.	Run <b>powermt check_registration</b> to determine the problem with the PowerPath license.
Warning: management of device device will be fully disabled upon reboot or after closing all applications using the device and issuing another powermt unmanage.  Warning: Cannot remove device that is in use: <device>  Warning: Device(s) will be unmanaged on Reboot.  ERROR: unmanage of device <device> failed	All UNIX, Linux	powermt unmanage	The unmanage of a device will not take effect until all applications using the device are closed or the host is rebooted.	Either: <ul style="list-style-type: none"> <li>• Close all applications using the specified device and re-issue the command.</li> <li>• Reboot the host.</li> </ul>

**Table 36** PowerPath error messages (page 27 of 28)

Message	Platform	Source	Explanation	Suggested action
<p>Warning: management of storage array class <code>class</code> will be fully disabled upon reboot or after closing all applications using devices belonging to the storage array class and issuing another <code>powermt unmanage</code>.</p> <p>Warning: Cannot remove device that is in use: <code>&lt;device&gt;</code></p> <p>Warning: management of storage array class <code>vx</code> will be fully disabled upon reboot.</p>	All UNIX, Linux	powermt unmanage class	The unmanage of a class will not take effect until all applications using devices belonging to the storage system class are closed or the host is rebooted.	<p>Either:</p> <ul style="list-style-type: none"> <li>Close all applications using devices belonging to the storage array class and re-issue the command.</li> <li>Reboot the host.</li> </ul>
<p>Warning: management of the path <code>path</code> will be configured upon reboot or after closing all the applications using the path and issuing another <code>powermt manage</code>.</p> <p>Warning: management of the path <code>%s</code> will be configured upon reboot.</p>	All UNIX, Linux	powermt manage dev		<p>Either:</p> <ul style="list-style-type: none"> <li>Close all applications using the path and re-issue the command.</li> <li>Reboot the host.</li> </ul>
<p>Warning: MPxIO is enabled. Verify that no arrays are managed by both MPxIO and PowerPath.</p>	Solaris	PowerPath Installation	MPxIO is enabled on the host where you are installing PowerPath. Co-management of devices by PowerPath and MPxIO is unsupported.	Disable MPxIO or remove from PowerPath management ( <b>powermt unmanage class=<i>class</i></b> ) the array class being managed by MPxIO.
<p>Warning: powermt load failed.</p>	Solaris	PowerPath installation	This can occur after an upgrade. The loading of previously configured PowerPath settings failed.	<p>The <b>powermt-custom</b> file, from which the load is attempted, is saved in the format <b>/etc/powermt.custom.yymmdd:hhmmss.load_failed</b> (for example, <b>/etc/powermt.custom.030408:142021.load_failed</b>). Try to load this file manually by running <b>powermt load</b> from the command line. If this fails, contact EMC Customer Support.</p> <hr/> <p><b>Note:</b> This naming convention differs from the convention used in earlier versions of PowerPath.</p> <hr/>
<p>Warning: powermt save failed.</p>	Solaris	PowerPath installation	This can occur after an upgrade. The attempt to save currently configured PowerPath settings failed.	Run <b>powermt save</b> from the command line. If that fails, ensure that the <code>powermt-custom.xml</code> file exists and check its permissions; if the file is bad, remove it and run <b>powermt save</b> again.

**Table 36** PowerPath error messages (page 28 of 28)

Message	Platform	Source	Explanation	Suggested action
Warning: <storage_system> device path <path_name> is dead.  Warning: Device <device> bus <bus ID> path <path> is currently dead.	All	powermt restore	The path failed the last path test and has been taken out of service.	For more information, refer to “powermt restore” on page 84.
Warning: <storage_system> path <path_name> is dead. Do you want to remove it (y/n/a/q)?	All UNIX and Linux	powermt check	The path failed the last path test.	For more information, refer to “powermt check” on page 26.
Warning: Unable to determine the prerequisite patch(es) - Please manually check and install the following patch(es) after installing PowerPath. Required Patch(es): patch_number Continue installation? [y,n,q,?] (default: y):	Solaris	PowerPath installation	The installation cannot determine whether the specified required patch is installed (for example, some patches cannot be detected with the standard Solaris <b>patchadd</b> utility).	Verify that the specified patch is installed; if the patch is not installed, install it.  Note: If you cannot check for (and, if necessary, install) the patch immediately after this installation completes, abort this installation, verify that the patch is installed, then install PowerPath 5.0.
Warning: unable to open the configuration file, /etc/powermt.custom.  Agile device for Mpx volume does not match previous Agile. Indicates possible misconfigured storage device not responding to SCSI Inquiry page 83.  Agile devices for possible misconfigured storage	HP-UX		HP-UX 11.31 has created multiple Agile devices per Symmetrix volume causing path management issues.	When connecting an HP-UX 11.31 host to a Symmetrix volume, use the following settings: a. Common_Serial_Number(C) : Enabled b. SPC2_Protocol_Version(SPC2) : Enabled c. Volume_Set_Addresssing(V) : Enabled d. Init_Point_to_Point(PP) : Enabled e. Unique_WWN(UWN) : Enabled
write failed for lack of space, check for filesystem full.	AIX	PowerPath installation	The file system containing /etc does not contain enough free space.	Increase the amount of free space in the file system containing /etc.
You must be root to run emcpregr.	All	emcpregr	You are trying to run <b>emcpregr</b> without being logged in as root.	Log in as root, and then run <b>emcpregr</b> as desired.

## Windows system event IDs

[Table 37 on page 226](#) lists all messages that appear on Windows platforms, along with their IDs. For an explanation and suggested actions for each message, refer to the indicated page numbers to locate the message in [Table 36 on page 198](#) which starts on [page 198](#).

**Note:** If a PowerPath message exceeds the maximum message size allowed by Windows, the Event Viewer truncates the message and places the data fields in the dump portion of the event (the hex information at the bottom of the viewer). For example, the message “Path c3t1d10 to 60001440B06028A300A4A48002000010 is alive.” exceeds the allowed maximum size. The message would be displayed as “Path @1 to @2 is alive.” with the data for @1 and @2 included at the bottom.

**Table 37** PowerPath messages: Windows platforms

Msg ID	Message	Refer to page
2	PowerPath driver loaded.	<a href="#">page 217</a>
5	PowerPath pseudo bus <bus_name> cannot be configured. Uninstall and reinstall PowerPath.	<a href="#">page 217</a>
59	Disk <X> will not be used because it is a redundant path for disk <Y>.	<a href="#">page 216</a>
100	Path <path_name> to <device_id> is dead.	<a href="#">page 216</a>
101	Path <path_name> to <device_id > is alive.	<a href="#">page 216</a>
102	<device_id> is dead.	<a href="#">page 204</a>
103	<device_id> is alive.	<a href="#">page 204</a>
104	All paths to <device_id> are dead.	<a href="#">page 198</a>
105	Not all paths to <device_id> are dead.	<a href="#">page 215</a>
106	Bus <bus_id> to <storage_system_name> port <port_name> is dead.	<a href="#">page 200</a>
108	Volume <clariion_volume> is unbound.	<a href="#">page 221</a>
109	Restored volume <clariion_volume> to default: SP<clariion_storage_processor>.	<a href="#">page 218</a>
110	Unable to restore volume <clariion_volume> to default SP.	<a href="#">page 220</a>
111	Volume <clariion_volume> followed to SP<clariion_storage_processor>.	<a href="#">page 221</a>
112	Unable to follow volume <clariion_volume> to SP<clariion_storage_processor>.	<a href="#">page 220</a>
113	Trespassed volume <clariion_volume> to SP<clariion_storage_processor>.	<a href="#">page 220</a>
116	Storage system <storage_system_name> PowerPath mode disabled.	<a href="#">page 219</a>
120	Bus <bus_id> to <storage_system_name> port <port_name> is alive.	<a href="#">page 200</a>
303	Encountered unexpected error 0x<code> on Dek <operation> command for device <id> on handle <association_handle>.	<a href="#">page 205</a>

## Windows application event IDs

[Table 38 on page 227](#) lists the Windows application event messages that appear on Windows platform, along with their IDs.

**Table 38** Event codes

Event ID	Description	Corrective measure
1,2,5	EmcAdminSvr could not initialize PowerPath driver	Verify if PowerPath is installed correctly, or reinstall PowerPath.
6	The <b>EmcPowSrv</b> is not running.	Run <b>EmcPowSrv</b> at system startup to ensure PowerPath is initialized properly.
7	<b>EmcAdminSvr</b> processing thread has thrown an exception.	Close the GUI (MMC snap-in) and try opening again.
8, 9	<b>EMCPowerPathAdmin</b> could not remove the <b>device_id</b> from the current configuration.	Remove the device again and try again.
10, 11	<b>EMCPowerPathAdmin</b> could not reconfigure the <b>device_id</b> paths.	Reconfigure the device and try again.
12, 13	<b>EMCPowerPathAdmin</b> could not remove the <b>path_id</b> from the current configuration.	Remove the path and try again.
14, 15	<b>EMCPowerPathAdmin</b> could not reconfigure the <b>path_id</b> .	Reconfigure the path again.
16,17	<b>EMCPowerPathAdmin</b> could not remove the adapter from the current configuration.	Remove the adapter and try again.
18, 19	<b>EMCPowerPathAdmin</b> could not reconfigure the adapter paths.	Reconfigure the adapter paths and try again.
20, 21	<b>EMCPowerPathAdmin</b> could not reconfigure paths.	Reconfigure the paths and try again.
22, 23, 24	<b>EMCPowerPathAdmin</b> could not perform a Restore operation.	Perform the Restore operation again.
25	<b>EmcAdminSvr</b> could not set the volume properties.	Set the volume properties and try again.
26	<b>EmcAdminSvr</b> could not set the path properties.	Reset the path properties and try again.
27,28	<b>EmcAdminSvr</b> could not set adapter properties.	Set the adapter properties and try again.
29,30,31,32	<b>EmcAdminSvr</b> could not load the driver configuration.	Load the driver configuration and try again.
33,34,35,36	<b>EmcAdminSvr</b> could not save the driver configuration file.	Save the driver configuration and try again.
37,38,39	<b>EmcAdminSvr</b> could not get AutoRestore settings.	Retrieve or modify the Auto Restore settings and try again.
40,41,42,43,44	<b>EmcAdminSvr</b> could not set the AutoRestore settings.	Reset the Auto Restore and try again.
47	Could not remove Service. Make sure the service has been stopped.	Stop the service before it can be unregistered or removed.

# CHAPTER 5

## PowerPath Migration Enabler Messages

This chapter contains the following topics:

- ◆ [Audit messages](#) ..... 229
- ◆ [Responding to error messages](#) ..... 233

## Audit messages

This section describes audit and error log messages logged to a common file, and error messages that appear when **powermig** commands fail. Audit messages show migration state transitions, logical-unit fault conditions, and generally record the migration workflow. Error and warning messages record unexpected events.

The following sections cover these topics in more detail:

- ◆ [“Logging messages to a common file” on page 229](#)
- ◆ [“Audit and error log messages” on page 229](#)
- ◆ [“Responding to error messages” on page 233](#)

## Logging messages to a common file

This section describes how to configure the host operating system to capture Migration Enabler error and log messages.

### AIX hosts

To log messages in `/tmp/emcpsyslog.log`:

1. Add the following entry to the `/etc/syslog.conf` file:

```
*.info /tmp/emcpsyslog.log
```

Optionally, configure `/etc/syslog.conf` to rotate the log file.

For example, to rotate the file daily, keep one week’s worth of files, and compress files to save space, add the following entry:

```
*.info /tmp/emcpsyslog.log rotate files 7 time 1d compress
```

2. Create the `/tmp/emcpsyslog.log` file. For example, touch `/tmp/emcpsyslog.log`
3. Type **refresh -s syslogd** to enable logging.

## Audit and error log messages

Two types of log messages are described in this section: audit log messages and error log messages.

### Audit log messages

PowerPath Migration Enabler records all events that change the migration state, including:

- ◆ command execution
- ◆ transition to the **SourceSelected** state when a synchronization completes
- ◆ detection of a logical-unit fault

On UNIX hosts, Migration Enabler uses **syslog(2)** to provide audit logging functionality. On Windows hosts, audit messages appear in the **emcpaudit.log** file and in the system event log.

### Message format and description

On UNIX hosts, audit log messages have the following format:



```
<timestamp> <hostname> emcpAudit: [ID <pid> local0.info] PPME: <message>
```

### Sample message

The following sample message from a UNIX host shows a migration in the **setup** state.

```
Feb 27 13:03:42 Host1a emcpAudit: [ID 702911 local0.info] PPME: Info: handle=1,
event=stateChanged, state=setup, cmd=setup, tech=OR, src=emcpower0a (000187910018:0022),
tgt=emcpower1a (000187910018:0014)
```

[Table 39 on page 230](#) describes the audit log message components.

**Table 39** Audit message description (page 1 of 2)

Message	Description
Feb 27 13:03:42	Date and time the event occurred.
Host1a	Name of the host where PowerPath Migration Enabler is running. <b>Note:</b> The host name does not appear in audit messages on Windows hosts.
emcpAudit: EMC PowerPath Audit (on Windows hosts)	The program performing the auditing function. This is the tag used for audit messages for PowerPath.
ID 702911 local0.info	The operating system's process ID (ID 702911) and the message priority. The message priority includes the facility (local0) and the level (info). <b>Note:</b> The process ID does not appear in audit messages on Windows hosts.
PPME: Info Information: (on Windows hosts)	Product and message severity.
handle=1	Handle for the migration.
event=stateChanged	Event that caused the log to be written. There are two possible values: <b>stateChanged</b> and <b>faultDetected</b> <ul style="list-style-type: none"> <li><b>stateChanged</b> — indicates the migration session transitioned to another state.</li> <li><b>faultDetected</b> — indicates a logical-unit fault was detected.</li> </ul> The <b>targetLUfault</b> or <b>sourceLUfault</b> values also appear when a logical-unit fault occurs.
state=setup	State of the migration. The migration state appears only when the <b>event=stateChanged</b> .
cmd=setup	Command that caused the state transition. This information appears only when the <b>event=stateChanged</b> and the change was the result of a command execution (as opposed to an automatic change, as in the transition from <b>Syncing</b> to <b>SourceSelected</b> ).

**Table 39** Audit message description (page 2 of 2)

Message	Description
tech=OR	The underlying technology used in conjunction with PowerPath Migration Enabler. Possible values are <b>OR</b> (Open Replicator) and <b>INVE</b> (Virtual encapsulation), <b>SYMCLONE</b> (migration within a Symmetrix) and <b>HOSTCOPY</b> (host-based copy). This information appears only when the <b>event=stateChanged</b> and <b>cmd=setup</b> .
src=emcpower4a (000187910018:0022)	Source-device name involved in the migration. This information appears only when the <b>event=stateChanged</b> and <b>cmd=setup</b> . If the source logical unit is a Symmetrix device, then the Symmetrix serial number appears in parenthesis along with the device name. If the source logical unit is any other type of device, the storage array WWN appears.
tgt=emcpower9a (000187910018:0014)	Target-device name involved in the migration. This information appears only when the <b>event=stateChanged</b> and <b>cmd=setup</b> . If the target logical unit is a Symmetrix device, then the Symmetrix serial number appears in parenthesis along with the device name. If the target is any other type of device, the storage array WWN appears in parenthesis.

### Sample messages: Migration process and logical-unit fault

The following sample audit messages from a UNIX host show a logical-unit fault and a successful migration. The fault occurs on the target logical unit.

#### Migration process

These messages show **powermig** commands completing successfully:

```
Feb 27 13:03:42 Host1a emcpAudit: [ID 702911 local0.info] PPME: Info: handle=1,
event=stateChanged, state=setup, cmd=setup, tech=OR, src=emcpower0a (000187910018:0022),
tgt=emcpower1a (000187910018:0014)
Feb 27 13:03:43 Host1a emcpAudit: [ID 702911 local0.info] PPME: Info: handle=1,
Event=stateChanged, state=syncing, cmd=sync
Feb 27 13:03:45 Host1a emcpAudit: [ID 702911 local0.info] PPME: Info: handle=1,
event=stateChanged, state=sourceSelected
Feb 27 13:03:47 Host1a emcpAudit: [ID 702911 local0.info] PPME: Info: handle=1,
event=stateChanged, state=targetSelected, cmd=selectTarget
```

#### Target logical-unit fault

This message indicates a target fault occurred:

```
Feb 27 13:03:48 Host1a emcpAudit: [ID 702911 local0.info] PPME: Info: handle=1,
event=faultDetected, targetLUfault
```

#### Abort the migration

When a target fault occurs, the correct course of action is to abort the migration. When you abort the migration, a message similar to the following is logged:

```
Feb 27 13:03:48 Host1a emcpAudit: [ID 702911 local0.info] PPME: Info: handle=1,
event=stateChanged, state=setup, cmd=abort, targetLUfault
```

#### Logical-unit fault correction and successful migration

After aborting the session, the migration returns to the **Setup** state. You can then restart the synchronization after correcting the fault. The following messages (from a UNIX host) are logged during a successfully completed migration:

```
Feb 27 13:03:49 Host1a emcpAudit: [ID 702911 local0.info] PPME: Info: handle=1,
event=stateChanged, state=syncing, cmd=sync
```

```
Feb 27 13:03:52 Host1a emcpAudit: [ID 702911 local0.info] PPME: Info: handle=1,
event=stateChanged, state=sourceSelected
Feb 27 13:04:02 Host1a emcpAudit: [ID 702911 local0.info] PPME: Info: handle=1,
event=stateChanged, state=targetSelected, cmd=selectTarget
Feb 27 13:04:50 Host1a emcpAudit: [ID 702911 local0.info] PPME: Info: handle=1,
event=stateChanged, state=committed, cmd=commit
Feb 27 13:05:28 Host1a emcpAudit: [ID 702911 local0.info] PPME: Info: handle=1,
event=stateChanged, state=cleanup, cmd=cleanup
```

## Error log messages

Error log messages appear along with audit log messages in the log file when it is configured for common logging. The error log messages capture unexpected events that occur during the course of a migration. Some error log messages convey information that also appears on-screen when a **powermig** command fails. This section describes the basic error-log message format.

### Error log message format

On UNIX hosts, error log messages have the following format:

```
<timestamp> <hostname> PPME: [ID <pid> local0.<module>]: <message>
```

The log components are described in [Table 40 on page 232](#).

### Sample message

The following sample error log message from a UNIX host shows a failed attempt to set up a migration.

```
Aug 1 16:51:13 lcl1a075 PPME: [ID 702911 local0.error] API: Error: Setup
failed. rc=100
```

[Table 40 on page 232](#) describes the components of the sample message.

**Table 40** Error log message description

Message	Description
Aug 1 01:19:49	Date and time the event occurred.
Host1a	Name of the host where PowerPath Migration Enabler is running. <b>Note:</b> The host name does not appear in log messages on Windows hosts.
PPME:	The program performing the auditing function. This is the tag used for PowerPath Migration Enabler error messages.
ID 702911 local0.error	The operating system's process ID (ID 702911) and the message priority. The message priority includes the facility (local0) and the level (info). <b>Note:</b> The process ID does not appear in log messages on Windows hosts.
API:	The module where the error occurred. This can be API, OR (Open Replicator), or SYS (system).
Error	Message type. This can be <b>Info</b> , <b>Error</b> , or <b>Warning</b> .
Successfully recovered handle=1	The error message. For an explanation of error messages you may see, refer to <a href="#">Table 41 on page 234</a> .

## Host Copy log file

Hostcopy migrations have a separate log file called that shows the progress of a migration. On Solaris hosts, the **emcpmigd** daemon logs messages to **/etc/emc/ppme/emcpmigd.log**. On Windows hosts, PowerPath Migration Service (PowerMigSvc) orchestrates the copying of data and logs messages to the **hostcopy.log** in **%InstallDir%\EMC\PowerPath\Logs**.

The log file includes messages about the amount of data copied during each copy cycle, the throttle values, and **powermig** commands received and processed by the daemon.

## Responding to error messages

An on-screen error message indicates a condition that prevents a command from completing successfully. Messages generated by PowerPath Migration Enabler are documented in this section. Migration Enabler messages are identified by one of the following prefixes:

- ◆ PPME error — This designates a general error.
- ◆ PPME\_PP error — This designates an error related to PowerPath functionality.
- ◆ PPME\_OR error — Designates an error related to Open Replicator. For more information, refer to the Solutions Enabler documentation, available on the EMC Online Support site.
- ◆ PPME\_SYMCLONE error — Designates an error related to TimeFinder/Clone. For more information, refer to the Solutions Enabler documentation, available on the EMC Online Support site.
- ◆ PPME\_INVE error — Designates an error related to Virtual Encapsulation.
- ◆ PPME\_HOSTCOPY error — Designates an error related to a Host Copy migration.
- ◆ PPME\_DB error — Designates an error related to userspace metadata (UMD).

PowerPath error messages have the prefixes **PowerPath:Error**, **emcp: mod=Mpx**, and **emcp**.

## Error message descriptions

[Table 41 on page 234](#) lists the error messages you may see when a **powermig** command fails to complete. The table includes probable causes and solutions to correct the error condition.

**Table 41** PowerPath Migration Enabler error messages (page 1 of 6)

Message	Description	Suggested action
A function call exceeded allocated time	This applies only to Open Replicator.	The Solutions Enabler documentation provides additional information.
Already in the state requested	The migration is in the state to which you tried to transition.	Proceed to the next step in the migration process.
Already recovered	The <b>powermig recover</b> command was run and the migration is not in <b>needsRecovery</b> state.	<ol style="list-style-type: none"> <li>1. Run the <b>powermig info</b> command to identify the state of the migration.</li> <li>2. Resume the migration at the appropriate point in the migration process.</li> </ol>
An error occurred while using the base daemon, please see the log file.	This applies to Open Replicator and TimeFinder/Clone.	The Solutions Enabler documentation provides additional information.
An internal error occurred during a SYMAPI operation. Please report to EMC.	During a SYMAPI operation with Solution Enabler, the host has crashed causing the symapi database corruption.	<p>Add the following line in the symapi options file (/var/symapi/config/options) to flush all the symapi database writes to disk:</p> <p>SYMAPI_DB_FSYNC_MODE = ENABLE</p>
Command usage is displayed	Either no command or an invalid command was entered.	<ol style="list-style-type: none"> <li>1. Check you have the correct command syntax.</li> <li>2. Re-run the command.</li> </ol>
Command/option not supported for this technology	The technology type you are using in conjunction with Migration Enabler does not support the command or option that you entered.	Enter commands and options that are supported with the technology type.
Couldn't read disk label	This applies only to Solaris and occurs when the source or target EFI or VTOC label cannot be obtained. The system error log displays whether it was a source or target and EFI or VTOC problem.	Use the <b>format</b> or <b>powerformat</b> utility to verify that the source or target logical unit involved in the migration has a valid disk label and are of the same label type (EFI or VTOC).
Couldn't write disk label, be sure target is at least as large as source	This applies only to Solaris and occurs when a copy of the label fails from source to target during <b>powermig setup</b> . The copy could fail because the target is not large enough to accept the label.	Use the <b>format</b> or <b>powerformat</b> utility to verify the target logical unit is accessible, is as large as the source, and has a valid disk label of the same type as the source.
Daemon communication error	This applies only to Host Copy and occurs when read or write fails to migration daemon socket (UNIX) or migration Service named pipe (Windows).	<ul style="list-style-type: none"> <li>• Run <b>powermig recover</b> to retry the failed command.</li> <li>• Check that the emcpmigd daemon or the Service is running.</li> <li>• If the problem persists, contact EMC Customer Support.</li> </ul>

**Table 41** PowerPath Migration Enabler error messages (page 2 of 6)

Message	Description	Suggested action
Daemon failed to create thread	This applies only to Host Copy and occurs during the <b>syncing</b> state when thread creation fails in the migration daemon.	<ul style="list-style-type: none"> <li>• Run <b>powermig recover</b> to retry the failed command.</li> <li>• If the problem persists, contact EMC Customer Support.</li> </ul>
Daemon semaphore error	This applies only to Host Copy and occurs during the <b>syncing</b> state when Migration Enabler is unable to obtain semaphore for obtaining Host Copy copy slot semaphore.	<ul style="list-style-type: none"> <li>• Run <b>powermig recover</b> to retry the failed command.</li> <li>• If the problem persists, contact EMC Customer Support.</li> </ul>
Daemon socket connection failed	This occurs during the <b>syncing</b> state when failing to set up connection to the migration daemon socket (UNIX) or migration Service named pipe (Windows).	If the problem persists, contact EMC Customer Support.
Device is inaccessible	Migration Enabler is not able to access the source or target device.	Review the system error log to determine which device is inaccessible and to verify that the logical unit is not opened exclusively by another application.
Devices already involved in a migration	The source or target (pseudo or native) is being used in another migration. A logical unit can only be involved in one migration at a time.	<ul style="list-style-type: none"> <li>• Review the system error log to determine which device is part of another migration.</li> <li>• Ensure you entered the correct device name.</li> </ul>
Encrypted native device not supported	This occurs during <b>setup</b> when a pseudo device is encrypted and one of its native paths is specified as source or target.	Specify a pseudo device during <b>setup</b> .
Encryption device not allowed	This occurs during the <b>setup</b> state when either the source or target is encrypted and when using a technology type other than Host Copy.	Use Host Copy to migrate encrypted devices.
Failed to fix disk label of the target device	PowerPath Migration Enabler did not update the ASCII text label or disk signature, which includes the disk vendor, product, and revision. This update is usually done as part of a migration when the migration is a heterogeneous migration in a Solaris host environment.	<ul style="list-style-type: none"> <li>• Use the <b>powermig cleanup -force</b> command to ignore the failure.</li> <li>• For Solaris, use the <b>powerformat</b> utility to update the ASCII disk label information.</li> <li>• Review the system error log to determine which error is the cause.</li> </ul>
Failed to get disk layout of device	For Solaris hosts, after reading the label during the <b>setup</b> state, the disk geometry cannot be obtained. For Linux and Windows hosts, this may occur during <b>setup</b> if the disk geometry or disk size cannot be retrieved, or during <b>cleanup</b> (without the <b>-format</b> option), if the disk size cannot be retrieved as part of erasing the disk layout.	<p>For Solaris hosts, use the <b>format</b> or <b>powerformat</b> utility to verify the disk label.</p> <p>For Linux and Windows hosts, verify the disk is accessible.</p>

**Table 41** PowerPath Migration Enabler error messages (page 3 of 6)

Message	Description	Suggested action
Failed to get target size	Unable to obtain the size of the target. For Solaris hosts, there was a failure to obtain the size of the target from SCSI read_capacity IOCTL.	Check the accessibility of the target.
Failed to load EFI support library	For Solaris hosts, the EFI shared library failed to open. For AIX hosts, this may occur after running <b>powermig cleanup</b> when opening attributes class in ODM.	Contact EMC Customer Support.
Failed to set disk layout of device	For Solaris hosts, unable to set the target disk geometry during <b>setup</b> . For Windows hosts, this may occur during <b>setup</b> if the disk geometry or disk size cannot be retrieved, or during <b>cleanup</b> (without the <b>-format</b> option), if the disk size cannot be retrieved as part of erasing the disk layout.	For Solaris hosts, use the <b>format</b> or <b>powerformat</b> utility to verify the disk label. For Windows hosts, verify the disk is accessible.
Failed to validate geometry of target device	This applies only to Solaris. During <b>setup</b> , after the source VTOC label is copied to the target, the target geometry validation fails. Such as when Migration Enabler is unable to determine that the source label matches the target.	<ul style="list-style-type: none"> <li>• If you see this message when setting up a migration, re-run the <b>powermig setup</b> command (the problem may be intermittent).</li> <li>• If the error still occurs, provision a larger target device.</li> </ul>
Handle not found	The specified handle could not be found in any migration in the UMD.	<ul style="list-style-type: none"> <li>• Verify you are entering the correct handle for the source and target device-name pair.</li> <li>• Use the <b>powermig getHandle</b> command to determine the correct handle. However, once a migration has been cleaned up, the handle is no longer valid.</li> </ul>
Host can still access source device	This applies only to Invista encapsulation and occurs when the source is still available to the host after successful encapsulation.	Remove host access to the source logical unit before running the <b>powermig cleanup</b> command.
Invista device not in proper state for this operation	This applies only to Invista encapsulation and indicates that the target is in an unexpected state.	Review the system log and then contact EMC Customer Support.
Migration stopped due to IO error - will attempt to restart	This applies only to Host Copy and will occur when there is no application I/O and the source or target is unavailable.	Resolve the I/O problem or abort the migration.
Needs recovery	A process crash or error occurred during the execution of a <b>powermig</b> command, and the command needs to be recovered.	Resolve the problem and then run the <b>powermig recover</b> command.
no migration is in the proper state to perform the requested operation	This may occur when using the <b>-all</b> option if there is no migration in a valid state for the command.	Choose a different command that is valid for the state of at least one or more of the migrations.

**Table 41** PowerPath Migration Enabler error messages (page 4 of 6)

Message	Description	Suggested action
Not in proper state to perform this operation	The migration is not in the required state for the command being executed.	<ul style="list-style-type: none"> <li>To determine the state of the migration, run the <b>powermig info</b> or <b>powermig query</b> command.</li> <li>Run the proper command that is valid for the current state of the migration.</li> </ul>
Open failed for unexpected reason when trying to verify source removal	This applies only to Invista encapsulation. The underlying migration technology requires that host access to the source logical unit be removed before you run <b>powermig cleanup</b> . However, an operating system error prevented verification of this requirement.	Refer to the system log for more information to help you resolve this problem, and then retry the <b>powermig cleanup</b> command.
Out of memory	Memory resources requested by the API are not available.	Resolve the problem and then re-run the command.
Remote SE environment variable(s) not allowed in setup	This may occur during the <b>setup</b> state if either environment variable <b>PPME_REMOTE_SE_SERVER</b> or <b>PPME_REMOTE_SE_PORT</b> is set.	<ul style="list-style-type: none"> <li>Use the <b>-techHost</b> or <b>-techPort</b> command line option during <b>setup</b>.</li> <li>Delete the environment variables.</li> </ul>
Source and target are the same device	The same device name was specified as the source and target when you set up the migration ( <b>powermig setup</b> ).	Ensure that the source and target names point to different devices, and then retry the migration setup.
Source and target must be from same array	This applies only to TimeFinder/Clone and occurs when the devices specified are from different arrays. TimeFinder/Clone requires that devices are from the same Symmetrix.	Ensure that the source and target names point to devices on the same array.
Source device is in use (-force will override)	This may occur after running the <b>powermig undoRedirect</b> command if the source device is still in use.	<ol style="list-style-type: none"> <li>Bring down any application sending I/O to the target.</li> <li>Run the <b>powermig undoRedirect</b> command.</li> <li>Reconfigure the application to use the target device name.</li> </ol> <p>Otherwise, use the <b>-force</b> option to stop the redirection of I/O from the source to the target even if the source device is in use.</p>
Source faulted, must commit migration	This occurs during the <b>sourceSelected</b> or <b>targetSelected</b> state when a source logical-unit fault occurred after the synchronization completes.	<ul style="list-style-type: none"> <li>If you are in the <b>sourceSelected</b> state, run the <b>powermig selectTarget</b> command and then the <b>powermig commit</b> command.</li> <li>If you are in the <b>targetSelected</b> state, run the <b>powermig commit</b> command.</li> </ul>
Source faulted, must abort migration	This occurs during the <b>syncing</b> state when a source logical-unit fault occurred before the synchronization completes.	Run the <b>powermig abort</b> command to abort the migration. After you resolve the fault situation, restart the synchronization.



**Table 41** PowerPath Migration Enabler error messages (page 5 of 6)

Message	Description	Suggested action
Target device is in use	This occurs during <b>setup</b> when the target is accessed by an application.	<ul style="list-style-type: none"> <li>• Ensure that no application is accessing the target. Such as in a VxVM environment.</li> <li>• On UNIX hosts, run <b>fuser</b> to discover which process is holding open the target.</li> <li>• Use <b>-force</b> if you are sure that the device can be used for a migration.</li> </ul>
Target does not encapsulate source	This applies only to Invista encapsulation.	Follow the instructions on configuring the underlying technology in the <i>EMC PowerPath Migration Enabler User Guide</i> , and then retry the migration.
Target expected to be Symmetrix device	This applies only to Open Replicator and occurs during <b>setup</b> when the target is not a Symmetrix device.	Ensure that the target is a Symmetrix device.
Target faulted, must abort migration	This occurs during the <b>syncing, sourceSelected, or targetSelected</b> state when a target logical-unit fault occurred.	Run the <b>powermig abort</b> command to abort the migration. After you resolve the fault situation, restart the migration session.
Target is not accessible	This occurs before transitioning to the <b>commit</b> state if Migration Enabler cannot read one block from the target. The migration remains in the <b>targetSelected</b> state.	<ul style="list-style-type: none"> <li>• Check the availability of the target.</li> <li>• If the problem cannot be resolved, run the <b>powermig abort</b> command to abort the migration.</li> </ul>
Target must be a pseudo	A native device name was entered as the target, and a pseudo device name was entered as the source in a <b>powermig setup</b> command. Pseudo-named to native-named device migrations are not supported.	Specify a pseudo target device.
Target must be at least as large as the source	During <b>setup</b> , the target is smaller than the source.	<ul style="list-style-type: none"> <li>• Specify a target that is the same size or larger than the source.</li> <li>• For encrypted targets, ensure that the size accounts for metadata requirements (65 KB).</li> </ul>
Target should be in TEST state.	This applies only to Invista encapsulation.	Collect an <b>emcgrab</b> file and then contact EMC Customer Support.
Target max_transfer attribute is too small	This applies only to AIX and Windows and occurs during <b>setup</b> when the max_transfer attribute of the target is smaller than that of the source.	Modify the target's max_transfer size to be at least equal to that of the source.
The feature being requested is not currently licensed	This is not a PowerPath Migration Enabler licensing error. The technology that you are using with Migration Enabler is not licensed.	Ensure that all technologies that you use with Migration Enabler are licensed. The documentation for the technology provides more information.

**Table 41** PowerPath Migration Enabler error messages (page 6 of 6)

Message	Description	Suggested action
The gatekeeper device (while using the Base Daemon) has an error (Please see the log file)	This is a Solutions Enabler error that applies to Open Replicator and TimeFinder/Clone. You may need to modify the semaphore parameters for SYMCLI. The semaphore parameters control access to the database file and gatekeepers.	The Solutions Enabler documentation provides additional information.
UMD database not initialized	This may occur on any <b>powermig</b> command if the userspace metadata (UMD) is not opened properly.	Contact EMC Customer Support.
Unable to determine encryption state of one of the devices	This occurs during <b>setup</b> if Migration Enabler is unable to determine the encryption state of the source or target.	<ul style="list-style-type: none"> <li>Check for errors in the system log to determine which device is causing the issue. A PowerPath Encryption error will also be included in the log.</li> <li>Run the <b>powervt xcrypt</b> command for additional information.</li> </ul>
Unable to open target	This applies only to Windows and InVista encapsulation and may occur during a graceful shutdown. Migration Enabler was unable to open the target logical unit.	<ol style="list-style-type: none"> <li>Check the system log.</li> <li>Take any necessary corrective action.</li> <li>Retry the command.</li> </ol>
Unable to query InVista state	This applies only to InVista encapsulation.	Collect an <b>emcgrab</b> file and then contact EMC Customer Support.
Unable to scrub the abandoned device	This occurs during <b>cleanup</b> if there is a problem scrubbing the abandoned device.	Review the system log for more detail.
Unable to scrub the abandoned device. Device is offline	This applies only to AIX and Solaris and may occur during <b>cleanup</b> if there is a problem scrubbing the abandoned device. Device is offline or cannot be opened.	Check the state of the abandoned device.
Unable to set InVista state to COMMIT	This applies only to InVista encapsulation.	Collect an <b>emcgrab</b> file and then contact EMC Customer Support.
Unable to set InVista state to NOT_READY		
Unable to set InVista state to TEST		
Unknown device	This may occur during <b>setup</b> when the source or target device name entered cannot be found.	Enter a valid device name, and then retry the command.

# CHAPTER 6

## PowerPath Encryption Messages

This chapter contains the following topics:

- ◆ [Audit messages](#) ..... 241
- ◆ [Error messages](#) ..... 242

## Audit messages

Audit messages record administrative actions. On Linux and Solaris hosts, PowerPath Encryption uses the syslog facility to provide audit logging functionality

### Audit message format

[Table 42 on page 241](#) describes the audit message format:

**Table 42** Audit message component

Component	Description
Timestamp	Date and time the event occurred. For example, Mar 03 13:03:42
Hostname	Name of the host where PowerPath Encryption is running. For example, host1.lss.abc.com
Message source	Name of the module that generated the audit message: <ul style="list-style-type: none"> <li>• <b>emcp_xcryptd</b> daemon</li> <li>• PPVT (<b>powervt</b> command)</li> <li>• <b>xcrypt_config</b> script or <b>RKM_Config.bat</b> batch file</li> </ul>
Process ID and priority	The operating system process ID (for example, <b>ID 815963</b> ) and message priority. The message priority includes the facility (for example, <b>local0</b> ) and the log level (for example, <b>info</b> ).
Audit program name	The program performing the auditing function. The <b>emcpAudit</b> tag is used for audit messages generated by PowerPath Encryption.
Log level	Type of message: <b>Info</b> , <b>Warning</b> , or <b>Error</b> .
Command	The command being logged. For example, <b>powervt xcrypt -on -dev emcpower10c</b> .
Real user ID	Real user ID. For example, <b>root</b> .
Effective user ID	Effective user ID. For example, <b>root</b> .

A sample audit message on a Linux host is shown below:

```
Mar 18 03:31:53 host1.lss.abc.com emcp_xcryptd: [ID 678949 local0.info]
emcpAudit: Info: cmd=KeyLookUpSuccess: Lookup Key success for device emcpower10c
```

[Table 43 on page 241](#) lists audit messages logged by PowerPath Encryption.

**Table 43** Audit messages (page 1 of 2)

Source	Message	Description
emcp_xcryptd daemon	KeyLookUpSuccess: Lookup Key success for device <device>.	The <b>emcp_xcryptd</b> daemon was able to retrieve the encryption key for the specified device.
	KeyLookUpFailure: Lookup Key failed for device <device>.	The <b>emcp_xcryptd</b> daemon was not able to retrieve the encryption key for the specified device. Check DPM appliance connectivity and availability.

**Table 43** Audit messages (page 2 of 2)

Source	Message	Description
powervt command	Turn on encryption on <i>&lt;logical_unit&gt;</i>	Encryption was enabled on the specified device.
	Turn off encryption on <i>&lt;logical_unit&gt;</i>	Encryption was disabled on the specified device.
	Resize <i>&lt;logical_unit&gt;</i>	The specified device was resized.
xcrypt_config command	Failed installing files for powerxcrypt service.	The <b>xcrypt_config</b> command was not able to copy files from the EMCpower package directories to AIX system directories. Verify that the EMCpower package is installed correctly.
	Configuring cst.	The lockbox is being configured on the host.
	Configuring ckm.	The key manager client is being configured on the host.
	A valid license key for PowerPath Encryption with RSA was not found. Please register a valid license and re-run xcrypt_config command.	There is no valid PowerPath Encryption license on the host. Use the emcpreg command to register a valid PowerPath Encryption license and re-run the <b>xcrypt_config</b> command.

## Error messages

Error messages capture unexpected events:

- ◆ On Linux and Solaris hosts, PowerPath Encryption reports errors to the system console and (in some cases) to the **syslog** file defined in `/etc/syslog.conf`.
- ◆ On Windows host, error messages are written to **Event Viewer - Application** with a source of **EmcPowerPathEncryption**.

## Error message sources

The following components of a PowerPath Encryption system generate error messages:

- ◆ RSA Data Protection Manager (DPM) server
- ◆ RSA Key Manager client on the PowerPath host
- ◆ Lockbox service on the PowerPath host
- ◆ PowerPath kernel
- ◆ ckmadm command
- ◆ cstadmin command
- ◆ emcp\_xcryptd daemon
- ◆ powervt command
- ◆ xcrypt\_config configuration script

## RSA Data Protection Manager (DPM) server messages

Table 44 on page 243 lists DPM server error messages, descriptions, and possible remedies.

**Table 44** DPM server error messages (page 1 of 2)

Message	Description	Suggested action
Access violation error occurred where the user application had no permission to access the requested key.	There is a problem with the identity or group defined on the DPM server. For example, you can get this error if there is no Identity configured with the client certificate that you are using.	Check the identity and group definitions on the DPM server.
Could not authenticate the server certificate with given CA file or problem reading CA file.	There is a problem with the root certificate on the PowerPath host.	Verify that the trusted root file is present on the host, that it has not expired, and that it resides in the location specified by the <b>clientTrustedRoots</b> attribute in the <code>rkm_init.conf</code> file.
Detected that a required argument passed used a format that is not supported or incorrect.	The command included a required argument but its format was wrong.	Correct the command syntax and re-run the command.
Error occurred in accessing the cache database.	The SQLite database on the PowerPath host cannot be opened.	Check that the cache file exists and is readable.
Generic IO error occurred or cache password is not correct for an existing persistent cache file.	XCR_RKM error (10002).	Perform the following: <ol style="list-style-type: none"> <li>1. Stop the PowerPath Encryption service.</li> <li>2. Delete the cache files.</li> <li>3. Restart the PowerPath Encryption service.</li> </ol> Otherwise, remove PowerPath and then reinstall.
Invalid configuration format specified in the service configuration. The following service configuration elements are mandatory: <code>transportSvc</code> , <code>cacheSvc</code> .	The <code>rkm_svc.conf</code> file is missing either the <b>transportSvc</b> or <b>cacheSvc</b> attribute. Both attributes must be included in the <code>rkm_svc.conf</code> file.	Correct the <code>rkm_svc.conf</code> file and then re-run the command.
Invalid configuration specified in the initialization configuration.	The <code>rkm_init.conf</code> file contains invalid configuration information.	Correct the <code>rkm_init.conf</code> file and re-run the command.
Invalid service type specified in the initialization configuration. Service types must be one of the following values: <code>transportSvc</code> , <code>cacheSvc</code> , <code>logSvc</code> .	The <code>rkm_init.conf</code> file contains an invalid value for the <code>svcType</code> attribute. Valid values are <b>transportSvc</b> , <b>cacheSvc</b> , and <b>logSvc</b> .	
Key Class supplied to the key request operation is invalid.	The key class specified by this command is not defined on the DPM server.	Verify that the key class set on the PowerPath Encryption host corresponds to a key class defined on the DPM server. If necessary, change the default key class specified in the <code>rkm_keyclass.conf</code> file and run the <b>ckmadm setkeyclass</b> command.
Key ID supplied to the key request operation is invalid.		

**Table 44** DPM server error messages (page 2 of 2)

Message	Description	Suggested action
Key Manager Client could not complete the request at the Server. Protocol error detected during transport.	An error occurred when the key manager client on the PowerPath Encryption host attempted to communicate with the DPM server.	Check the key manager client configuration. Verify the contents of the following configuration files: <ul style="list-style-type: none"> <li>• rkm_init.conf</li> <li>• rkm_svc.conf</li> </ul>
Key Manager Client support library cannot be loaded by the Key Manager Client.	The KM_SUPPORT_LIB_PATH environment variable is not set on the PowerPath host. This variable should be set by the PowerPath installation process.	Set the KM_SUPPORT_LIB_PATH environment variable to point to the /etc/emc/rsa/rkm_client/lib/libkmsvcshlib.so library.
Key Manager Client support library cannot be located by the Key Manager Client; environment variable, KM_SUPPORT_LIB_PATH, must be defined for the current client installation.		
Local certificate problem.	There is a problem with the client credential file on the PowerPath host.	Verify that the client credential file is present on the host, that it has not expired, and that it resides in the location specified by the <b>clientCredentialFile</b> attribute in the rkm_init.conf file.
Memory allocation operation failure occurred.	The host does not have sufficient memory to perform the operation.	Free some memory on the host, or add more memory to it.
Pre-requisite for the requested operation is not found.	There was a problem finding one of the crypto headers (at the beginning of the encrypted data), or an algorithm is not available (such as in FIPS mode).	Check that the default key class on the PowerPath host uses a supported algorithm and mode.
Pre-requisite for the requested operation is unavailable.		
Protocol error detected during transport.	XCR_RKM error (10003). An error occurred when the key manager client on the PowerPath Encryption host attempted to communicate with the DPM server, or you are using a new RKM client version with an old DPM server version.	Check the key manager client configuration. Verify the contents of the following configuration files: <ul style="list-style-type: none"> <li>• rkm_init.conf</li> <li>• rkm_svc.conf</li> </ul> Upgrade to the latest DPM server version.
Required argument contains an incorrect or invalid value.	XCR_RKM (10022). The command syntax is incorrect.	Correct the command syntax and then re-run the command.
Required argument contains NULL.	The command did not include a required argument.	

## Key manager client messages

[Table 45 on page 245](#) lists key manager client error messages, descriptions, and possible remedies.

**Table 45** Key manager client error messages

Message	Description	Suggested action
Failed to load the configuration file.	The rkm_init.conf file could not be loaded.	Correct the rkm_init.conf file and re-run the command.
File open failed.	One of the following applies: <ul style="list-style-type: none"> <li>The <b>ckmadm setup</b> command could not open the rkm_init.conf file.</li> <li>The <b>ckmadm setconfig</b> command could not open the rkm_svc.conf file.</li> <li>The <b>ckmadm setkeyclass</b> command could not open the rkm_keyclass.conf file.</li> </ul>	<ul style="list-style-type: none"> <li>Make sure the configuration file exists in the /etc/emc/rsa/rkm_client/config directory.</li> <li>Specify the full path to the configuration file.</li> </ul>
Input line was too long.	The command syntax is incorrect.	Correct the command syntax and re-run the command.
Invalid argument.	The command included an invalid argument.	
Invalid line in configuration file.	The rkm_init.conf file contains invalid configuration information.	Correct the rkm_init.conf file and re-run the command.
Key Manager Client could not complete the request at the Server.	This occurs while attempting to encrypt and use WordPad to edit the files leading to a configuration script failure.	Whenever the conf.tmpl files are edited, either use notepad or if any other text editing software is used, then save the file as text document.
Operation not supported.	The command specified an unsupported operation.	Correct the command syntax and re-run the command.

## Lockbox service messages

[Table 46 on page 245](#) lists lockbox error messages, descriptions, and possible remedies.

**Table 46** Lockbox service error messages (page 1 of 2)

Message	Description	Suggested action
File open failed.	One of the following applies: <ul style="list-style-type: none"> <li>The <b>ckmadm setup</b> command could not open the rkm_init.conf file.</li> <li>The <b>ckmadm setconfig</b> command could not open the rkm_svc.conf file.</li> <li>The <b>ckmadm setkeyclass</b> command could not open the rkm_keyclass.conf file.</li> </ul>	<ul style="list-style-type: none"> <li>Make sure the configuration file exists in the /etc/emc/rsa/rkm_client/config directory.</li> <li>Specify the full path to the configuration file.</li> </ul>



**Table 46** Lockbox service error messages (page 2 of 2)

Message	Description	Suggested action
Insufficient memory.	The host does not have sufficient memory to perform the operation.	Free some memory on the host, or add more memory to it.
Invalid argument.	The command included an invalid argument.	Correct the command syntax and re-run the command.
LLT link down.	When VCS 5.0 is installed and configured in a 2 node cluster, 1 node hangs during the LLT configuration stage. It displays the error message and the system goes to state from where even dump could not be initiated.	Before configuring VCS, ensure you have installed the tar file.

## ckmadm error messages

[Table 47 on page 246](#) lists **ckmadm** command error messages, descriptions, and possible remedies.

**Table 47** ckmadm error messages (page 1 of 2)

Message	Description	Suggested action
Ckm initialization error -9. Could not open lockbox.	Running the powermt config command for the first time after a fresh installation of PowerPath, results in an error message logged by EV_AGENT (navisphere agent).	Stop (before installing PowerPath) and re-start the Naviagent after installing PowerPath to avoid seeing the error message
Error with CST encryption service.	The lockbox encryption service could not encrypt the password entered during the <b>ckmadm setup</b> command execution. This message is followed by a message from the lockbox service.	Refer to the following sections for a description of the message and recommended suggested actions: <ul style="list-style-type: none"> <li>• <a href="#">“RSA Data Protection Manager (DPM) server messages” on page 243</a></li> <li>• <a href="#">“Key manager client messages” on page 245</a></li> <li>• <a href="#">“Lockbox service messages” on page 245</a></li> </ul>
Failed to store key class info.	The <b>ckmadm setkeyclass</b> command was not able to store key class configuration information in the lockbox. This message is followed by a message from the DPM server, key manager client, or lockbox service that indicates the source of the problem.	
Failed to store svc info.	The <b>ckmadm setconfig</b> command was not able to store services information in the lockbox. This message is followed by a message from the DPM server, key manager client, or lockbox service that indicates the source of the problem.	

**Table 47** ckmdm error messages (page 2 of 2)

Message	Description	Suggested action
file <i>filename</i> larger than expected.	Configuration files can have a maximum of 100 entries. Either the rkm_init.conf, rkm_svc.conf, or rkm_keyclass.conf file has more than the maximum number of entries.	Edit the configuration file causing the error to contain 100 entries or less.
file <i>filename</i> open failed.	One of the following applies: <ul style="list-style-type: none"> <li>• The <b>ckmdm setup</b> command could not open the rkm_init.conf file.</li> <li>• The <b>ckmdm setconfig</b> command could not open the rkm_svc.conf file.</li> <li>• The <b>ckmdm setkeyclass</b> command could not open the rkm_keyclass.conf file.</li> </ul>	<ul style="list-style-type: none"> <li>• Make sure the configuration file exists in the /etc/emc/rsa/rkm_client/config directory.</li> <li>• Specify the full path to the configuration file.</li> </ul>
More key configs than supported.	The rkm_keyclass.conf file contains multiple key class entries. This version of PowerPath Encryption supports a single key class.	Edit the rkm_keyclass.conf file to contain a single key class.

## cstadmin error messages

[Table 48 on page 247](#) lists **cstadmin** command error messages, descriptions, and possible remedies.

**Table 48** cstadmin error messages (page 1 of 2)

Message	Description	Suggested action
Authorization to open the Lockbox failed.	<ul style="list-style-type: none"> <li>• Lockbox passphrase is invalid.</li> <li>• Lockbox passphrase callback is missing.</li> <li>• Lockbox passphrase callback failed.</li> </ul>	None.
Could not open file for reading.	The csp-app.xml file is missing.	Use any one of these solutions in the order shown below: <ol style="list-style-type: none"> <li>1. Check for the presence or permissions on /etc/emc/rsa/cst/config/csp-app.xml</li> <li>2. Re-create the entire configuration.</li> </ol>
Could not open file for write.	Configuration file is not writable.	None.
Could not open lockbox.	<ul style="list-style-type: none"> <li>• Directory does not exist.</li> <li>• Lockbox does not exist.</li> <li>• Lockbox exists but is not writable.</li> </ul>	Use any one of these solutions in the order shown below: <ol style="list-style-type: none"> <li>1. Check for the presence or permissions on /etc/emc/rsa/cst/config/csp.clb</li> <li>2. Recover the lockbox from backup on # cp csp.clb.bak csp.clb</li> <li>3. Re-create the entire configuration.</li> </ol>
Could not read lockbox contents.	Lockbox is corrupt.	None.

**Table 48** cstadmin error messages (page 2 of 2)

Message	Description	Suggested action
Error converting string to object.	<ul style="list-style-type: none"> <li>Configuration file is malformed.</li> <li>Bootstrap file is malformed.</li> <li>Manifest file is malformed.</li> </ul>	None.
Failed to load class from library.	Method is missing in the DLL.	None.
Failed to load shared library.	<ul style="list-style-type: none"> <li>DLL does not exist.</li> <li>DLL is corrupt.</li> <li>Dependent DLL is missing.</li> </ul>	<ul style="list-style-type: none"> <li>Look for libCSP* in /etc/emc/rsa/cst/lib.</li> <li>Reinstall libraries from PowerPath distribution into /etc/emc/rsa/cst/lib.</li> </ul>
File does not match signature.	<ul style="list-style-type: none"> <li>Configuration file signature does not match the signature in the Lockbox.</li> <li>Manifest file signature does not match the signature in the Lockbox.</li> </ul>	Re-create the entire configuration.
Initialization failed (-21).Caught exception while posting log record. The system hostname key is missing from the lockbox	This occurs while attempting to configure the RKM client on a Windows 2008 host, and if the installation of PowerPath is completed prior to the installation of ArcServe from CA.	Ensure that RKM Client configuration is done in the following sequence: <ol style="list-style-type: none"> <li>1. Install ArcServe from CA.</li> <li>2. Install PowerPath.</li> <li>3. Execute the D@RE configuration.</li> </ol>
ld.so.1: powervt: fatal: libCSP-C.so: open failed: No such file or directory killed	<ul style="list-style-type: none"> <li>A dynamic link library is missing or an incorrect setting exists for LD_LIBRARY_PATH.</li> <li>Either libCSP.so, libCSP-RT.so or libCSP-C.so is not found.</li> </ul>	<ul style="list-style-type: none"> <li>Run <b>ldd xcryptd</b> to see if libraries are present in /etc/emc/rsa/cst/lib.</li> <li>Verify that environment variable LD_LIBRARY_PATH (on AIX) is set correctly.</li> <li>Reinstall libraries from PowerPath distribution into /etc/emc/rsa/cst/lib.</li> </ul>
Lockbox file is corrupt and failed to load.	Lockbox is corrupt.	Use any one of these solutions in the order shown below: <ol style="list-style-type: none"> <li>1. Recover the lockbox from backup on # cp csp.clb.bak csp.clb.</li> <li>2. Re-create the entire configuration.</li> </ol>
The bootstrap path must be assigned prior to performing this operation.	Directory was not specified.	None.
The item name does not exist in the Lockbox.	<ul style="list-style-type: none"> <li>Lockbox does not contain the lockbox encryption keys.</li> <li>Bootstrap signature does not exist.</li> <li>Manifest file signature does not exist.</li> <li>Configuration file signature does not exist.</li> </ul>	None.
The signature of the bootstrap file does not match the expected value.	Bootstrap file signature does not match the signature in Lockbox	None.

## emcp\_xcryptd daemon error messages

Table 49 on page 249 lists emcp\_xcryptd daemon error messages, descriptions, and possible remedies.

**Table 49** emcp\_xcryptd error messages (page 1 of 2)

Message	Description	Suggested action
Ckm initialization error. <DPM_server_or_key_manager_client_error_message>	The key manager client on the PowerPath host could not be initialized. This message is followed by a message from the DPM server, key manager client, or lockbox service that indicates the source of the problem.	Refer to the following sections for a description of the message and recommended suggested actions: <ul style="list-style-type: none"> <li>“RSA Data Protection Manager (DPM) server messages” on page 243</li> <li>“Key manager client messages” on page 245</li> <li>“Lockbox service messages” on page 245</li> </ul>
Daemon creation failed <err_no>	PowerPath <b>emcp_xcryptd</b> daemon was not able to start.	Check that the DPM server is available, properly configured, and able to communicate with the key manager client on the PowerPath host.
Failed to refresh the info for device <device> and error is <PowerPath_error_message>.	PowerPath <b>emcp_xcryptd</b> daemon was not able to refresh the encryption information.	Take the suggested action recommended by the PowerPath message.
Failed to start the daemon with err <error>.	PowerPath <b>emcp_xcryptd</b> daemon was not able to start.	Check that the DPM server is available, properly configured, and able to communicate with the key manager client on the PowerPath host.
Lookup Key Failed for device <device_name>. <DPM_server_or_key_manager_client_error_message>	PowerPath Encryption could not retrieve a key. This message is followed by a message from the DPM server, key manager client, or lockbox service that indicates the source of the problem.	Refer to the following sections for a description of the message and recommended suggested actions: <ul style="list-style-type: none"> <li>“RSA Data Protection Manager (DPM) server messages” on page 243</li> <li>“Key manager client messages” on page 245</li> <li>“Lockbox service messages” on page 245</li> </ul>
Lookup Key Failed for device Major <dev_no> Minor <dev_no>. <DPM_server_or_key_manager_client_error_message>		
Lookup Key Failed for user dev <dev_no>. <DPM_server_or_key_manager_client_error_message>		

**Table 49** emcp\_xcryptd error messages (page 2 of 2)

Message	Description	Suggested action
Request to wait for key lookup failed. <PowerPath_error_message>	A request to wait for a key lookup failed.	Take the suggested action recommended by the PowerPath message.
Sending empty key to the driver failed. <PowerPath_error_message>		
Sending key to the driver failed. <PowerPath_error_message>		
Unable to allocate memory for keyBlob.	The host does not have sufficient memory to perform the operation.	Free some memory on the host, or add more memory to it.
Unable to get key class. <DPM_server_or_key_manager_client_error_message>	PowerPath Encryption could not retrieve a key class. This message is followed by a message from the DPM server, key manager client, or lockbox service that indicates the source of the problem.	Refer to the following sections for a description of the message and recommended suggested actions: <ul style="list-style-type: none"> <li>• “RSA Data Protection Manager (DPM) server messages” on page 243</li> <li>• “Key manager client messages” on page 245</li> <li>• “Lockbox service messages” on page 245</li> </ul>

## powervt error messages

Table 50 on page 250 lists powervt command error messages, descriptions, and possible remedies.

**Table 50** powervt error messages (page 1 of 3)

Message	Description	Suggested action
-on, -off or -info cannot happen at the same time.	The <b>powervt xcrypt</b> command included more than one operand.	Re-run the <b>powervt xcrypt</b> command, specifying one of the following required operands: -on, -off, -info
-on, -off or -info required	The <b>powervt xcrypt</b> command did not include a required operand.	
Command or option not supported.	The <b>powervt</b> command specified an invalid subcommand or argument.	Run the <b>powervt -h</b> command to check command syntax, and then re-run the command.
Device is already encrypted.	The <b>powervt xcrypt -on</b> command attempted to enable encryption on a previously encrypted device, which is prohibited.	Run the <b>powervt xcrypt -info</b> command to determine which devices have already been encrypted.
<Device> is encrypted but no key available.	The <b>powervt xcrypt -info -dev all</b> command returns the error after the host name or IP address has been changed.	Perform the following: <ol style="list-style-type: none"> <li>1. Run <b>cstadmin sign-file -cstdir="&lt;install-dir&gt;\EMC\RSA\CST\config" "*.xml"</b> to re-sign the CST lockbox.</li> <li>2. Run <b>powervt xcrypt -info -dev all -refresh</b> to refresh the cache file.</li> </ol>

**Table 50** powervt error messages (page 2 of 3)

Message	Description	Suggested action
Device is not virtualized.	The <b>powervt resize</b> command failed to resize the virtualized device because the device is not encrypted.	Perform the following: <ol style="list-style-type: none"> <li>1. Run the <b>powervt xcrypt -info</b> command to obtain the encryption status of the device.</li> <li>2. Run the <b>powervt xcrypt -on</b> command to encrypt the device.</li> <li>3. Run the <b>powervt resize</b> command.</li> </ol>
Device required.	The <b>powervt</b> command either did not specify a device name or specified an invalid device.	Re-run the <b>powervt</b> command, using a PowerPath pseudo device name to identify a valid device.
Device state is unknown.	The <b>powervt xcrypt -on</b> command was not able to enable encryption on a device because it could not determine the device state.	Check that device is online and ready. This message is displayed for not-ready devices such as BCVs that are established and not split.
Failed to access the device.	The <b>powervt resize</b> command failed because the device is inaccessible.	Ensure that the device is online and ready.
Failed to change the device state.	The <b>powervt xcrypt -on</b> command could not enable encryption on a device.	Use the AIX format utility to add a VTOC label to the device, and then re-run the <b>powervt xcrypt -on</b> command.
Failed to disable encryption on <i>&lt;device_name&gt;. &lt;additional_message&gt;</i>	The <b>powervt xcrypt -off</b> command could not disable encryption on the specified device. This message is followed by a secondary message that describes the cause of the problem.	Refer to the following sections for a description of the message and recommended suggested actions: <ul style="list-style-type: none"> <li>• <a href="#">“RSA Data Protection Manager (DPM) server messages” on page 243</a></li> <li>• <a href="#">“Key manager client messages” on page 245</a></li> <li>• <a href="#">“Lockbox service messages” on page 245</a></li> </ul>
Failed to enable encryption on <i>&lt;device_name&gt;. &lt;additional_message&gt;</i>	The <b>powervt xcrypt -on</b> command could not enable encryption on the specified device. This message is followed by a secondary message that describes the cause of the problem.	Refer to the following sections for a description of the message and recommended suggested actions: <ul style="list-style-type: none"> <li>• <a href="#">“RSA Data Protection Manager (DPM) server messages” on page 243</a></li> <li>• <a href="#">“Key manager client messages” on page 245</a></li> <li>• <a href="#">“Lockbox service messages” on page 245</a></li> </ul>
Failed to get device from kernel. <i>&lt;PowerPath_kernel_message&gt;</i>	The <b>powervt</b> command could not get a device from the kernel. This message is followed by an additional message from the PowerPath kernel that indicates the source of the problem.	Refer to the following sections for a description of the message and recommended suggested actions: <ul style="list-style-type: none"> <li>• <a href="#">“RSA Data Protection Manager (DPM) server messages” on page 243</a></li> <li>• <a href="#">“Key manager client messages” on page 245</a></li> <li>• <a href="#">“Lockbox service messages” on page 245</a></li> </ul>
Failed to get encryption information on device <i>&lt;device_name&gt;. &lt;additional_message&gt;</i>	The <b>powervt</b> command could not determine the encryption status of the target device. This message is followed by an additional message that indicates the source of the problem.	Refer to the following sections for a description of the message and recommended suggested actions: <ul style="list-style-type: none"> <li>• <a href="#">“RSA Data Protection Manager (DPM) server messages” on page 243</a></li> <li>• <a href="#">“Key manager client messages” on page 245</a></li> <li>• <a href="#">“Lockbox service messages” on page 245</a></li> </ul>

Table 50 powervt error messages (page 3 of 3)

Message	Description	Suggested action
Failed to get geometry of the device.	The <b>powervt resize</b> command failed because it could not determine the geometry of the target device.	Ensure that the device is online and ready.
Failed to resize <i>&lt;device_name&gt;</i> . <i>&lt;additional_message&gt;</i>	The <b>powervt resize</b> command was not able to resize the specified device. This message is followed by an additional message that indicates the source of the problem.	Search this chapter for the secondary message and take the recommended suggested action.
File contains no valid devices.	The file containing the list of device names supplied with the command contained no valid devices.	Make sure that the devices named in the file are valid and then re-run the <b>powervt xcrypt</b> command.
Insufficient privilege.	A user who does not have root privileges attempted to run a <b>powervt</b> command.	Re-run the <b>powervt</b> command from an account with root privileges.
Invalid device or not under PowerPath control.	The <b>powervt</b> command failed because the target device is invalid or not under PowerPath control. For example, the device may be under VxVM control.	Re-run the <b>powervt</b> command, specifying a valid device. If the device is under the control of a volume manager, release it before running the <b>powervt</b> command. For a device under VxVM control, for example, run the <b>vxdisk offline</b> command before running the <b>powervt</b> command.
No such extension	ppvt error (0x8100)	Remove PowerPath and then reinstall.
Required license is missing or expired.	A <b>powervt</b> command exited because a valid PowerPath Encryption license is not registered on the host.	Register a valid PowerPath Encryption license on the host.
Root privileges required.	A user who does not have root privileges attempted to run a <b>powervt</b> command.	Re-run the <b>powervt</b> command from an account with root privileges.
Specific device names are required.	The <b>powervt</b> command either did not specify a device name or specified an invalid device.	Re-run the <b>powervt</b> command, using a PowerPath pseudo device name to identify a valid device.
The device was already encrypted.	The <b>powervt xcrypt -on</b> command attempted to enable encryption on a previously encrypted device, which is prohibited.	Use the <b>powervt xcrypt -info</b> command to determine which devices have already been encrypted.
This action will destroy data on <i>&lt;device_name&gt;</i> , continue? yes/[no]	Turning encryption on or off will destroy all data on the device, or the PowerPath Encryption Driver does not get installed.	Turn encryption on or off, or reinstall the encryption driver.
Unable to open <i>&lt;filename&gt;</i> for reading.	The file containing the list of device names supplied with the command could not be opened.	Verify the file name and then re-run the <b>powervt xcrypt</b> command.
Valid devices required.	The <b>powervt</b> command either did not specify a device name or specified an invalid device.	Re-run the <b>powervt</b> command, using a PowerPath pseudo device name to identify a valid device.

## xcrypt\_config error messages

Table 51 on page 253 lists xcrypt\_config command error messages, descriptions, and possible remedies.

**Table 51** xcrypt\_config error messages

Message	Description	Suggested action
A valid license key for PowerPath Data Encryption was not found. Please register a valid license and re-run <b>xcrypt_config</b> .	The <b>xcrypt_config</b> script exited because a valid PowerPath Encryption license is not registered on the host.	Register the PowerPath Encryption license on the host. Re-run the <b>xcrypt_config</b> command.
Can't find PowerPath package base.	PowerPath is not properly installed on the host.	Check the PowerPath installation and correct any problems that you find. Refer to your platform specific Installation and Administration Guide for more information.
Encryption is already configured. Exiting.	The <b>xcrypt_config</b> script has already been run on this host.	<ul style="list-style-type: none"> <li>• Use the <b>ckmadm</b> command to update the key manager client configuration after initial PowerPath Encryption configuration.</li> <li>• Use the <b>cstadmin new-passphrase</b> command to change the lockbox passphrase after initial PowerPath Encryption configuration.</li> </ul>
Failed to change the device state.	XCR error (0x5130). The LUN is not write-enabled.	Ensure you write-enable the LUN prior to encrypting a device in a AIX environment.



# INDEX

## Symbols

/etc/syslog.conf 242

## A

abandoned LU cleanup 142  
active path mode 52, 56, 92  
adaptive failover policy 100  
alive path state 52, 57  
audit log messages  
    description 229  
    example 230  
    format and description 229  
    migration examples 231  
autostandby (asb)  
    aging period 90  
    IOsPerFailure limit 91  
    mode 42, 52, 56

## B

basic failover policy 100

## C

ceilingValue, Host Copy ceiling 150  
checking consistency 26  
ckmadm commands 167  
ckmadm getconfig 169  
ckmadm getkeyclass 170  
ckmadm help 168  
ckmadm setconfig 171  
ckmadm setkeyclass 173  
ckmadm setup 172  
ckmadm upgrade 174  
CLARiiON LUN names 40, 55, 58, 111  
CLARiiON optimization failover policy 100  
CLARiiON user-assignable LUN names 40, 55, 58, 111  
cleaning up a migration 142  
committing a migration 145  
configuration, custom  
    loading 77, 78  
configuration, removing dead paths 26  
configuring PowerPath Encryption software 188  
cstadmin initialize 176  
cstadmin commands 175  
cstadmin new-passphrase 177  
cstadmin sign-file 178  
custom configuration  
    loading 77, 78

## D

dead path state 52, 57  
dead paths, removing 26  
degraded HBA path status 37, 48

device  
    display, field explanations 40  
    monitoring 33, 45, 65, 71, 72  
    parameter specification for powermt commands 23  
displaying migration information 152  
displaying the software version 167

## E

emcpadm 118  
emcpadm check\_mappings 119  
emcpadm debugging 133  
emcpadm examples 123  
emcpadm export\_mappings 120  
emcpadm import\_mappings 120  
emcpadm print\_mappings 121  
emcpadm renamepseudo 128, 129, 131  
EMCPADM\_DBG environment variable 123  
EMCPADM\_ERR environment variable 123  
emcpcfg backup 134  
emcpcfg cleanup 134  
emcpcfg list\_cfgs 134  
emcpcfg restore 134  
emcpcfg script 134  
emcpcfg set\_cfgs\_limit 134  
emcphostid 115  
emcpmigd.log file 233  
EmcPowerpathEncryption 242  
emcpreg 113  
encryption error messages  
    ckmadm command 246  
    cstadmin command 247  
    DPM (Data Protection Manager) server 243  
    emcp\_xcryptd daemon 249  
    key manager client 245  
    lockbox service 245  
    messages sources 242  
    powervt command 250  
    xcrypt\_config command 253  
environment variable  
    EMCPADM\_DBG 123  
    EMCPADM\_ERR 123  
    KM\_SUPPORT\_LIB\_PATH 244  
    LD\_LIBRARY\_PATH 248  
    PP\_DISPLAY\_TIME\_STAMP 33, 59, 60, 72  
    PP\_DISPLAY\_WIDTH 34, 46, 49, 53, 60, 73  
    PP\_SHOW\_ALUA\_FAILOVER\_MODE 41  
    PP\_SHOW\_CLAR\_LUN\_NAMES 40, 55, 59  
    PP\_SHOW\_VERSION\_STYLE 112  
error log messages 232  
    configuring logging 229  
    description 232  
    example 232  
    format 232  
    responding to 233

Event Viewer - Application 242

## F

failed HBA path status 37, 48  
failover policy. *See also* load-balancing and failover policy 99

## H

help commands  
    ckmadm 168  
    powermig 147  
    powervt 180  
host bus adapter (HBA)  
    monitoring 33, 45, 65, 71, 72  
    path status 37, 48  
host copy migration log file 233  
hostcopy\_ceiling option 150

## I

I/O path mode 52, 56  
I/O path state 52, 57  
Invista encapsulation. *See* virtual encapsulation (INVE)

## K

key class  
    changing the default 173  
    returning the default 170  
key manager client, configuring and managing 167  
KM\_SUPPORT\_LIB\_PATH environment variable 244

## L

latency (path)  
    displaying 49  
    monitoring 94  
    threshold 95  
LD\_LIBRARY\_PATH environment variable 248  
least blocks failover policy 100  
least I/Os failover policy 100  
license, checking 28  
licensing errors 196  
load-balancing and failover policy  
    adaptive policy 100  
    basic failover policy 100  
    CLARiiON optimization policy 100  
    least blocks policy 100  
    least I/Os policy 100  
    no redirect policy 101  
    request policy 101  
    round robin policy 101  
    setting 99  
    streamIO policy 101  
    Symmetrix optimization policy 101  
lockbox, managing of 175  
logical unit, obtaining encryption status 186  
LUN names 40, 55, 58, 111

## M

message sources

    encryption 242  
    multipathing 197  
messages. *See* audit log messages, encryption error messages, error log messages, message sources, Migration Enabler error messages, PowerPath messages, Windows event ID messages  
Migration Enabler error messages 234  
migration handle, description of 142  
migration handle, retrieving 146  
mode, setting for paths 92  
monitoring HBAs and devices 33, 45, 65, 71, 72

## N

nice names. *See* CLARiiON LUN names  
no redirect policy 101

## O

optimal HBA path status 37, 48

## P

passphrase, changing the lockbox 177  
path latency, information 49  
paths  
    HBA status 37, 48  
    latency monitoring on 94  
    latency threshold 95  
    restoration 84  
    setting the mode 92  
pausing a migration 151  
policy. *See* load-balancing and failover policy  
powermig cleanup 142  
powermig commit 145  
powermig getHandle 146  
powermig help 147  
powermig info 148  
powermig options command 150  
powermig pause 151  
powermig query 152  
powermig recover 154  
powermig resume 155  
powermig selectSource 156  
powermig selectTarget 157  
powermig setup 158  
powermig sync 162  
powermig throttle 163  
powermig undoRedirect 166  
powermig version 167  
powermigcl 168  
    powermigcl config 171  
    powermigcl display 169  
    powermigcl help 174  
    powermigcl unconfig 172  
    powermigcl version 173  
powermt check 26  
powermt check\_registration 28  
powermt commands 17  
powermt disable hba 32  
powermt display 33, 45, 65, 71, 72  
powermt display alua 45

powermt display latency 49  
 powermt display options 53, 58, 68, 75, 80, 110  
 powermt display paths 60  
 powermt display perf 60, 65  
 powermt display perf dev 60, 65  
 powermt display perf paths 60, 65  
 powermt display ports 71, 72  
 powermt load 77, 78  
 powermt restore 84  
 powermt set autostandby 88  
 powermt set autostandby\_agingperiod 90  
 powermt set autostandby\_iopflimit 91  
 powermt set mode 92  
 powermt set path\_latency\_monitor 94  
 powermt set path\_latency\_threshold 95  
 powermt set perf\_mon 97  
 powermt set periodic\_autorestore 98  
 powermt set policy 99  
 powermt set ppfailfast 105  
 powermt set reactive\_autorestore 105  
 powermt set streamio\_threshold 107  
 powermt set transparent. *See also* powermt set policy 15  
 powermt set write\_throttle 108  
 powermt set write\_throttle\_queue 109  
 powermt utility  
   executing 22  
   install directory 21  
   required privileges 22  
 powermt utility installation directory 21  
 powermt version 112  
 PowerPath Encryption, configuring 188  
 PowerPath messages  
   licensing errors 196  
   message prefixes 196  
   message sources 197  
   PowerPath messages 197  
   Windows application event IDs 227  
   Windows system event IDs 225  
 powervt commands 179  
 powervt help 180  
 powervt resize 181  
 powervt update 182  
 powervt version 183  
 powervt xcrypt -info 186  
 powervt xcrypt -off 185  
 powervt xcrypt -on 184  
 PP\_DISPLAY\_TIME\_STAMP environment variable 33, 59, 60,  
   72  
 PP\_DISPLAY\_WIDTH environment variable 34, 46, 49, 53, 60,  
   73  
 PP\_SHOW\_ALUA\_FAILOVER\_MODE environment variable 41  
 PP\_SHOW\_CLAR\_LUN\_NAMES environment variable 40, 55,  
   59  
 PP\_SHOW\_VERSION\_STYLE environment variable 112  
 emcpadm getfree 121

## R

reconfiguring the PowerPath Encryption host 189  
 redirection of I/O, stopping 166  
 renaming pseudo devices 128, 129, 131

request failover policy 101  
 restoring paths 84  
 resuming a paused migration 155  
 retrieving available pseudo devices 121  
 RKM (RSA Key Manager). *See* DPM (Data Protection Manager)  
   243  
 RKM\_Config.bat batch file 172, 188  
 RKM\_Config.exe (Windows) 172  
 round robin failover policy 101

## S

selecting the source 156  
 selecting the target 157  
 setting the host ID 115  
 setting up a migration 158  
 standby path mode 52, 56, 92  
 stopping the redirection of I/O 166  
 streamIO failover policy 101  
 Symclone. *See also* TimeFinder/Clone 158, 159, 163, 164  
 Symmetrix optimization failover policy 101  
 synchronizing the source and target 162  
 syslog 229, 242

## T

throttle value 163  
 TimeFinder/Clone  
   QoS (Quality of Service) values 165  
   technology type 158  
   throttle value 165

## U

unlicensed path mode 52, 56  
 user-assignable LUN names 40, 55, 58, 111

## V

verbose version output 112  
 version  
   Migration Enabler 167  
   multipathing 112  
   powermt command 112  
   powervt command 183  
 virtual encapsulation (INVE) 53, 158, 231, 233  
 virtual logical unit  
   disabling encryption on 185  
   managing of 179  
   resizing of 181

## W

Windows application event ID messages 227  
 Windows system event ID messages 225  
 write throttling  
   enabling 108  
   queue depth, setting 109

## X

xcrypt\_config script 172, 188

