EMC® PowerPath/VE for VMware vSphere

Version 5.9 and Minor Releases

Remote CLI Guide

P/N 302-000-237
REV 02

EMC²

CONTENTS

Contents

PREFACE

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

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

Note: This document was accurate at publication time. 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/VE 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.

Related documentation

The following EMC publications provide additional information:

◆ *PowerPath/VE 5.9 and Minor Releases for VMware vSphere Release Notes*

◆ *PowerPath/VE 5.9 and Minor Releases for VMware vSphere Installation and Administration Guide*

These PowerPath manuals are updated periodically. Electronic versions of the updated manuals are available on the EMC Online Support site: https://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.

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:
  http://www.EMC.com/clariionsupport

◆ VNX OE Storage System Support website:
  www.emc.com/vnxsupport

Limited PowerPath functions are supported by the Unisphere™ application for VNX™ Operating Environment (OE) systems. Refer to the VNX OE Storage System Support website (www.emc.com/vnxsupport).

Limited PowerPath functions are supported by the Navisphere® and the Unisphere™ applications for CLARiiON systems. Refer to CLARiiON Storage-System Support website (www.emc.com/clariionsupport).

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 1 Revision History

Revision	Date	Description
02	December 18, 2013	First release of PowerPath/VE 5.9 Service Pack 1 for VMware vSphere.
01	September 13, 2013	First release of PowerPath/VE 5.9 for VMware vSphere.

Conventions used in this document

EMC uses the following conventions for special notices:

⚠️**CAUTION**

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

NOTICE

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

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

IMPORTANT

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

Typographical conventions

EMC uses the following type style conventions in this document:

Normal	Used in running (nonprocedural) text for: • Names of interface elements, such as names of windows, dialog boxes, buttons, fields, and menus • Names of resources, attributes, pools, Boolean expressions, buttons, DQL statements, keywords, clauses, environment variables, functions, and utilities • URLs, pathnames, filenames, directory names, computer names, links, groups, service keys, file systems, and notifications
Bold	Used in running (nonprocedural) text for names of commands, daemons, options, programs, processes, services, applications, utilities, kernels, notifications, system calls, and man pages Used in procedures for: • Names of interface elements, such as names of windows, dialog boxes, buttons, fields, and menus • What the user specifically selects, clicks, presses, or types
Italic	Used in all text (including procedures) for: • Full titles of publications referenced in text • Emphasis, for example, a new term • Variables
`Courier`	Used for: • System output, such as an error message or script • URLs, complete paths, filenames, prompts, and syntax when shown outside of running text
`Courier bold`	Used for specific user input, such as commands
`Courier italic`	Used in procedures for: • Variables on the command line • User input variables
< >	Angle brackets enclose parameter or variable values supplied by the user
[ ]	Square brackets enclose optional values
|	Vertical bar indicates alternate selections — the bar means "or"
{ }	Braces enclose content that the user must specify, such as x or y or z
…	Ellipses indicate nonessential information omitted from the example

Where to get help

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

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

https:// support.emc.com

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

Your comments

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

techpubcomment@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.

CHAPTER 1
Introduction

EMC® PowerPath/VE for VMware supports the remote PowerPath (rpowermt) command line interface (CLI). Use the rpowermt command on a remote PowerPath server to configure and manage PowerPath/VE for VMware on vSphere ESXi hosts.

rpowermt command summary

Table 2 lists the PowerPath/VE rpowermt commands that you can use to manage PowerPath and paths under its control on vSphere ESXi hosts.

Table 2 rpowermt command summary (page 1 of 2)

Command	Description
"rpowermt check_registration"	Checks the state of the PowerPath license for a vSphere host.
"rpowermt display"	Displays the state of HBAs and devices managed by PowerPath.
"rpowermt display alua"	Displays ALUA state for all paths to an ALUA device.
"rpowermt display bus"	Displays all paths (bus-port specific information) between the HBAs on a host and the storage system ports that the HBAs can access.
"rpowermt display latency"	Displays storage-system path information, including statistics on path latency monitoring.
"rpowermt display options"	Displays option settings for storage-system classes.
"rpowermt display perf"	Displays storage device performance metrics for all logical devices.
"rpowermt display perf bus"	Displays storage device performance metrics for each bus.
"rpowermt display ports"	Displays port-specific information.
"rpowermt help"	Displays rpowermt command syntax.
"rpowermt register"	Registers a PowerPath license for a vSphere host.
"rpowermt restore"	Tests and restores paths.
"rpowermt server lockbox"	
"rpowermt set autostandby"	Specifies a path that has been placed into standby by PowerPath/VE automatically either using the proximity-based autostandby algorithm (asb:prox) or the IOSPerFailure-based autostandby algorithm (asb:iopf).
"rpowermt set autostandby agingperiod"	Sets the length of time that a path stays in autostandby due to an iopf (IOsPerFailure) trigger before it returns to active.
"rpowermt set autostandby iopflimit"	Specifies the average number of I/Os between failures to automatically change that path to autostandby mode.
"rpowermt set mode"	Sets paths to either active or standby mode.
"rpowermt set path_latency_monitor"	Enables or disables path latency monitoring on all paths.
"rpowermt set path_latency_threshold"	Sets a time limit for the completion of I/O on all paths.
"rpowermt set perfmon"	Enables or disables data collection for performance monitoring of all devices.
"rpowermt set periodic_autorestore"	Enables or disables periodic autorestore.
"rpowermt set policy"	Changes the load-balancing and failover policy for devices managed by PowerPath.

rpowermt syntax summary

Following is a summary of the syntax of all rpowermt commands:

```
 rpowermt <command>
[class={all|netapp|symm|vnx|clariion|vplex|invista|netapp|hitachi|hpxp|ess|xtremio|generic}]
 host=<hostname>
 [username=<username> [password=<password>|no_password]]
 [cim_sessionid=<CIM sessionid>]
 rpowermt check_registration
 rpowermt display [every=<#seconds>] [width=<#col>]
 rpowermt display alua dev=<device>|all [class=<class>|all] [every=<#seconds>]
[width=<#col>]
 rpowermt display bus [class=<class>|all] [every=<#seconds>]
[width=<#col>]
 rpowermt display latency [dev=<device>|all] [class=<class>|all]
       [every=<#seconds>] [width=<#col>]
 rpowermt display options
 rpowermt display perf dev=<device>|all [continuous] [verbose] [xml] [nowait]
 rpowermt display perf bus [continuous] [verbose] [xml] [nowait]
 rpowermt display ports [dev=<device>|all] [class=<class>|all]
       [every=<#seconds>] [width=<#col>]
 rpowermt help [<command>]
 rpowermt register [force]
 rpowermt restore [hba=<hba#>|all] dev=<path>|<device>|all [class=<class>|all]
 rpowermt set autostandby={on|off|reinitialize} [trigger={prox|iopf}]
 rpowermt set autostandby iopflimit=<value>
 rpowermt set autostandby agingperiod=<#days>
 rpowermt set mode={active|standby} [hba=<hba#>|all]
       [dev=<path>|<device>|all] [class=<class>|all] [force]
 rpowermt set path_latency_monitor={on|off} [force]
 rpowermt set path_latency_threshold=<#seconds>|<#milliseconds>ms [force]
 rpowermt set perfmon={on [interval=<#seconds> | off}
 rpowermt set periodic_autorestore={on|off} [class=<class>|all] [force]
 rpowermt set policy={ad|co|lb|li|rr|so|si}  [dev=<device>|all]
[class=<class>|all] [force]
 rpowermt set reactive_autorestore={on|off} [class=<class>|all] [force]
 rpowermt set streamio_threshold=<threshold_count> [dev=<device>|all]
```

```
[class=<class>|all]
rpowermt setup add_host [host_file=<filename>]
      [username=<username>] [password=<password>]
rpowermt setup list_hosts [host_file=<filename>]
rpowermt setup remove_host [host_file=<filename>]
rpowermt setup update_host [host_file=<filename>]
      [username=<username>] [password=<password>]
rpowermt setup verify_hosts [host_file=<filename]
rpowermt unregister [force]
rpowermt update lun_names
rpowermt version host=<hostname> username=<username>
      [password=<password>|no_password] [cim_sessionid=<CIM sessionid>]
```

rpowermt help

Purpose   Displays a summary of the command syntax. If a specific command is specified, it displays the syntax for that command. The usage menu is also displayed when invalid commands or options are entered.

Syntax   `rpowermt help <command>`

Arguments   `<command>`

A valid rpowermt command.

Example   Following is a sample output when no command is specified.

```
Usage:
rpowermt <command>
[class={all|netapp|symm|vnx|clariion|vplex|invista|netapp|hitachi|hpxp|ess|xtremio|generic}]
host=<hostname>
[username=<username> [password=<password>|no_password]]
[cim_sessionid=<CIM sessionid>]
rpowermt check_registration
rpowermt display [every=<#seconds>] [width=<#col>]
rpowermt display alua dev=<device>|all [class=<class>|all] [every=<#seconds>]
[width=<#col>]
rpowermt display bus [class=<class>|all] [every=<#seconds>]
[width=<#col>]
rpowermt display latency [dev=<device>|all] [class=<class>|all]
      [every=<#seconds>] [width=<#col>]
rpowermt display options
rpowermt display perf dev=<device>|all [continuous] [verbose] [xml] [nowait]
rpowermt display perf bus [continuous] [verbose] [xml] [nowait]
rpowermt display ports [dev=<device>|all] [class=<class>|all]
      [every=<#seconds>] [width=<#col>]
rpowermt help [<command>]
rpowermt register [force]
rpowermt restore [hba=<hba#>|all] dev=<path>|<device>|all [class=<class>|all]
rpowermt set autostandby={on|off|reinitialize} [trigger={prox|iopf}]
rpowermt set autostandby iopflimit=<value>
rpowermt set autostandby agingperiod=<#days>
rpowermt set mode={active|standby} [hba=<hba#>|all]
      [dev=<path>|<device>|all] [class=<class>|all] [force]
rpowermt set path_latency_monitor={on|off} [force]
rpowermt set path_latency_threshold=<#seconds>|<#milliseconds>ms [force]
rpowermt set perfmon={on [interval=<#seconds> | off}
rpowermt set periodic_autorestore={on|off} [class=<class>|all] [force]
rpowermt set policy={ad|co|lb|li|rr|so|si}  [dev=<device>|all]
[class=<class>|all] [force]
rpowermt set reactive_autorestore={on|off} [class=<class>|all] [force]
```

```
rpowermt set streamio_threshold=<threshold_count> [dev=<device>|all]
[class=<class>|all]
rpowermt setup add_host [host_file=<filename>] [username=<username>]
      [password=<password>]
rpowermt setup list_hosts [host_file=<filename>]
rpowermt setup remove_host [host_file=<filename>]
rpowermt setup update_host [host_file=<filename>]
      [username=<username>] [password=<password>]
rpowermt setup verify_hosts [host_file=<filename]
rpowermt unregister [force]
rpowermt update lun_names
rpowermt version host=<hostname> username=<username>
      [password=<password>|no_password] [cim_sessionid=<CIM sessionid>]
```

Following is a sample output when a command (in this example, display) is specified.

```
C:\Documents and Settings\Administrator>rpowermt help display
display - Display multipathing information.
Usage:
  rpowermt <command> [class=all|symm|clariion|invista|netapp|hitachi|hpxp|ess|xtremio|generic]
            host=<hostname>
            [username=<username>] [[password=<password>|no_password]]
            [cim_sessionid=<cim sessionid>]
    rpowermt display
            [dev=<device|all> | class=<class|all>]
            [every=<#seconds>] [width=<#col>]
    rpowermt display latency
            [dev=<device|all> | class=<class|all>]
            [every=<#seconds>] [width=<#col>]
    rpowermt display options
    rpowermt display paths
            [class=<class|all>] [every=<#seconds>] [width=<#col>]
    rpowermt display ports
            [dev=<device|all> | class=<class|all>]
            [every=<#seconds>] [width=<#col>]
License search path: C:\Documents and Settings\Administrator\My
Documents\EMC\PowerPath\rpowermt;C:\
Documents and Settings\All Users\My Documents\EMC\PowerPath\rpowermt
Host file: C:\Documents and Settings\Administrator\My
Documents\EMC\PowerPath\rpowermt\lockbox.clb
```

Automatic license registration

PowerPath/VE supports the automatic license registration, or autoregistration, feature on served and unserved licenses. Any rpowermt command licenses the vSphere host with a valid PowerPath/VE for VMware vSphere license on the rpowermt server.

Commands excluded from autoregistration functionality are:

◆ rpowermt check_registration

◆ rpowermt register

◆ rpowermt setup

◆ rpowermt unregister

◆ rpowermt version

On the vSphere host, the PowerPath driver is functional from initial PowerPath/VE installation. There is no concept of unlicensed functionality on the vSphere host with respect to multipathing.

The rpowermt CLI determines the license state on each rpowermt command and if the state is unlicensed, the rpowermt CLI automatically attempts to license.

In a successful case, the following output appears:

```
# rpowermt display dev=all host=<hostname>
Remote host is not licensed. Attempted to automatically register.
PowerPath license is registered using automatic registration.
Pseudo name=emcpower6
Symmetrix ID=000192601669
Logical device ID=24BE
Standard UID=naa.60000970000192601669533032344245
state=alive; policy=SymmOpt; queued-IOs=0
...
```

Autoregistration can be disabled by setting an environment variable *PPMT_DISABLE_LICAUTO* to 1.

Manual registration is required if you have disabled autoregistration. Use the `rpowermt register` command.

CHAPTER 2
rpowermt command line utilities

This chapter describes the PowerPath/VE rpowermt command line utilities.

rpowermt check_registration

Purpose Displays the state of the PowerPath/VE license on a vSphere host.

Syntax
```
rpowermt check_registration host=<hostname>
    [username=<username> [password=<password>|no_password]]
    [cim_sessionid=<CIM sessionid>]
```

Arguments
`host=<hostname>`
Specifies the IP address or hostname of the vSphere host.

`username=<username>`
Specifies the user account on the vSphere host. If the username is not included in the command line, you are prompted to specify the user name.

`password=<password>`
Specifies the password associated with the specified username. If the password is not included on the command line, you are prompted to specify the user name.

`cim_sessionid=<CIM sessionid>`
Specifies the CIM session ID ticket output for vSphere host authorization.

Example Output similar to the following is returned for a vSphere host with a valid license:

```
rpowermt host=lcla224 check_registration
PowerPath License Information:
 ------------------------------

Host ID           : 4dc16302-1265-a5c4-0394-002219287b68
Type              : served (counted)
State             : licensed
Registered To     : EMC
Issue Date        : 24-may-2011
Feature           : PowerPathMP
Feature Version   : 5.4
Registering Svr   : rpowermt@<rpowermt_server>-><licensing_server>
License Count     : 12
Overdraft Count   : 0
License Server    : <licensing_server>
Days until expiration : 230
License search path: <license search path>/my_license.lic
License file(s): <license files>
```

rpowermt display

Purpose Displays information about HBAs and/or devices managed by PowerPath/VE.

The **`rpowermt display dev`** command includes an Array failover field, which displays the failover mode setting for VNX and CLARiiON system arrays. This setting determines the failover mode for a server's HBA port and is configured through the VNX and CLARiiON Unisphere and Navisphere software.

Note: After changing devices to ALUA mode, a host reboot is required for PowerPath/VE to recognize the change.

Syntax `rpowermt display [dev=<device>|all] [class=<class>|all]`

```
[every=<#seconds>] [width=<#col>] host=<hostname>
[username=<username> [password=<password>|no_password]]
[cim_sessionid=<cim sessionid>]
```

Arguments   **Note:** The every=‹#seconds› and width=‹#col› arguments can be used with any other option as required.

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

　　The supported device name formats for rpowermt display dev= are:

- PPVE pseudo device name

- Native path identifier- Format of the native device path is the ‹hwpath›:‹Cx:Ty:Lz› value.

`class={all|netapp|symm|vnx|clariion|vplex|invista|hitachi|hpxp|ess|xtremio|generic}`
　　Limits the command to the specified type of storage system. all specifies all storage-system types. The default is all.

　　Note: XtremIO devices of version 2.2 are managed under the `xtremio` class. Version 1.05 devices continue to be managed by `generic` storage class.

　　VNXe devices are included in the generic device class.

`every=<#seconds>`
　　Integer in the range 1 to 86400 that is specified in seconds, and the frequency with which 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.

`width=<#col>`
　　Sets the display width to a constant value within the range of 80 to 160 columns. This value applies only to the specific output and is not persistent. You can set a persistent value for the width=columns option with the PP_DISPLAY_WIDTH environment variable. This variable changes the width setting from 80 (the default) to the value you specify.

　　If the environment variable and command line values differ, the command line value takes precedence.

`host=<hostname>`
　　Specifies the IP address or hostname of the vSphere host.

`username=<username>`
　　Specifies the user account on the vSphere host. If the username is not included in the command line, you are prompted to specify the user name.

`password=<password>`
　　Specifies the password associated with the username. If the password is not included on the command line, you are prompted to specify it.

`cim_sessionid=<cim sessionid>`
　　Specifies the CIM session ID ticket output for vSphere host authorization.

Example
```
#rpowermt display dev=vmhba3:C0:T7:L1 host=192.0.2.0
Pseudo name=emcpower9
VPLEX ID=FCCNH999901457 , FNNNH9297777
Logical device ID=61111440222200010A000CF15444F95A5
Standard UID=naa.6006016061702c00cc3c077eb5d4e011
state=alive; policy=Adaptive; queued-IOs=0
===========================================================================

  -------------- Host -------------- Stor -  -- I/O Path --  -- Stats ---

### HW Path             I/O Paths    Interf.  Mode    State   Q-IOs   Errors

===========================================================================

1 vmhba3                C0:T7:L1     04       active   alive    0       0

2 vmhba2                C0:T4:L1     07       active   alive    0       0

2 vmhba2                C0:T5:L1     0A       active   alive    0       0

1 vmhba3                C0:T4:L1     07       active   alive    0       0

2 vmhba2                C0:T6:L1     0C       active   alive    0       0

1 vmhba3                C0:T5:L1     0A       active   alive    0       0

2 vmhba2                C0:T7:L1     04       active   alive    0       0

1 vmhba3                C0:T6:L1     0C       active   alive    0       0
```

The first form of the command (rpowermt display) yields the HBA display. The second form of the command (rpowermt display dev) yields the devices display.

For a more detailed description of these displays, including example displays and a discussion of the effects of zoning on the displays, see the *PowerPath/VE 5.9 for VMware vSphere Installation and Administration Guide*.

Table 3 explains the fields in the HBA display.

Table 3 Fields in the HBA display (page 1 of 2)

Field	Value	Description
‹storage_ system_ type› 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 r**powermt display** table contains information about one type of storage system, such as a Symmetrix, VNX, CLARiiON or third-party array.
Host Bus Adapters ###	Non-negative integer	PowerPath number for the HBA. When the **dev** option is used, the output of r**powermt display** 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	vSphere-assigned name for the HBA.
I/O Paths Summary	optimal degraded failed	Status of the paths originating from this HBA: • *optimal* means all paths are alive (usable). • *degraded* means one or more, but not all, paths from this HBA are dead (not usable). • *failed* means all paths are dead and no data is passing through this HBA.

Table 3 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 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/O Paths Total*	Total number of paths originating from this HBA that are dead (not usable).
Stats IO/sec	Non-negative integer	This field is blank for **rpowermt display,** unless it is used with the every parameter. Subsequent **rpowermt display** 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. **Note:** 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. PowerPath/VE periodic autorestore 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.

Storage System Interface	Alphanumeric string	Storage-system interface.
		 For Symmetrix systems, this has three parts:
		 • Interface type: Fibre Channel (FA) or SCSI (SA)
		 • Interface address: integer in the range 1 to 16
		 • Interface port: [abcd][AB]
		 For VNX and CLARiiON systems, this has one part:
		 • Interface port: SP [A-B][0-3]
		 For HP StorageWorks EVA this has one part:
		 • Interface port: [A-B][1-2]
		    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.

Table 8 describes the fields in the devices display.

Table 4 Fields in the devices display (page 1 of 3)

Field	Value	Description
Pseudo name	Alphanumeric string	Platform-specific value assigned by PowerPath to the PowerPath device.
‹*storage_ system_ type*› ID	Hexadecimal value or alphanumeric string	Identification number for the storage system on which the logical device is located. 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 not the same as the standard device identifier used by Navisphere. For VNX and CLARiiON systems only, the user-assignable LUN name, if available, is by default displayed in parenthesis after this identification number. To suppress 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".)
Logical device ID	Hexadecimal value or alphanumeric string	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.
Owner	default=SP *x*, current=SP *x*	Default and current owners of the logical device (VNX and CLARiiON systems only).
policy	Adaptive (ad) CLAROpt (co) LeastBlocks (lb) Least I/O (li) RoundRobin (rr) Stream I/O (si) SymmOpt (so)	Current load-balancing and failover policy for the device. Use **"rpowermt set policy"** to change the policy.
Array failover mode	1 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. **Note:** 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: • 1 - passive not ready (PNR) mode • 4 - asymmetric logical unit access (ALUA) mode VNX and CLARiiON failover mode appears only when hosts are connected to CLARiiON and VNX arrays. Stop all I/O on the host before changing the failover mode. After changing the failover mode, rebooting the host is required for PowerPath/VE to recognize the change. To prevent the failover mode from appearing in the rpowermt display output, set the PP_SHOW_ALUA_FAILOVER_MODE environment variable to "FALSE" or "false".
Host Bus Adapters ###	Non-negative integer	PowerPath number for the HBA. T When the **dev** option is used, the output of **rpowermt display** identifies the HBA by this same HBA number. This number is preserved across boots but is not preserved after configuration changes.

Table 4 Fields in the devices display (page 2 of 3)

Field	Value	Description
Host HW Path	Alphanumeric string	vSphere-assigned name for the HBA.
Host I/O Path	Alphanumeric string	The platform-specific device name for the path.

Stor Interf.	Alphanumeric string	Storage-system interface.
		For Symmetrix systems, this has three parts:
		• Interface type: Fibre Channel (FA) or SCSI (SA), or iSCSI (SE)
		• Interface address: integer in the range 1 to 16
		• Interface port: [abcd][AB]
		For VNX and CLARiiON systems, this has one part:
		• Interface port: SP [A-B][0-3]
		For HP StorageWorks EVA this has one part:
		• Interface port: [A-B][1-2]
		   A, B denotes the controller
		   1, 2 denotes the port
		For EMC Invista and VPLEX 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.)

I/O Path Mode	• active • standby • asb:prox • asb:iopf	Current path mode: • **active** 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. • **standby** 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 adjusted to preclude its use in normal operations. A standby path still can be selected if it is the best path for a request. • **asb** (for autostandby) specifies a path that has been placed into standby automatically either using the VPLEX proximity autostandby algorithm (**asb:prox**) or the intermittent I/O failure algorithm (**asb:iopf**). This path will not be used unless all other active paths are dead. **asb** path mode settings do not persist across reboots.
I/O Path State	• alive • dead	Current path state: • **alive** indicates the path is usable: PowerPath can direct I/O to this path. • **dead** indicates the path is not usable: PowerPath/VE will not direct I/O to this path. After marking the path dead and incrementing the Errors count, PowerPath/VE tests the path to see whether it is usable. If the test succeeds, PowerPath/VE marks the path alive; the path is then available for I/O. If the test fails, the path remains dead, and PowerPath/VE ignores it for subsequent I/O operations. If all the paths to a logical device are dead, PowerPath/VE retests each path.
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. **Note:** 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. This is because Q-IOs is inclusive of all types of I/Os filtered by PowerPath/VE, 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.

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

Field	Value	Description
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 rpowermt restore executes. PowerPath/VE periodic autorestore does not clear the error count. If both the dev and ports options are specified, a compressed version of the devices display is output. The compressed display includes a column listing the write throttle setting (on or off) of each storage system port and a column listing the write throttle queue depth for each port. To accommodate the extra information, the I/O Paths Mode and I/O Paths State columns of the normal devices display are merged into a Path Status column. Valid values for Path Status are as follows: • a- active: See Path mode, above • s- standby: See path mode, above • v- alive: see path state, above • d- dead: see path state, above

rpowermt display alua

Purpose   Displays ALUA state for all paths to an ALUA device, that is, storage systems that support the SCSI-standard ALUA. It 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/VE distributes I/O among paths in the active/optimized asymmetric access state only for optimal performance. Paths in the active/non-optimized asymmetric access state are used only when all paths in the active/optimized asymmetric access state have failed.

Syntax   ```
rpowermt display [dev=<device>|all] [class=<class>|all]
     [every=<#seconds>] [width=<#col>] host=<hostname>
     [username=<username> [password=<password>|no_password]]
     [cim_sessionid=<cim sessionid>]
```

**Arguments**　**Note:** The every=‹#seconds› and width=‹#col› arguments can be used with any other option as required.

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

　The supported device name formats for rpowermt display dev= are:

　• PPVE pseudo device name

　• Native path identifier - Format of the native device path is, ‹hwpath›:‹Cx:Ty:Lz› value

`class={all|netapp|symm|vnx|clariion|vplex|invista|hitachi|hpxp|ess|xtremio|generic}`
　Limits the command to the specified type of storage system. all specifies all storage-system types. The default is all.

　**Note:** XtremIO devices of version 2.2 are managed under the `xtremio` class. Version 1.05 devices continue to be managed by `generic` storage class.

　VNXe devices are included in the generic device class. If class is specified along with dev, the command is applied to paths that meet all specified constraints.

`every=<#seconds>`
　Specifies the frequency with which alua mode information is displayed. 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.

`width=<#col>`
　Sets the display width to a constant value within the range of 80 to 160 columns. This value applies only to the specific output and is not persistent. You can set a persistent value for the width=columns option with the PP_DISPLAY_WIDTH environment variable. This variable changes the width setting from 80 (the default) to the value you specify.

　If the environment variable and command line values differ, the command line value takes precedence.

`host=<hostname>`
　Specifies the IP address or hostname of the vSphere host.

`username=<username>`
　Specifies the user account on the vSphere host. If the username is not included in the command line, you are prompted to specify the user name.

`password=<password>`
　Specifies the password associated with the specified username. If the password is not included on the command line, you are prompted to specify the user name.

`cim_sessionid=<CIM sessionid>`
　Specifies the CIM session ID ticket output for vSphere host authorization.

# rpowermt display bus

**Purpose**
Displays all paths (bus-port specific information) between the HBAs on a host and the storage system ports that the HBAs can access.

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 the table applies.

**Syntax**
```
rpowermt display bus [class=<class>|all]
    [every=<#seconds>] [width=<#col>] host=<hostname>
    [username=<username> [password=<password>|no_password]]
    [cim_sessionid=<cim sessionid>]
```

**Arguments**
bus
Displays one line per bus.

class={all|netapp|symm|vnx|clariion|vplex|invista|hitachi|hpxp|ess|xtremio|generic}
Limits the command to the specified type of storage system. all specifies all storage-system types. The default is all.

> **Note:** XtremIO devices of version 2.2 are managed under the xtremio class. Version 1.05 devices continue to be managed by generic storage class.

VNXe devices are included in the generic device class.

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.

width=<#col>
Sets the display width to a constant value within the range of 80 to 160 columns. This value applies only to the specific output and is not persistent. You can set a persistent value for the width=columns option with the PP_DISPLAY_WIDTH environment variable. This variable changes the width setting from 80 (the default) to the value you specify.

If the environment variable and command line values differ, the command line value takes precedence.

If you set the  every=‹#seconds› option along with the width=‹#col› option, the rpowermt display output widens or narrows to accommodate changes you make to the window width.

host=<hostname>
Specifies the IP address or hostname of the vSphere host.

username=<username>
Specifies the user account on the vSphere host. If the username is not included in the command line, you are prompted to specify the user name.

password=<password>
Specifies the password associated with the specified username. If the password is not included on the command line, you are prompted to specify the user name.

```
cim_sessionid=<CIM sessionid>
```
Specifies the CIM session ID ticket output for vSphere host authorization.

# rpowermt display latency

**Purpose**   Displays information on path latency monitoring, including the last and longest I/O completion times for each path.

Valid latency values are displayed only if "rpowermt set path_latency_monitor" is set to on. A value of 0 is displayed otherwise.

The rpowermt 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

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

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

**Syntax**
```
rpowermt display latency [dev=<device>|all] [class=<class>|all]
     [every=<#seconds>] [width=<#col>] host=<hostname>
     [username=<username> [password=<password>|no_password]]
     [cim_sessionid=<cim sessionid>]
```

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

The supported device name formats for rpowermt display dev= are:

- PPVE pseudo device name

- Native path identifier - Format of the native device path is, ‹hwpath›:‹Cx:Ty:Lz› value.

`class={all|netapp|symm|vnx|clariion|vplex|invista|hitachi|hpxp|ess|xtremio|generic}`
Limits the command to the specified type of storage system. all specifies all storage-system types. The default is all.

> **Note:** XtremIO devices of version 2.2 are managed under the `xtremio` class. Version 1.05 devices continue to be managed by `generic` storage class.

VNXe devices are included in the generic device class.

`every=<#seconds>`
Specifies the frequency with which path_latency_monitor 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.

`width=<#col>`
Sets the display width to a constant value within the range of 80 to 160 columns. This value applies only to the specific output and is not persistent. You can set a persistent value for the width=columns option with the PP_DISPLAY_WIDTH environment variable. This variable changes the width setting from 80 (the default) to the value you specify.

If the environment variable and command line values differ, the command line value takes precedence.

`host=<hostname>`
Specifies the IP address or hostname of the vSphere host.

`username=<username>`
Specifies the user account on the vSphere host. If the username is not included in the command line, you are prompted to specify the user name.

`password=<password>`
Specifies the password associated with the specified username. If the password is not included on the command line, you are prompted to specify the user name.

`cim_sessionid=<CIM sessionid>`
Specifies the CIM session ID ticket output for vSphere host authorization.

Table 5 describes the fields in the latency display.

**Table 5**  Fields in the latency display (page 1 of 2)

| Field | Value | Description |
|---|---|---|
| Host Bus Adapters ### | Non-negative integer | PowerPath number for the HBA.<br>When the **dev** option is used, the output of **rpowermt display** identifies the HBA by this same HBA number.<br>This number is preserved across boots but is not preserved after configuration changes. |
| Host Bus Adapters HW Path | Alphanumeric string | vSphere-assigned name of the HBA. |
| Storage System ID | Hexadecimal value or alphanumeric string | Identification number for the storage system on which the logical device is located. |

**Table 5** Fields in the latency display (page 2 of 2)

| Field | Value | Description |
|-------|-------|-------------|
| Storage System Interface | Alphanumeric string | Storage-system interface.<br>For Symmetrix systems, this has three parts:<br>• Interface type: Fibre Channel (FA) or SCSI (SA)<br>• Interface address: integer in the range 1 to 16<br>• Interface port: [abcd][AB]<br>For VNX and CLARiiON systems, this has one part:<br>• Interface port: SP [A-B][0-3]<br>For HP StorageWorks EVA this has one part:<br>• Interface port: [A-B][1-2]<br>  A, B denotes the controller.<br>  1, 2 denotes the port.<br>For Hitachi Lightning, Hitachi TagmaStore and HP-UX systems, this is the port name on the array.<br>For IBM systems, this is a two-digit hexadecimal port ID that can be decoded to a Bay/Adapter/Port on the system. (IBM command-line interface documentation describes how to decode the port ID.)<br>For EMC Invista systems, this is a two-digit hexadecimal number that can be decoded to a virtual target on the Invista system.<br>For EMC VPLEX systems, this is a two-digit hexadecimal number that can be decoded to an array port WWN on the VPLEX system. |
| I/O Path Mode | • active<br>• standby<br>• asb:prox<br>• asb:iopf | Current path mode:<br>• **active** 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.<br>• **standby** 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 adjusted to preclude its use in normal operations. A standby path still can be selected if it is the best path for a request.<br>• **asb** (for autostandby) specifies a path that has been placed into standby automatically either using the VPLEX proximity autostandby algorithm (**asb:prox**) or the intermittent I/O failure algorithm (**asb:iopf**). This path will not be used unless all other active paths are dead. **asb** path mode settings do not persist across reboots. |
| I/O Path State | • alive<br>• dead | Current path state:<br>• **alive** indicates the path is usable: PowerPath can direct I/O to this path.<br>• **dead** indicates the path is not usable: PowerPath/VE will not direct I/O to this path. After marking the path dead and incrementing the Errors count, PowerPath/VE tests the path to see whether it is usable. If the test succeeds, PowerPath/VE marks the path alive; the path is then available for I/O. If the test fails, the path remains dead, and PowerPath/VE ignores it for subsequent I/O operations. If all the paths to a logical device are dead, PowerPath/VE retests each path. |
| 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. |

# rpowermt display options

**Purpose**    Shows the following option settings for each storage system class:

◆ Whether path latency monitoring and the path latency threshold are enabled, and if one of the value is set

◆ Whether CLARiiON user-assignable LUN names are displayed

◆ Whether periodic autorestore and reactive autorestore are enabled

- Autostandby settings

- Performance data collection settings

To change settings, use the rpowermt set command.

You can add a timestamp to the rpowermt display every 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, year, time in hours, minutes, and seconds that increment until you exit the rpowermt display command.

TIME_SECONDS starts a counter from zero that increments until you exit the rpowermt display command.

If the environment variable is not defined or if it is defined but not assigned a proper value, then the timestamp is not displayed.

**Syntax**
```
rpowermt display options host=<hostname>
    [username=<username> [password=<password>|no_password]]
    [cim_sessionid=<CIM sessionid>]
```

**Arguments**

host=<hostname>
Specifies the IP address or hostname of the vSphere host.

username=<username>
Specifies the user account on the vSphere host. If the username is not included in the command line, you are prompted to specify the user name.

password=<password>
Specifies the password associated with the specified username. If the password is not included on the command line, you are prompted to specify the user name.

cim_sessionid=<CIM sessionid>
Specifies the CIM session ID ticket output for vSphere host authorization.

# rpowermt display paths

**Purpose**   Displays path-specific information.

**Syntax**
```
rpowermt display paths [class=<class>|all]
    [every=<#seconds>] [width=<#col>] host=<hostname>
    [username=<username> [password=<password>|no_password]]
    [cim_sessionid=<cim sessionid>]
```

**Arguments**
```
class={all|netapp|symm|vnx|clariion|vplex|invista|hitachi|hpxp|ess|xtremio|generic}
```
Limits the command to the specified type of storage system. all specifies all storage-system types. The default is all.

Note: XtremIO devices of version 2.2 are managed under the `xtremio` class. Version 1.05 devices continue to be managed by `generic` storage class.

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

every=<#seconds>
: Specifies the frequency with which port 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.

width=<#col>
: Sets the display width to a constant value within the range of 80 to 160 columns. This value applies only to the specific output; it is not persistent. You can set a persistent value for the width=columns option with the PP_DISPLAY_WIDTH environment variable. This variable changes the width setting from 80 (the default) to the value you specify.

If the environment variable and command line values differ, the command line value takes precedence.

host=<hostname>
: Specifies the IP address or hostname of the vSphere host.

username=<username>
: Specifies the user account on the vSphere host. If the username is not included in the command line, you are prompted to specify the user name.

password=<password>
: Specifies the password associated with the specified username. If the password is not included on the command line, you are prompted to specify the user name.

cim_sessionid=<CIM sessionid>
: Specifies the CIM session ID ticket output for vSphere host authorization.

**Example**    With no arguments, the display shows path information for each storage class:

```
# rpowermt display paths host=111.222.222.111
Symmetrix logical device count=2
================================================================================
----- Host Bus Adapters --------- ------ Storage System -----   - I/O Paths -
###  HW Path                       ID           Interface     Total     Dead
================================================================================
   1 vmhba3                        000190100286   FA   9aA        2        0
   1 vmhba3                        000190100286   FA   9bA        2        0
   2 vmhba4                        000190100286   FA   9aA        2        0
   2 vmhba4                        000190100286   FA   9bA        2        0
CLARiiON logical device count=3
================================================================================
----- Host Bus Adapters --------- ------ Storage System -----   - I/O Paths -
###  HW Path                       ID           Interface     Total     Dead
================================================================================
   1 vmhba3                        HK190807490051  SP A4         3        0
   1 vmhba3                        HK190807490051  SP B4         3        0
   2 vmhba4                        HK190807490051  SP A4         3        0
   2 vmhba4                        HK190807490051  SP B4         3        0
```

Table 6 describes the fields in the path display.

Table 6 Fields in the path display

Field	Value	Description
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. Each rpowermt display 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 **dev** option is used, the output of **rpowermt display** 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	vSphere-assigned name of the HBA.
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:
		• Interface type: Fibre Channel (FA) or SCSI (SA)
		• Interface address: integer in the range 1 to 16
		• Interface port: [abcd][AB]
		For VNX and CLARiiON systems, this has one part:
		• Interface port: SP [A-B][0-3]
		For HP StorageWorks EVA this has one part:
		• Interface port: [A-B][1-2]
		  A, B denotes the controller.
		  1, 2 denotes the port.
		For Hitachi Lightning, Hitachi TagmaStore and HP-UX 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.

I/O Paths Total	Non-negative integer	Total number of paths that originate from this HBA. 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/O Paths Total	Total number of paths originating from this HBA that are dead (not usable).

With a valid storage class, the display shows path information for the specified storage class:

```
# rpowermt display paths class=symm host=111.222.222.111
Symmetrix logical device count=2
===============================================================================
----- Host Bus Adapters --------- ------ Storage System -----    - I/O Paths -
###  HW Path                         ID          Interface    Total    Dead
===============================================================================
   1 vmhba3                        000190100286     FA  9aA       2        0
```

```
1 vmhba3                          000190100286    FA  9bA        2        0
2 vmhba4                          000190100286    FA  9aA        2        0
2 vmhba4                          000190100286    FA  9bA        2        0
```

rpowermt display perf

Purpose   Displays storage device performance metrics for all logical devices. The output has one line per path.

Syntax
```
rpowermt display perf dev=<device>|all
[continuous] [verbose] [xml]  [nowait] host=<hostname>
[username=<username> [password=<password>|no_password]]
[cim_sessionid=<cim sessionid>]
```

Arguments   dev=<device>|all
   Displays information about the specified devices. all specifies all devices.

   The supported device name formats for rpowermt display perf dev= are:

   - PPVE pseudo device name

   - Native path identifier - Format of the native device path is, ‹hwpath›:‹Cx:Ty:Lz› value.

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

verbose
   Displays latency bins 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 an XML format after the normal output. To redirect the performance data to an XML file, add xml 2› ‹filename.xml› and optionally, a full path name, to the rpowermt display perf dev command. For example, running rpowermt display perf dev=emcpower22 xml 2› perf.xml displays the performance data onscreen and creates a file named perf.xml that contains the data in the XML format. If you redirect the XML metrics to a file, then the XML data is not displayed on the screen.

nowait
   If this option is specified, rpowermt 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.

host=<hostname>
   Specifies the IP address or hostname of the vSphere host.

username=<username>
   Specifies the user account on the vSphere host. If the username is not included in the command line, you are prompted to specify the user name.

password=<password>

Specifies the password associated with the specified username. If the password is not included on the command line, you are prompted to specify the user name.

```
cim_sessionid=<CIM sessionid>
```
Specifies the CIM session ID ticket output for vSphere host authorization.

Example   The following output is the example of rpowermt display perf dev=all.

Note: The dash is displayed for output not observed.

```
# rpowermt host=<rpowermt server> display perf dev=1 verbose
Timestamp = 00:51:27 UTC, 12 Aug 2013
Sample Interval = 60
Pseudo name=emcpower1
VNX ID=FNM00130800788
Standard UID=naa.60060160071034007 8dcafd126f4e211
state=alive; policy=CLAROpt; queued-IOs=0
                              KB<=4    4<KB<=8  8<KB<=128    KB>128        All
Read  bytes/s                   -        -         -           -          -
Write bytes/s                   -        -         -           -          -
Total bytes/s                                                             -
Read  Avg Response ms           -        -         -           -          -
Write Avg Response ms           -        -         -           -          -
All   Avg Response ms                                                     -
Read  IOs/s                     -        -         -           -          -
Write IOs/s                     -        -         -           -          -
Total IOs/s                                                               -
==============================================================================
--------------- Host --------------     Metrics  --------------  Retry   Error
###   HW Path              I/O Paths              Reads---Writes  delta   delta
==============================================================================
   1  vmhba4               C0:T6:L0
                              Low Latency (ms)      -       -        0       0
                              High Latency (ms)     -       -
                              Lat <=    1ms (#)     0       0
                     1ms < Lat <=   10ms (#)        0       0
                    10ms < Lat <= 100ms (#)         0       0
                              Lat >  100ms (#)      0       0
                                    IOs/sec (#)     -       -
                                    Bytes/sec       -       -
   1  vmhba4               C0:T7:L0
                              Low Latency (ms)      -       -        0       0
                              High Latency (ms)     -       -
                              Lat <=    1ms (#)     0       0
                     1ms < Lat <=   10ms (#)        0       0
                    10ms < Lat <= 100ms (#)         0       0
                              Lat >  100ms (#)      0       0
                                    IOs/sec (#)     -       -
                                    Bytes/sec       -       -
   2  vmhba5               C0:T3:L0
                              Low Latency (ms)      -       -        0       0
                              High Latency (ms)     -       -
                              Lat <=    1ms (#)     0       0
                     1ms < Lat <=   10ms (#)        0       0
                    10ms < Lat <= 100ms (#)         0       0
                              Lat >  100ms (#)      0       0
                                    IOs/sec (#)     -       -
                                    Bytes/sec       -       -
   2  vmhba5               C0:T7:L0
                              Low Latency (ms)      -       -        0       0
                              High Latency (ms)     -       -
                              Lat <=    1ms (#)     0       0
                     1ms < Lat <=   10ms (#)        0       0
                    10ms < Lat <= 100ms (#)         0       0
```

```
                          Lat >  100ms (#)        0        0
                               IOs/sec (#)        -        -
                               Bytes/sec          -        -
[root@vmq216048 ~]#
```

The following output is the example of rpowermt display perf dev=all verbose.

Note: The dash is displayed for output not observed.

```
# rpowermt host=<rpowermt server> display perf dev=all verbose
Timestamp = 00:53:27 UTC, 12 Aug 2013
Sample Interval = 60
Pseudo name=emcpower0
Symmetrix ID=000187910018
Logical device ID=06C7
Standard UID=mpx.emc_Symm_000187910018_vol_06c7
state=alive; policy=SymmOpt; queued-IOs=0
                          KB<=4    4<KB<=8  8<KB<=128    KB>128       All
Read  bytes/s              -        -        -          -        -
Write bytes/s              -        -        -          -        -
Total bytes/s                                                      -
Read  Avg Response ms      -        -        -          -        -
Write Avg Response ms      -        -        -          -        -
All   Avg Response ms                                             -
Read  IOs/s                -        -        -          -        -
Write IOs/s                -        -        -          -        -
Total IOs/s                                                        -
===========================================================================
--------------- Host --------------    Metrics -------------- Retry  Error
###  HW Path            I/O Paths           Reads---Writes    delta  delta
===========================================================================
1 vmhba4                C0:T2:L1
                           Low Latency (ms)      -        -        0        0
                           High Latency (ms)     -        -
                           Lat <=    1ms (#)      0        0
                     1ms < Lat <=   10ms (#)      0        0
                     10ms < Lat <= 100ms (#)      0        0
                           Lat >  100ms (#)       0        0
                                IOs/sec (#)       -        -
                                Bytes/sec         -        -
  1 vmhba4                C0:T5:L1
                           Low Latency (ms)      -        -        0        0
                           High Latency (ms)     -        -
                           Lat <=    1ms (#)      0        0
                     1ms < Lat <=   10ms (#)      0        0
                     10ms < Lat <= 100ms (#)      0        0
                           Lat >  100ms (#)       0        0
                                IOs/sec (#)       -        -
                                Bytes/sec         -        -
  2 vmhba5                C0:T2:L1
                           Low Latency (ms)      -        -        0        0
                           High Latency (ms)     -        -
                           Lat <=    1ms (#)      0        0
                     1ms < Lat <=   10ms (#)      0        0
                     10ms < Lat <= 100ms (#)      0        0
                           Lat >  100ms (#)       0        0
                                IOs/sec (#)       -        -
                                Bytes/sec         -        -
  2 vmhba5                C0:T6:L1
                           Low Latency (ms)      -        -        0        0
                           High Latency (ms)     -        -
                           Lat <=    1ms (#)      0        0
                     1ms < Lat <=   10ms (#)      0        0
                     10ms < Lat <= 100ms (#)      0        0
```

```
                                    Lat >  100ms (#)       0         0
                                       IOs/sec (#)         -         -
                                       Bytes/sec           -         -
    . . .
```

Table 7 explains the fields in the performance metrics display.

Table 7 Fields in the performance metrics display (page 1 of 2)

Field	Description
Timestamp	The period since the last metrics were computed.
Sample Interval	The period of time interval in seconds. It is the interval is from which the data was gathered.
Pseudo name	Platform-specific value assigned by PowerPath to the PowerPath device.
state	The state of the storage device from which the metrics are gathered. The state can be either dead or alive.
policy	The selected load-balancing policy for the storage device. Based on the storage device, it could be: • Adaptive (ad) • CLAROpt (co) • LeastBlocks (lb) • Least I/O (li) • RoundRobin (rr) • Stream I/O (si) • SymmOpt (so)
queued- IOs	Number of I/O requests queued to this PowerPath device.
Sz‹=4K 4K‹Sz‹=8K 8K‹Sz‹=128K Sz›128K All	The I/O size buckets for each storage device and for all storage devices put together. In KB or MB.
Read bytes/s	The read per second for different I/O size buckets for each storage device, in KB/second or MB/second.
Write bytes/s	The write per second for different I/O size buckets for each storage device, in KB/second or MB/second.
Total bytes/s	The total throughput per second value of all reads and writes for all storage devices, in KB/second or MB/second.
Read Avg Response (ms)	The average Read response time for all Read and Writes in the Sample Interval for all different I/O size buckets for each storage device, in milliseconds.
Write Avg Response (ms)	The average write response time for all Read and Writes in the Sample Interval for different I/O size buckets for each storage device, in milliseconds.
All Avg Response (ms)	The total average response time for all Read and Writes in the Sample Interval for different I/O size buckets for each storage device, in milliseconds.
Read I/Os	The read I/Os per second rate value for each consolidated I/O size buckets, in KB/second.
Write I/Os	The write I/Os per second rate value for each consolidated I/O size buckets, in KB/second.
Total I/Os	The total I/Os per second for all reads and writes, in KB/second.
IOs/sec	The number of IOs per second on the path.
Bytes/sec	The path level throughput to PMI statistics.
Host HW Path	The platform-specific hardware descriptive name for the path.

Table 7 Fields in the performance metrics display (page 2 of 2)

Field	Description
Host I/O Path	The 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 latency bin. The latency bins in the verbose output can be used to verify that dead, standby (both manual and autostandby), and non-optimal paths are not handling application reads and writes under normal circumstances. Even standby and non-optimal paths can have application reads and writes routed to them under extremely heavy loads, and such over low I/O would show up in the latency bins.
Latency Reads	The high and low water marks for Read counts per each device path in the interval, in microseconds.
Latency writes	The high and low water marks for Write counts per each device path in the interval, in microseconds.
Retry delta	The number of retry attempts to restore a failed I/O through another path to the storage device during the interval.
Error delta	The number of alive to dead paths transition during the interval.

When there is no data for a device or there is no device specified, the following error messages are displayed:

```
[root@RHEL ~]# rpowermt host=<rpowermt server> display perf dev
ERROR: No value specified with dev.  dev=<value> required.
Usage:
  rpowermt display perf dev=<device>|all [continuous] [verbose] [xml] [nowait]
          host=<hostname>
          [username=<username> [password=<password>|no_password]]
          [cim_sessionid=<cim sessionid>]
  rpowermt display perf bus [continuous] [verbose] [xml] [nowait]
          host=<hostname>
          [username=<username> [password=<password>|no_password]]
          [cim_sessionid=<cim sessionid>]
[root@RHEL ~]# rpowermt host=<rpowermt server> display perf dev=
ERROR: "dev" value is not valid.
Usage:
  rpowermt display perf dev=<device>|all [continuous] [verbose] [xml] [nowait]
          host=<hostname>
          [username=<username> [password=<password>|no_password]]
          [cim_sessionid=<cim sessionid>]
  rpowermt display perf bus [continuous] [verbose] [xml] [nowait]
          host=<hostname>
          [username=<username> [password=<password>|no_password]]
          [cim_sessionid=<cim sessionid>]
[root@RHEL ~]# rpowermt host=<rpowermt server> display perf
ERROR: Invalid options.
Usage:
  rpowermt display perf dev=<device>|all [continuous] [verbose] [xml] [nowait]
          host=<hostname>
          [username=<username> [password=<password>|no_password]]
          [cim_sessionid=<cim sessionid>]
  rpowermt display perf bus [continuous] [verbose] [xml] [nowait]
          host=<hostname>
          [username=<username> [password=<password>|no_password]]
          [cim_sessionid=<cim sessionid>]
[root@RHEL ~]#
```

rpowermt display perf bus

Purpose Displays storage device performance metrics for each bus. If performance monitoring has just been turned on and continuous is specified, then the first output might take up to two sample intervals to get the performance data. The subsequent output will also wait until the specified interval time.

Syntax
```
rpowermt display perf bus
[continuous] [verbose] [xml] [nowait] host=<hostname>
[username=<username> [password=<password>|no_password]]
[cim_sessionid=<cim sessionid>]
```

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

verbose
   Displays latency bins 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 an XML format after the normal output. To redirect the performance data to an XML file, add xml 2› ‹filename.xml› and optionally, a full path name, to the rpowermt display perf dev command. For example, running the rpowermt display perf bus xml 2› perf.xml displays the performance data onscreen 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, rpowermt 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.

host=<hostname>
   Specifies the IP address or hostname of the vSphere host.

username=<username>
   Specifies the user account on the vSphere host. If the username is not included in the command line, you are prompted to specify the user name.

password=<password>
   Specifies the password associated with the specified username. If the password is not included on the command line, you are prompted to specify the user name.

cim_sessionid=<CIM sessionid>
   Specifies the CIM session ID ticket output for vSphere host authorization.

Example The following output is an example of rpowermt display perf bus for each bus or ITL nexus.

```
# rpowermt host=<rpowermt server> display perf bus
Timestamp = 00:59:27 UTC, 12 Aug 2013
Sample Interval = 60
Symmetrix logical device count=27
==============================================================================
----- Host Bus Adapters --------- ------ Storage System -----    Retry    Error
```

```
### HW Path                            ID              Interface    delta   delta
================================================================================
   1 vmhba4

                                  000187910018  FA   4aA              0       0
                                     Latency----Reads--Writes
                      Low Latency (ms)           -        -
                     High Latency (ms)           -        -
   1 vmhba4
                                  000187910018  FA   4bA              0       0
                                     Latency----Reads--Writes
                      Low Latency (ms)           -        -
                     High Latency (ms)           -        -
   1 vmhba4
                                  000195700879  FA   9fB              0       0
                                     Latency----Reads--Writes
                      Low Latency (ms)           -        -
                     High Latency (ms)           -        -
   1 vmhba4
                                  000195700879  FA  10fB              0       0
                                     Latency----Reads--Writes
                      Low Latency (ms)           -        -
                     High Latency (ms)           -        -
   2 vmhba5
                                  000187910018  FA   4bA              0       0
                                     Latency----Reads--Writes
                      Low Latency (ms)           -        -
                     High Latency (ms)           -        -
   2 vmhba5
                                  000187910018  FA   4aA              0       0
                                     Latency----Reads--Writes
                      Low Latency (ms)           -        -
                     High Latency (ms)           -        -
   2 vmhba5
                                  000195700879  FA  10fB              0       0
                                     Latency----Reads--Writes
                      Low Latency (ms)           -        -
                     High Latency (ms)           -        -
   2 vmhba5
                                  000195700879  FA   9fB              0       0
                                     Latency----Reads--Writes
                      Low Latency (ms)           -        -
                     High Latency (ms)           -        -
```

The following output is an example of rpowermt display perf bus verbose for each bus or ITL nexus:

```
~ # /opt/emc/powerpath/bin/powermt display perf bus verbose
Timestamp = 16:36:56 UTC, 07 Feb 2013
Sample Interval = 60

Symmetrix logical device count=11
================================================================================
----- Host Bus Adapters --------- ------ Storage System -----   Retry   Error
### HW Path                            ID              Interface    delta   delta
================================================================================
   1 vmhba2
                                  000192601669  FA  12fA              0       0
                                     Latency----Reads--Writes
                          Low (ms)               -        -
                         High (ms)               -        -
                  Lat <=   1ms (#)               0        0
           1ms < Lat <=  10ms (#)               0        0
          10ms < Lat <= 100ms (#)               0        0
                  Lat >  100ms (#)               0        0
```

```
                                       IOs/sec (#)         -         -
                                       Bytes/sec           -         -
         1 vmhba2
                                 000192601699  FA 12fA             0         0
                                    Latency----Reads--Writes
                                    Low (ms)    0.125    0.177
                                    High (ms)     100     15.8
                            Lat <=    1ms (#)   12308    14004
                     1ms < Lat <=   10ms (#)    1715       33
                    10ms < Lat <= 100ms (#)     116        3
                            Lat >  100ms (#)       1        0
                                    IOs/sec (#)   235      234
                                    Bytes/sec     965K     958K
         2 vmhba3
                                 000192601669  FA 12fA             0         0
                                    Latency----Reads--Writes
                                    Low (ms)      -        -
                                    High (ms)     -        -
                            Lat <=    1ms (#)       0        0
                     1ms < Lat <=   10ms (#)       0        0
                    10ms < Lat <= 100ms (#)        0        0
                            Lat >  100ms (#)        0        0
                                    IOs/sec (#)    -        -
                                    Bytes/sec      -        -
         2 vmhba3
                                 000192601699  FA 12fA             0         0
                                    Latency----Reads--Writes
                                    Low (ms)    0.128    0.173
                                    High (ms)    63.2     18.3
                            Lat <=    1ms (#)   12345    14082
                     1ms < Lat <=   10ms (#)    1610       17
                    10ms < Lat <= 100ms (#)     132        1
                            Lat >  100ms (#)       0        0
                                    IOs/sec (#)   234      235
                                    Bytes/sec     962K     962K
         2 vmhba3
                                 000192601669  FA 11fA             0         0
                                    Latency----Reads--Writes
                                    Low (ms)      -        -
                                    High (ms)     -        -
                            Lat <=    1ms (#)       0        0
                     1ms < Lat <=   10ms (#)       0        0
                    10ms < Lat <= 100ms (#)        0        0
                            Lat >  100ms (#)        0        0
                                    IOs/sec (#)    -        -
                                    Bytes/sec      -        -
         2 vmhba3
                                 000192601699  FA 11fA             0         0
                                    Latency----Reads--Writes
                                    Low (ms)    0.129    0.178
                                    High (ms)    75.1     55.4
                            Lat <=    1ms (#)   12219    14075
                     1ms < Lat <=   10ms (#)    1683       32
                    10ms < Lat <= 100ms (#)     152       15
                            Lat >  100ms (#)       0        0
                                    IOs/sec (#)   234      235
                                    Bytes/sec     959K     964K
```

Table 8 explains the fields in the performance metrics display.

Table 8 Fields in the performance metrics display

Field	Description
Timestamp	The period since the last metrics were computed.
Sample Interval	The period of time interval in seconds. It is the interval is from which the data was gathered.
Device class name count	Storage device class name and the total number of devices accessed through the buses.
HW Path	The platform-specific hardware descriptive name for the path.
ID	The platform-specific ID for the bus.
Latency	When the verbose option is specified, the Read and Write I/Os per second count for each device path is shown under each latency bin. The latency bins in the verbose output can be used to verify that dead, standby (both manual and autostandby), and non-optimal paths are not handling application reads and writes under normal circumstances. Even standby and non-optimal paths can have application reads and writes routed to them under extremely heavy loads, and such over low I/O would show up in the latency bins.
IOs/sec	The number of IOs per second on the path.
Bytes/sec	The path level throughput to PMI statistics.
Latency Reads	The high and low water marks for Read counts per each device path in the interval, in microseconds.
Latency writes	The high and low water marks for Write counts per each device path in the interval, in microseconds.
Retry delta	The number of retry attempts to restore a failed I/O through another path to the storage device during the interval.
Error delta	The number of alive to dead paths transition during the interval.

When there is no data for a device or there is no device specified, the following error message is returned:

```
Device(s) not found.
```

rpowermt display ports

Purpose Displays port information for the specified device or for all devices.

The device number can either be the PowerPath/VE pseudo device name (for example, emcpower4 or simply, 4) or the standard UUID (for example, naa.600601606f3023009069be840dell), which is the vSphere-assigned identifier for the PowerPath/VE device.

rpowermt display ports includes the Array failover mode field, which displays 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 the device display fields (see "rpowermt display" on page 14).

After changing devices to ALUA mode, reboot the host for PowerPath/VE to recognize the change.

Limited PowerPath/VE functions are supported by the Unisphere application for VNX systems and Navisphere application for CLARiiON systems. Refer to the storage system websites for more information:

◆ http://www.emc.com/vnxsupport

◆ http://www.emc.com/clarrionsupport

The rpowermt display ports command displays the same fields as the rpowermt display dev=device command. however, when you add the dev option to the rpowermt display ports command, it displays the additional path status field.

The Path status field has two parts:

◆ Mode, which has a value of:

  • a: active

  • s: standby

    – asb:prox

    – asb:iopf

◆ State, which has a value of:

  • v: alive

  • d: dead

Syntax

```
rpowermt display ports [dev=<device>|all] [class=<class>|all]
[every=<#seconds>] [width=<#col>] host=<hostname>
[username=<username> [password=<password>|no_password]]
[cim_sessionid=<cim sessionid>]
```

Arguments

`dev=<device>|all`
Displays port information, one line per port.

The supported device name formats for rpowermt display dev= are:

• PPVE pseudo device name

• Native path identifier- Format of the native device path is, ‹hwpath›:‹Cx:Ty:Lz› value.

`class={all|netapp|symm|vnx|clariion|vplex|invista|hitachi|hpxp|ess|xtremio|generic}`
Limits the command to the specified type of storage system. all specifies all storage-system types. The default is all.

> **Note:** XtremIO devices of version 2.2 are managed under the `xtremio` class. Version 1.05 devices continue to be managed by `generic` storage class.

VNXe devices are included in the generic device class.

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

`every=<#seconds>`
Specifies the frequency with which port information is displayed. 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.

You can add a timestamp to the rpowermt 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, year, time in hours, minutes, and seconds that increment until you exit the rpowermt display command. TIME_SECONDS starts a counter rom zero that increments until you exit the rpowermt display command.

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

`width=<#cols>`
Sets the display width to a constant value within the range of 80 to 160 columns. This value applies only to the specific output; it is not persistent. You can set a persistent value for the width=#cols option with the PP_DISPLAY_WIDTH environment variable. This variable changes the width setting from 80 (the default) to the value you specify.

If the environment variable and command line values differ, the command line value takes precedence.

`host=<hostname>`
Specifies the IP address or hostname of the vSphere host.

`username=<username>`
Specifies the user account on the vSphere host. If the username is not included in the command line, you are prompted to specify the user name.

`password=<password>`
Specifies the password associated with the specified username. If the password is not included on the command line, you are prompted to specify the user name.

`cim_sessionid=<CIM sessionid>`
Specifies the CIM session ID ticket output for vSphere host authorization.

rpowermt help

Purpose
Displays a summary of the command syntax. If a specific command is specified, it displays the syntax for that command. The usage menu is also displayed when invalid commands or options are entered.

Syntax
`rpowermt help <command>`

Arguments
`<command>`

A valid rpowermt command.

Example
Following is a sample output when no command is specified.

```
Usage:
rpowermt <command>
[class={all|netapp|symm|vnx|clariion|vplex|invista|hitachi|hpxp|ess|xtremio|generic}]
host=<hostname>
```

```
[username=<username> [password=<password>|no_password]]
[cim_sessionid=<CIM sessionid>]
rpowermt check_registration
rpowermt display [every=<#seconds>] [width=<#col>]
rpowermt display alua dev=<device>|all [class=<class>|all] [every=<#seconds>]
[width=<#col>]
rpowermt display bus [class=<class>|all] [every=<#seconds>]
[width=<#col>]
rpowermt display latency [dev=<device>|all] [class=<class>|all]
       [every=<#seconds>] [width=<#col>]
rpowermt display options
rpowermt display perf dev=<device>|all [continuous] [verbose] [xml] [nowait]
rpowermt display perf bus [continuous] [verbose] [xml] [nowait]
rpowermt display ports [dev=<device>|all] [class=<class>|all]
       [every=<#seconds>] [width=<#col>]
rpowermt help [<command>]
rpowermt register [force]
rpowermt restore [hba=<hba#>|all] dev=<path>|<device>|all [class=<class>|all]
rpowermt set autostandby={on|off|reinitialize} [trigger={prox|iopf}]
rpowermt set autostandby iopflimit=<value>
rpowermt set autostandby agingperiod=<#days>
rpowermt set mode={active|standby} [hba=<hba#>|all]
       [dev=<path>|<device>|all] [class=<class>|all] [force]
rpowermt set path_latency_monitor={on|off} [force]
rpowermt set path_latency_threshold=<#seconds>|<#milliseconds>ms [force]
rpowermt set perfmon={on [interval=<#seconds> | off}
rpowermt set periodic_autorestore={on|off} [class=<class>|all] [force]
rpowermt set policy={ad|co|lb|li|rr|so|si} [dev=<device>|all]
[class=<class>|all] [force]
rpowermt set reactive_autorestore={on|off} [class=<class>|all] [force]
rpowermt set streamio_threshold=<threshold_count> [dev=<device>|all]
[class=<class>|all]
rpowermt setup add_host [host_file=<filename>] [username=<username>]
       [password=<password>]
rpowermt setup list_hosts [host_file=<filename>]
rpowermt setup remove_host [host_file=<filename>]
rpowermt setup update_host [host_file=<filename>]
       [username=<username>] [password=<password>]
rpowermt setup verify_hosts [host_file=<filename>]
rpowermt unregister [force]
rpowermt update lun_names
rpowermt version host=<hostname> username=<username>
       [password=<password>|no_password] [cim_sessionid=<CIM sessionid>]
```

Following is a sample output when a command (in this example, display) is specified.

```
C:\Documents and Settings\Administrator>rpowermt help display
display - Display multipathing information.
Usage:
  rpowermt <command>
[class={all|netapp|symm|vnx|clariion|vplex|invista|hitachi|hpxp|ess|xtremio|generic}]
           host=<hostname>
           [username=<username>] [[password=<password>|no_password]]
           [cim_sessionid=<cim sessionid>]
   rpowermt display
           [dev=<device|all> | class=<class|all>]
           [every=<#seconds>] [width=<#col>]
   rpowermt display latency
           [dev=<device|all> | class=<class|all>]
           [every=<#seconds>] [width=<#col>]
   rpowermt display options
   rpowermt display paths
           [class=<class|all>] [every=<#seconds>] [width=<#col>]
   rpowermt display ports
           [dev=<device|all> | class=<class|all>]
           [every=<#seconds>] [width=<#col>]
```

```
License search path: C:\Documents and Settings\Administrator\My
Documents\EMC\PowerPath\rpowermt;C:\
Documents and Settings\All Users\My Documents\EMC\PowerPath\rpowermt
Host file: C:\Documents and Settings\Administrator\My
Documents\EMC\PowerPath\rpowermt\lockbox.clb
```

rpowermt register

Purpose Registers a PowerPath/VE license for a vSphere host.

Syntax
```
rpowermt register [force] host=<hostname>
[username=<username> [password=<password>|no_password]]
[cim_sessionid=<CIM sessionid>]
```

Arguments `force`

> If a previous rpowermt register was not able to complete successfully, use the force argument to re-run the command.

`host=<hostname>`
> Specifies the IP address or hostname of the vSphere host.

`username=<username>`
> Specifies the user account on the vSphere host. If the username is not included in the command line, you are prompted to specify the user name.

`password=<password>`
> Specifies the password associated with the specified username. If the password is not included on the command line, you are prompted to specify the user name.

`cim_sessionid=<CIM sessionid>`
> Specifies the CIM session ID ticket output for vSphere host authorization.

Example For a vSphere host with a valid license, the following confirmation is displayed:

```
PowerPath license is registered.
```

If the vSphere host already has a valid license, the following warning is displayed:

```
Warning: PowerPath driver is already registered.
```

If the vSphere host does not have a valid license, the following error is displayed:

```
ERROR: License key or feature not found.
```

If the license file is corrupted, one of the following errors is displayed:

```
ERROR: License key or feature not found.
ERROR: License file format error or misspelling.
ERROR: License checkout failed.
```

rpowermt restore

Purpose Tests and restores specified paths. It issues test I/Os and responds to the test results as follow:

◆ If a *live* path passes the test, rpowermt restore does nothing.

◆ If a *dead* path passes the test, rpowermt restore marks it alive; PowerPath can now use it. In addition, other dead paths on the same HBA and/or storage system port may be tested.

◆ If a *live* path fails the test, rpowermt restore marks it dead and displays 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, rpowermt restore displays a warning (every time the path fails the test).

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

In addition to testing and restoring paths, rpowermt restore attempts to restore dead volumes. A volume 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 CLARiiON storage systems, rpowermt restore also relocates LUNs to their default storage processor.

Syntax

```
rpowermt restore [hba=<hba#>|all] [dev=<device>|all]
    [class=<class>|all] host=<hostname>
[username=<username> [password=<password>|no_password]]
[cim_sessionid=<CIM sessionid>]
```

Arguments

class={all|netapp|symm|vnx|clariion|vplex|invista|hitachi|hpxp|ess|xtremio|generic}
Limits the command to the specified type of storage system. all specifies all storage-system types. The default is all.

> **Note:** XtremIO devices of version 2.2 are managed under the `xtremio` class. Version 1.05 devices continue to be managed by `generic` storage class.

VNXe devices are included in the generic device class.

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

hba=<hba#>|all
Specifies the IP address or hostname of the vSphere host.

dev=<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.

host=<hostname>
Specifies the IP address or hostname of the vSphere host.

username=<username>
Specifies the user account on the vSphere host. If the username is not included in the command line, you are prompted to specify the user name.

password=<password>

Specifies the password associated with the specified username. If the password is not included on the command line, you are prompted to specify the user name.

```
cim_sessionid=<CIM sessionid>
```
Specifies the CIM session ID ticket output for vSphere host authorization.

Example   If there are new dead paths, an error message indicates each dead path:

```
Warning: Device emcpower0 bus 1 path C0:T0:L30 is currently dead.
```

rpowermt restore returns no output if there are no dead paths.

rpowermt server lockbox

Purpose   Uses a lockbox to securely store vSphere host IP addresses and their corresponding username/password. Storing this information eliminates the need for specifying a username/password for the vSphere host when running the rpowermt command.

Create lockbox

Any rpowermt command that includes the host argument attempts to retrieve the username/password for the specified vSphere host from the lockbox.

If the lockbox does not exist, rpowermt creates it, prompting you for the vSphere host's username/password and the new lockbox's passphrase:

```
Enter lockbox passphrase:
Confirm passphrase:
Enter server username: root
Enter server password:
```

The passphrase must meet the following requirements:

◆ Eight or more characters in length.

◆ Contain at least one numeric, one uppercase, one lowercase character, and one non-alphanumeric character (such as # or !).

You will need this passphrase if the lockbox is moved to a different rpowermt server.

If the lockbox exists but rpowermt cannot retrieve the username/password for the vSphere host (for example, no entry in the lockbox pertains to the specified host or the existing username/password stored for the specified vSphere host is no longer valid), rpowermt prompts you for the vSphere host credentials of the specified host:

```
Enter server username: root
Enter server password:
```

Default lockbox location

Linux — On Linux, the default lockbox is,

    : /etc/emc/lockbox/‹username›.clb

where ‹username› is the username of the rpowermt user.

Note: The username is obtained by the getpwuid_r() library call. Some systems may not support this library call. For those systems you can disable this internal call by setting the PPMT_DISABLE_PW_LOOKUP environment variable. The default lockbox on these system is /etc/emc/lockbox/lockbox.clb.

Windows — The default lockbox on Windows is:

Windows 2012:

C:\Users\Administrator\Documents\EMC\PowerPath\rpowermt\lockbox.clb

Windows 2008:

C:\Users\Administrator\ Documents\EMC\PowerPath\rpowermt\lockbox.clb

Windows 2003:

C:\Documents and Settings\‹username›\MyDocuments\EMC\PowerPath\rpowermt\lockbox.clb

Change lockbox name and location

Use the PPMT_LB_FILE environment variable to change the lockbox name and location.

Linux — To change the lockbox name to /tmp/pplockbox.clb on a Linux system, enter:

```
PPMT_LB_FILE=/tmp/pplockbox.clb; export PPMT_LB_FILE
```

Use the PPMT_LB_FILE environment variable to change the lockbox name and location.

Windows — To change the lockbox name to C:\temp\pplockbox.clb on a Windows system, enter:

```
set PPMT_LB_FILE=C:\temp\pplockbox.clb
```
To change the lockbox name to C:\temp\foobar.clb on a Windows system, type:

```
set PPMT_LB_FILE=C:\temp\foobar.clb
```

Disable lockbox

Use the PPMT_LB_DISABLE environment variable to disable writing to and reading from the lockbox.

Linux — To disable the lockbox on a Linux system, type:

```
export PPMT_LB_DISABLE=1
```
Windows — To disable the lockbox on a Windows system, enter:

```
set PPMT_LB_DISABLE=1
```
Alternatively, disable through the system settings.

1. Go to Control Panel › System › Advanced System Settings

2. Go to:

   a. Environment Variables and add to user variables.

   b. System Variables and add to systemwide.

After disabling the lockbox you are still prompted for username and password but these are not written to the file.

Bypass lockbox

Use the [use_lockbox=no] option to bypass the lockbox.

If both username=‹username› and password=‹password›|no password are entered as command line arguments, the lockbox is automatically bypassed.

Set rpowermt CST libraries environment variable

If the rpowermt CST libraries have been moved, you need to set an environment variable to point the rpowermt server to the correct location. This issue may occur if you have more than one EMC product installed on the rpowermt server and those other applications use different versions of the lockbox libraries.

The optional environment variable to point the rpowermt server to the correct directory is PP_LB_LIB.

On Windows, enter:

```
PP_LB_LIB=C:\program files\EMC\PowerPath\rpowermt
```

On Linux, enter:

```
export PP_LB_LIB=/usr/lib
```

On Linux hosts, if you receive an error message like the following:

```
ERROR: Failed to open host file.
```

The corrective action is to set the environment variable. Enter:

```
PP_LB_LIB=<location of CST libraries>
```
Alternatively, remove the environment variable definition for LD_LIBRARY_PATH. Enter:

```
unset LD_LIBRARY_PATH
```
This pertains to issue 339891.

rpowermt set autostandby

Purpose
Specifies a path that has been placed into standby by PowerPath/VE 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 multiframe 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/VE detects the most optimal paths (local paths) and places the remote paths in autostandby mode. When you turn on or re-initialize autostandby with the prox trigger, then PowerPath/VE 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 flaky paths. When the average number of I/O per path error falls between certain limits, the path changes to autostandby mode. Set the iopflimit

using the rpowermt 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 rpowermt set autostandby agingperiod =‹#days› command.

When you turn on autostandby with the iopf trigger, PowerPath/VE 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/VE places that path in autostandby. A higher limit makes the trigger more sensitive to errors. A lower limit makes the trigger more tolerant of errors.

This command turns the autostandby on or off globally within the PowerPath/VE. It does not pertain to a specific storage array or volume. The autostandby modes for each path do not persist across reboots. When the autostandby is turned on or off , the setting persists across reboots.

For example, if path set A is in autostandby mode and path B is in active mode, then on reboot path, path A can change to active mode and path 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 to put into autostandby mode and which path to put into active mode.

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

Note: For scenarios in which all paths to a volume are in asb:prox or asb:iopf mode, PowerPath/VE selects paths for I/O as if all paths are active. You can override autostandby for a path by running rpowermt set mode={standby|active} dev+‹path› force. If the path is already in the requested mode, then the command is ignored.

Syntax

```
rpowermt set autostandby={on|off|reinitialize}
    [trigger={prox|iopf} host=<hostname>
[username=<username> [password=<password>|no_password]]
[cim_sessionid=<CIM sessionid>]
```

Arguments

autostandby=on
    Enables autostandby. There is no output. If you do not select a trigger, then both prox and iopf are enabled. For example, rpowermt set autostandby=on 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/VE clears the global setting for enabling autostandby (or the selected type of autostandby) and persists that setting. PowerPath/VE 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, rpowermt set autostandby=off globally turns off autostandby with both proximity-based and the IOsPerFailure-based auto standby features.

autostandby=reinitialize
    By default, prox and iopf are enabled. If the proximity trigger is enabled, this option recalibrates the trigger and determines 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. When autostandby=reinitialize is run, prox and iopf are reinitialized.

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

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

`host=<hostname>`
Specifies the IP address or hostname of the vSphere host.

`username=<username>`
Specifies the user account on the vSphere host. If the username is not included in the command line, you are prompted to specify the user name.

`password=<password>`
Specifies the password associated with the specified username. If the password is not included on the command line, you are prompted to specify the user name.

`cim_sessionid=<CIM sessionid>`
Specifies the CIM session ID ticket output for vSphere host authorization.

rpowermt set autostandby agingperiod

Purpose
Sets the length of time that a path stays in autostandby due to an iopf (IOsPerFailure) trigger before it returns to active. When a path is placed in autostandby mode through the IOsPerFailure trigger, PowerPath/VE adds the aging amount to the current time. If the errors are corrected during this time, the aging period timeout allows PowerPath/VE to return the paths to full use without your intervention. Error monitoring is resumed when the paths are returned to active state. Each path ages separately. This may result in each path having a different time at which it returns to active mode.

When the aging period is changed, only paths moving to asb:iopf subsequent to the change will receive the new aging period. Those paths already in asb:iopf retain the aging period setting that was in force when they were moved to asb:iopf.

This command sets the aging period for autostandby iopf globally within PowerPath/VE and does not pertain to a specific storage array or volume. When you set aging period, this setting persists across reboots.

Syntax
```
rpowermt set autostandby agingperiod=<#days> host=<hostname>
    [username=<username> [password=<password>|no_password]]
    [cim_sessionid=<CIM sessionid>]
```

Arguments
`#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 period 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.

host=<hostname>
Specifies the IP address or hostname of the vSphere host.

username=<username>
Specifies the user account on the vSphere host. If the username is not included in the command line, you are prompted to specify the user name.

password=<password>
Specifies the password associated with the specified username. If the password is not included on the command line, you are prompted to specify the user name.

cim_sessionid=<CIM sessionid>
Specifies the CIM session ID ticket output for vSphere host authorization.

rpowermt set autostandby iopflimit

Purpose   Specifies the average number of I/Os between failures to automatically change that path to autostandby mode. This command sets autostandby iopflimit globally within PowerPath/VE and does not pertain to a specific storage array or volume.

Syntax
```
rpowermt set autostandby iopflimit=<value> host=<hostname>
[username=<username> [password=<password>|no_password]]
[cim_sessionid=<CIM sessionid>]
```

Arguments   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/VE 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/VE will not mark the path as having intermittent I/O errors. A higher limit makes the trigger more sensitive to errors. A lower limit makes the trigger more tolerant of errors.

host=<hostname>
Specifies the IP address or hostname of the vSphere host.

username=<username>
Specifies the user account on the vSphere host. If the username is not included in the command line, you are prompted to specify the user name.

password=<password>
Specifies the password associated with the specified username. If the password is not included on the command line, you are prompted to specify the user name.

cim_sessionid=<CIM sessionid>
Specifies the CIM session ID ticket output for vSphere host authorization.

rpowermt set mode

Purpose   Sets paths to active or standby mode.

For most applications, the best performance is achieved by designating all paths to a PowerPath/VE 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 can be overridden and moved to active or standby mode with this command. The force option is needed to override the 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.

```
ERROR: Path is in Autostandby, force option required
```

Syntax   rpowermt set mode={active|standby} [hba=<hba#>|all]
[dev=<path><device>|all] [class=<class>|all] [force] host=<hostname>
[username=<username> [password=<password>|no_password]]
[cim_sessionid=<CIM sessionid>]

The format of the native device path is, ‹hwpath›:‹Cx:Ty:Lz› value.

Arguments   mode=active|standby

Note: PowerPath/VE does not support the standby mode with iSCSI software initiators.

Determines whether a path is available for I/O or held in reserve. In active mode, 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. In standby mode, 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 prevent 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. Standby mode persists across reboots.

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

Limits the command to the specified type of storage system. all specifies all storage-system types. The default is all.

Note: XtremIO devices of version 2.2 are managed under the xtremio class. Version 1.05 devices continue to be managed by generic storage class.

VNXe devices are included in the generic device class.

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

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 paths output. all specifies all HBAs under PowerPath control. The default is all.

dev=<path><device>|all
>   Sets the mode for all paths to the specified path. all specifies all paths to all devices. The default is all.

>   The specified device cannot be a native device. The format of the native device path is ‹hwpath›:‹Cx:Ty:Lz› value.

force
>   If a previous rpowermt set mode command was not able to complete successfully, use the force argument to re-run the command.

host=<hostname>
>   Specifies the IP address or hostname of the vSphere host.

username=<username>
>   Specifies the user account on the vSphere host. If the username is not included in the command line, you are prompted to specify the user name.

password=<password>
>   Specifies the password associated with the specified username. If the password is not included on the command line, you are prompted to specify the user name.

cim_sessionid=<CIM sessionid>
>   Specifies the CIM session ID ticket output for vSphere host authorization.

rpowermt set path_latency_monitor

Purpose   Enables or disables path latency monitoring.

Before enabling path latency monitoring for the first time, read "rpowermt set path_latency_threshold" on page 51.

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

By sampling outputs from rpowermt 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.

Note: To avoid extraneous log messages, EMC recommends setting a threshold value before turning on latency monitoring.

Syntax   rpowermt set path_latency_monitor=on|off [force] host=<hostname>
[username=<username> [password=<password>|no_password]]
[cim_sessionid=<CIM sessionid>]

Arguments   on
>   Enables I/O latency monitoring on each path, resulting in meaningful outputs from rpowermt display latency, which are Current (most recent) and Max (high watermark) latencies for each path.

off
>   Disables I/O latency monitoring latency measurement for all paths and resets all paths Current and Max latencies to zero. The default is off.

force
    If a previous rpowermt set path_latency_monitor command was not able to complete successfully, use the force argument to re-run the command.

host=<hostname>
    Specifies the IP address or hostname of the vSphere host.

username=<username>
    Specifies the user account on the vSphere host. If the username is not included in the command line, you are prompted to specify the user name.

password=<password>
    Specifies the password associated with the specified username. If the password is not included on the command line, you are prompted to specify the user name.

cim_sessionid=<CIM sessionid>
    Specifies the CIM session ID ticket output for vSphere host authorization.

rpowermt set path_latency_threshold

Purpose   Sets a time interval in seconds within which I/Os should complete. When an I/O takes longer to complete than the set threshold value, PowerPath/VE sends a message to the log file. The threshold value applies to all paths in your environment.

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

For example, if the global path_latency_threshold is set to 2, the current Latency Max for the path (as seen in rpowermt 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, EMC recommends setting a threshold value before enabling path latency monitoring. Additionally, when changing a threshold value, disable path latency monitoring, set the new threshold value, and then re-enable path latency monitoring.

When changing the threshold, use the following 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 "rpowermt set path_latency_monitor" command. To view information on I/O completion times, use the "rpowermt display latency" command.

Example log message

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

```
<date> <time> <host_name> emcp: [ID 801593 kern.notice]
:Mpx:Latency Threshold exceeded by UserDev <bus_number>
Vol <volume number> with a new high latency of <latency_value> msec
2012-09-07T16:29:02.706Z cpu3:2051)PowerPath:Latency Threshold exceeded by Path
```

```
vmhba2:C0:T3:L1 Vol Symm 000192601699 vol 1ff5 with a new high latency of 13 msec
```

Syntax
```
rpowermt set path_latency_threshold=<#seconds|<#milliseconds>ms
    [force] host=<hostname>[username=<username>
    [password=<password>|no_password]][cim_sessionid=<CIM sessionid>]
```

Arguments   `<#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, or 0 to 3600000 milliseconds. Seconds is the default.

When the threshold is set to zero (also known as Discovery Mode), PowerPath/VE 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, ensure that you set an exceptionally high threshold (for example, 30 seconds) before enabling latency monitoring for the first time. You can then sample **rpowermt display latency** output to estimate expected latencies for your environment

`force`
If a previous rpowermt set path_latency_threshold command was not able to complete successfully, use the force argument to re-run the command.

`host=<hostname>`
Specifies the IP address or hostname of the vSphere host.

`username=<username>`
Specifies the user account on the vSphere host. If the username is not included in the command line, you are prompted to specify the user name.

`password=<password>`
Specifies the password associated with the specified username. If the password is not included on the command line, you are prompted to specify the user name.

`cim_sessionid=<CIM sessionid>`
Specifies the CIM session ID ticket output for vSphere host authorization.

rpowermt set periodic_autorestore

Purpose   Enables or disables the PowerPath/VE periodic autorestore facility.

When periodic autorestore is on, PowerPath/VE 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, and automatic path restoration continues to occur based on the results of that testing.

To determine the periodic autorestore setting, use rpowermt display options.

periodic_autorestore can be used with reactive_autorestore to define when paths are restored after path failure.

Syntax
```
rpowermt set periodic_autorestore={on|off} [class=<class>|all] [force]
    host=<hostname>
```

```
                        [username=<username> [password=<password>|no_password]]
                        [cim_sessionid=<CIM sessionid>]
```

Arguments   on
                    Enables periodic autorestore. The default is on.

                off
                    Disables periodic autorestore.

class={all|netapp|symm|vnx|clariion|vplex|invista|hitachi|hpxp|ess|xtremio|generic}
                    Limits the command to the specified type of storage system. all specifies all
                    storage-system types. The default is all.

> **Note:** XtremIO devices of version 2.2 are managed under the `xtremio` class. Version
> 1.05 devices continue to be managed by `generic` storage class.

                VNXe devices are included in the generic device class.

                force
                    If a previous rpowermt set periodic_autorestore command was not able to complete
                    successfully, use the force argument to re-run the command.

                host=<hostname>
                    Specifies the IP address or hostname of the vSphere host.

                username=<username>
                    Specifies the user account on the vSphere host. If the username is not included in the
                    command line, you are prompted to specify it.

                password=<password>
                    Specifies the password associated with the specified username. If the password is not
                    included on the command line, you are prompted to specify it.

                cim_sessionid=<CIM sessionid>
                    Specifies the CIM session ID ticket output for vSphere host authorization.

rpowermt set perfmon

Purpose   Enables or disables data collection for performance monitoring of all devices. Performance
              monitoring helps characterize I/O patterns and possibly aids in diagnosing I/O problems.

              When performance monitoring is enabled, rpowermt display dev=all and rpowermt
              display perf bus displays performance metrics for all paths to all devices and bus ports.
              See rpowermt display perf and "rpowermt display perf bus" for more information. To view
              whether performance display is enabled, use the "rpowermt display options" command.

Syntax   rpowermt set perfmon={on [interval=<#seconds>] | off} host=<hostname>
              [username=<username> [password=<password>|no_password]]
              [cim_sessionid=<CIM sessionid>]

Arguments   on
                    Enables performance monitoring for all devices. There is no output. This command
                    initializes the counters, clears all measurements including high and low watermarks
                    for latency, and accepts a sampling interval for data collection in seconds. The default
                    is 900 seconds (15 minutes).

Note: The legacy path latency monitoring threshold watermark is not cleared.

off

Disables performance monitoring and clears all measurements. To clear all measurements and restart performance monitoring, you must first set perfmon=off and then set perfmon=on. There is no output.

interval=<#seconds>
The sampling interval in seconds for data collection. The default is 900 seconds (15 minutes).

host=<hostname>
Specifies the IP address or hostname of the vSphere host.

username=<username>
Specifies the user account on the vSphere host. If the username is not included in the command line, you are prompted to specify the user name.

password=<password>
Specifies the password associated with the specified username. If the password is not included on the command line, you are prompted to specify the user name.

cim_sessionid=<CIM sessionid>
Specifies the CIM session ID ticket output for vSphere host authorization.

rpowermt set policy

Purpose   Sets the load-balancing and failover policy for devices managed by PowerPath.

Note: When PowerPath/VE is installed on the VMware vSphere host, the default load-balancing and failover policies are so (Symmetrix optimized) for Symmetrix devices, co (CLARiiON optimized) for VNX and CLARiiON devices, and ad (Adaptive) for Invista, VPLEX, VNXe, supported Celerra devices, and supported third-party devices. When PowerPath/VE is installed but not licensed, the rpowermt server is unable to display the information for the vSphere host. Upon installation of a valid PowerPath/VE for VMware vSphere license, host display and management capabilities are enabled through the rpowermt server. bf and nr are not supported as an argument to the rpowermt set policy command. You cannot manually set the policy for a device to **bf** or **nr.**

Syntax
```
rpowermt set policy={ad|co|lb|li|rr|so|si} [dev=<device>|all]
[class=<class>|all] [force] host=<hostname>
[username=<username> [password=<password>|no_password]]
[cim_sessionid=<CIM sessionid>]
```

Arguments   policy=ad|co|lb|li|rr|so|si

Table 9 describes the policies.

Table 9 Policies

Policy	Description
ad (Adaptive)	I/O requests are assigned to paths based on an algorithm that takes into account path load and logical device priority. This is the default for EMC VPLEX and Invista and supported third-party arrays on a vSphere host with a registered PowerPath license.
co (CLARiiON optimization)	I/O requests are assigned to paths based on an algorithm that takes into account path load and logical device priority. This policy is valid only for VNX and CLARiiON storage systems and is the default policy for CLARiiON storage systems on a vSphere host with a registered PowerPath license.
lb (Least blocks)	Load balance is based on the number of blocks in pending I/Os. I/O requests are routed to the path with the least queued blocks, regardless of the number of requests involved.
li (Least I/O)	Load balance is based on the number of pending I/Os. I/O requests are routed to the path with the least queued requests, regardless of total block volume.
rr (Round robin)	I/O requests are assigned to each available path in rotation.
si (StreamIO)	For each possible path for an I/O to a particular volume, this policy selects the same path as was selected for the previous I/O to the volume, unless the pending I/O count since the last path change exceeds the volume's threshold value. When the threshold is exceeded, the policy selects a new path based on the adaptive policy algorithm.
so (Symmetrix optimization)	I/O requests are routed to paths based on an algorithm that takes into account path load and logical device priority. Load is a function of the number, size, priority, and type of I/O queued on each path. This policy is valid only for Symmetrix storage systems and is the default policy for Symmetrix storage systems on a vSphere host with a registered PowerPath license.

`dev=<device>|all`

Limits the policy change to the specified devices.

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

> **Note:** Limits the command to the specified type of storage system. all specifies all storage-system types. The default is all. XtremIO devices of version 2.2 are managed under the `xtremio` class. Version 1.05 devices continue to be managed by `generic` storage class.

VNXe devices are included in the generic device class.

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

`force`

If a previous rpowermt set policy command was not able to complete successfully, use the force argument to re-run the command.

`host=<hostname>`
Specifies the IP address or hostname of the vSphere host.

`username=<username>`

Specifies the user account on the vSphere host. If the username is not included in the command line, you are prompted to specify it.

`password=<password>`
Specifies the password associated with the specified username. If the password is not included on the command line, you are prompted to specify it.

`cim_sessionid=<CIM sessionid>`
Specifies the CIM session ID ticket output for vSphere host authorization.

rpowermt set reactive_autorestore

Purpose   Enables or disables the PowerPath/VE reactive autorestore facility.

Reactive autorestore 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_autorestore is on, PowerPath/VE reactively tests dead paths and, if they pass the test, restores them to service. For non-disruptive upgrades (NDU), rpowermt set reactive_autorestore should be set to on.

To determine the reactive_autorestore setting, use "rpowermt display options".

reactive_autorestore can be used with periodic_autorestore to define when paths are restored after path failure.

- ◆ If reactive_autorestore=on and periodic_autorestore=on, then paths will be automatically restored at the next opportunity in either context. This is the default for these settings.

- ◆ If reactive_autorestore=on and periodic_autorestore=off, then 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. 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.

- ◆ If reactive_autorestore=off and periodic_autorestore=on, then the path will be restored when the path is selected for path test.

- ◆ If reactive_autorestore=off and periodic_autorestore=off, then paths that move into the dead state due to path test failures will never be automatically restored. EMC does not recommend this setting in your normal running environment. 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 autorestored and returned to the alive state during the next path test. This path test may be either periodic or in the context of completing I/O (reactive). To restore physically dead paths in this scenario, run rpowermt restore at the command line.

Syntax   `rpowermt set reactive_autorestore={on|off} [class=<class>|all] [force]`
`    host=<hostname>`
`[username=<username> [password=<password>|no_password]]`
`[cim_sessionid=<CIM sessionid>]`

Arguments   `on`

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

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 monitoring if you place a higher priority on these failover mechanisms than you place on PowerPath/VE failover

class={all|netapp|symm|vnx|clariion|vplex|invista|hitachi|hpxp|ess|xtremio|generic}
Limits the command to the specified type of storage system. all specifies all storage-system types. The default is all.

Note: XtremIO devices of version 2.2 are managed under the xtremio class. Version 1.05 devices continue to be managed by generic storage class.

VNXe devices are included in the generic device class.

force
If a previous rpowermt set periodic_autorestore command was not able to complete successfully, use the force argument to re-run the command.

host=<hostname>
Specifies the IP address or hostname of the vSphere host.

username=<username>
Specifies the user account on the vSphere host. If the username is not included in the command line, you are prompted to specify it.

password=<password>
Specifies the password associated with the specified username. If the password is not included on the command line, you are prompted to specify it.

cim_sessionid=<CIM sessionid>
Specifies the CIM session ID ticket output for vSphere host authorization.

rpowermt set streamio_threshold

Purpose   Sets the I/O threshold values to switch paths on devices where the StreamIO (si) load-balancing policy is in effect on PowerPath/VE.

The streamio value is accepted only if the load-balancing policy in effect on the device is si. If you specify a device for which the load-balancing policy is not si, this will generate an error message.

You can check whether streamio threshold is set by viewing the rpowermt display output. For example: state=alive; policy=StreamIo; streamio_threshold=1000; queued-IOs=o.

Syntax
```
rpowermt set streamio_threshold=<threshold_count>
    [dev=device|all] [class=<class>|all] host=<hostname>
    [username=<username>
    [password=<password>|no_password]]
    [cim_sessionid=<CIM sessionid>]
```

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.

These values are truncated to nearest power of 2, which is lower than the specified value. For example, if you try to set the streamio_threshold value as 127, the truncated value 64 is set as the streamio_threshold value.

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

`class={all|netapp|symm|vnx|clariion|vplex|invista|hitachi|hpxp|ess|xtremio|generic}`
Limits the command to the specified type of storage system. all specifies all storage-system types. The default is all.

> **Note:** XtremIO devices of version 2.2 are managed under the `xtremio` class. Version 1.05 devices continue to be managed by `generic` storage class.

VNXe devices are included in the generic device class.

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

`host=<hostname>`
Specifies the IP address or hostname of the vSphere host.

`username=<username>`
Specifies the user account on the vSphere host. If the username is not included in the command line, you are prompted to specify it.

`password=<password>`
Specifies the password associated with the specified username. If the password is not included on the command line, you are prompted to specify it.

`cim_sessionid=<CIM sessionid>`
Specifies the CIM session ID ticket output for vSphere host authorization.

rpowermt setup add_host

Purpose Adds a vSphere host and its corresponding username/password pair to the specified lockbox (if the host_file argument is specified) or to the default lockbox on the rpowermt server.

Syntax `rpowermt setup add_host [host_file=<filename>] host=<hostname>`
`[username=<username> [password=<password>]]`

Arguments `host_file=<filename>`
Lockbox file on the rpowermt server. If the specified lockbox does not exist, it is created.

If the host_file argument is omitted, the vSphere host and its username/password pair are added to the default lockbox.

`host=<hostname>`
Specifies the IP address or hostname of the vSphere host.

`username=<username>`

User account on the vSphere host. If the username is not included in the command line, you are prompted to specify it.

```
password=<password>
```
Password associated with the specified username. If the password is not included on the command line, you are prompted to specify it.

Example   If the specified lockbox does not exist, you are prompted for the vSphere host password (if not specified on command line) and new lockbox passphrase. The lockbox is created.

```
Enter lockbox passphrase:
Confirm passphrase:
Enter server password:
```

If the specified lockbox exists, but there is no host entry created for the specified host, you are prompted for the vSphere host's password (if not specified on command line). The username/password entry is added to the lockbox.

```
# rpowermt setup add_host host=1.1.1.1 username=root
Enter server password:
Confirm password:
```

If the specified lockbox exists and it has a host entry created for specified host, you are prompted for host's password (if not specified on command line) . The username/password entry is updated.

```
# rpowermt setup add_host host=1.1.1.1 username=root
Enter server password:
```

rpowermt setup list_hosts

Purpose   Lists each vSphere host that has a username/password entry in the specified lockbox (if the host_file argument is specified) or in the default lockbox on the rpowermt server.

Syntax   `rpowermt setup list_hosts [host_file=<filename>] [host=<hostname>]`

Arguments   `host_file=<filename>`
List the vSphere host in the specified lockbox on the rpowermt server. If this argument is omitted, the vSphere hosts in the default lockbox are listed.

`host=<hostname>`
Specifies the IP address or hostname of the vSphere host.

Example   If the specified lockbox exists and it has stored entries, a list of host IPs are displayed:

```
# rpowermt setup list_hosts
10.31.18.34 <host1.domainname.com>
111.111.111.111 <host2.domainname.com>
111.111.111.112 <host3.domainname.com>
```

rpowermt setup remove_host

Purpose   Removes a vSphere host and its username/password entry from the specified lockbox (if the host_file argument is specified) or from the default lockbox on the rpowermt server.

Syntax   `rpowermt setup remove_host [host_file=<filename>] host=<hostname>`

Arguments host_file=<filename>
Remove the vSphere host from the specified lockbox on the rpowermt server. If this argument is omitted, the vSphere host is removed from the default lockbox.

host=<hostname>
Specifies the IP address or hostname of the vSphere host.

Example If the specified lockbox exists and it has an entry for specified vSphere host, no output is returned.

rpowermt setup update_host

Purpose Updates the specified vSphere host's username/password entry in the specified lockbox (if the host_file argument is specified) or in the default lockbox on the rpowermt server.

Syntax `rpowermt setup update_host [host_file=<filename>] host=<hostname> [username=<username>] [password=<password>]`

Arguments host_file=<filename>
Lockbox file on the rpowermt server. If the specified lockbox does not exist, it is created. If the host_file argument is omitted, the vSphere host and its username/password pair are updated in the default lockbox.

host=<hostname>
Specifies the IP address or hostname of the vSphere host.

username=<username>
User account on the vSphere host. If the username is not included in the command line, you are prompted to specify it.

password=<password>
Password associated with the specified username. If the password is not included on the command line, you are prompted to specify it.

Example If the specified lockbox exists and it has an entry for specified host exists, you are prompted for vSphere host's password (if not specified on command line). The host entry is updated.

```
# rpowermt setup update_host host=1.1.1.1 username=root
Enter server password:
```

If the specified lockbox exists but it does not have an entry for specified vSphere host, you are prompted for the vSphere host's password (if not specified on command line). The host entry is added.

```
# rpowermt setup update_host host=1.1.1.1 username=root
Enter server password:
```

rpowermt setup verify_hosts

Purpose Lists the hosts/IPs from the host_file (that is, the rpowermt lockbox) that are reachable. The lockbox name/location can be the default, can be specified with the PPMT_LB_FILE environment variable, and can be specified through the command line (if the host_file argument is specified).

Syntax   `rpowermt setup verify_host [host_file=<filename>] [host=<hostname>]`

Arguments   `host_file=<filename>`
   Lockbox file on the rpowermt server. If the specified lockbox does not exist, it is
   created. If the host_file argument is omitted, the vSphere host and its
   username/password pair are updated in the default lockbox.

`host=<hostname>`
   Specifies the IP address or hostname of the vSphere host.

Example   If the specified lockbox exists and the verification is successful, the following output is
displayed:

```
# rpowermt setup verify_hosts host=lcla112
Host verified
111.111.111.112
```

If the specified lockbox exists and the verification is not successful, the following output
is displayed:

```
# rpowermt setup verify_hosts
Host(s) Verified
111.111.111.112
111.111.111.113                  Failed
```

If the specified lockbox exists, and stored entries are present, the output lists all host IPs
that are reachable:

rpowermt setup verify_hosts

```
IP address Hostname
111.111.111.112       <host1.domainname.com>
111.111.111.113       <host2.domainname.com>
10.12.13.14
```

If the specified lockbox exists but no entries are currently stored, the output shows that no
entries are found:

```
# rpowermt setup verify_hosts
No entries are found
```

If the specified lockbox does not exist the output displays an error message:

```
# rpowermt setup verify_hosts
ERROR: Failed to open host file.
```

If the specified lockbox exists and specified host is present, the output displays a list of
the host IP addresses stored if it's reachable:

```
# rpowermt setup verify_hosts host=111.111.111.112
Host(s) verified
111.111.111.112
```

If the specified lockbox exists and specified host is not present, an error message is
displayed:

```
# rpowermt setup verify_hosts host=111.111.111.112
```

```
                              ERROR: Host not found.
```

rpowermt unregister

Purpose   Removes a PowerPath license for a vSphere host.

Syntax
```
rpowermt unregister [force] host=<hostname>
[username=<username> [password=<password>|no_password]]
[cim_sessionid=<CIM sessionid>]
```

Arguments
`force`
   If a previous rpowermt unregister was not able to complete successfully, use the force argument to re-run the command.

`host=<hostname>`
   Specifies the IP address or hostname of the vSphere host.

`username=<username>`
   Specifies the user account on the vSphere host. If the username is not included in the command line, you are prompted to specify it.

`password=<password>`
   Specifies the password associated with the specified username. If the password is not included on the command line, you are prompted to specify it.

`cim_sessionid=<CIM sessionid>`
   Specifies the CIM session ID ticket output for vSphere host authorization.

Example   For a vSphere host with a current valid license:

```
Do you really want to unregister the PowerPath license? yes/[no]: yes
PowerPath license is unregistered.
```

For a vSphere host that does not have a valid license:

```
Warning: PowerPath driver is already unlicensed.
```

rpowermt update lun_names

Purpose   Retrieves the latest CLARiiON nice names. A CLARiiON nice name, also called a user-assignable LUN name, is a character string associated with a logical device on a VNX and CLARiiON array and assigned through Unisphere and Navisphere. Run this command to perform initial discovery of CLARiiON nice names and whenever the nice name is changed on the VNX and CLARiiON array through Unisphere and Navisphere.

Note: CLARiiON LUN nice names are updated automatically every hour. Use this command to manually update CLARiiON nice names.

Syntax
```
rpowermt update lun_names host=<hostname>
[username=<username> [password=<password>|no_password]]
[cim_sessionid=<CIM sessionid>]
```

Arguments
`host=<hostname>`
   Specifies the IP address or hostname of the vSphere host.

username=<username>
>   Specifies the user account on the vSphere host. If the username is not included in the command line, you are prompted to specify it.

password=<password>
>   Specifies the password associated with the specified username. If the password is not included on the command line, you are prompted to specify it.

cim_sessionid=<CIM sessionid>
>   Specifies the CIM session ID ticket output for vSphere host authorization.

Example   CLARiiON nice names are disabled in the following rpowermt display dev output:

```
Pseudo name=emcpower2
CLARiiON ID=HK190807490007 []
Standard UID=naa.60060160bab01e00930a8cdde601dd11 []
```

CLARiiON nice names are enabled in the following rpowermt display dev output:

```
Pseudo name=emcpower2
CLARiiON ID=HK190807490007 [storage group name]
Standard UID=naa.60060160bab01e00930a8cdde601dd11 [LUN 4]
```

Note the storage group name and the device name (LUN 4) are now displayed.

rpowermt version

Purpose   Displays the installed version of PowerPath, the license search path, the host file name and location on the vSphere host.

Syntax
```
rpowermt version host=<hostname> [username=<username>]
[[password=<password>|no_password]][cim_sessionid=<CIM sessionid>]
```

Arguments   host=<hostname>
>   Specifies the IP address or hostname of the vSphere host.

username=<username>
>   Specifies the user account on the vSphere host. If the username is not included in the command line, you are prompted to specify it.

password=<password>
>   Specifies the password associated with the specified username. If the password is not included on the command line, you are prompted to specify it.

cim_sessionid=<CIM sessionid>
>   Specifies the CIM session ID ticket output for vSphere host authorization.

Example   Output similar to the following appears:

```
c:\> rpowermt version host=<host name>
EMC rpowermt for PowerPath (c) client Version 5.7 (build 154)
EMC PowerPath (c) host=<host name>Version 5.7 (build 154)
License search path: C:\Users\<user>\Documents\EMC\PowerPath\rpowermt; C:\ProgramData\My
Documents\EMC\PowerPath\rpowermt
Host file: C:\Users\<user>\Documents\EMC\PowerPath\rpowermt\lockbox.clb
```

CHAPTER 3
rpowermt error and warning messages

The chapter describes the error messages returned by the rpowermt command.

Error messages

Table 10 lists the error messages

Table 10 Error messages (page 1 of 3)

Message	Error
ERROR: Performance Monitor is not enabled.	Running any perf command when the performance monitor is off.
ERROR: License autoregistration is disabled. Commands not allowed.	PPMT_DISABLE_LICAUTO is set to 1. Disabling autoregistration on an unlicensed host.
ERROR: ‹array_type› device(s) not found	The rpowermt set policy command specified a policy that is valid only for one class but no devices are configured for that class.
ERROR: Cannot communicate with host, version mismatch too great.	The rpowermt server from which you ran rpowermt check_registration is running a version of RTOOLS that is not compatible with the PowerPath/VE version on the host.
ERROR: Enable Autostandby first.	An attempt has been made to reinitialize the Autostandby setting while the global Autostandby setting is off.
ERROR: multiple commands specified: ‹command 1› and ‹command 2›.	Multiple commands were included on a single command line.
ERROR: Invalid option [‹invalid_command/option›] - cannot resolve	An invalid command or argument was specified.
ERROR: "every" value ‹invalid_value› is not in range (1 - 86400).	An invalid interval was specified. For values greater than 2^31, the overflowed value is indicated in the error message.
ERROR: Unsupported class name: ‹invalid_class›. class=‹all|netapp|symm|vnx|clariion|vplex|invista| netapp|hitachi|hpxp|ess|xtremio|generic›	An invalid class argument was specified.
ERROR: Bad dev value emcpower ‹invalid value›, or not under PowerPath control.	An invalid device number was specified.
ERROR: Bad dev value ‹invalid_value›, or not under PowerPath control.	A non-alpha character was specified.
ERROR: "width" value ‹invalid_value› is not in range (80 - 160).	An invalid width was specified. For values greater than 2^31, the overflowed value is indicated in the error message.
ERROR: NULL value specified with ‹option›=	Either the every, class, dev, width, or hba argument was left blank.
ERROR: Device(s) not found.	A management command (for example, a set command) was issued, but no devices are being managed by PowerPath.
ERROR: Bad hba value ‹invalid_value›, or not under PowerPath control.	An invalid HBA number was specified. For values greater than 2^31, the overflowed value is indicated in the error message.
ERROR: ‹class› device(s) not found.	Both the class and dev arguments were used, but the specified values do not match a currently configured PowerPath volume and device path.
ERROR: ‹class› device(s) not found.	A management command (for example, a set command) was issued for a vSphere host, but no devices were configured for the specified class.
ERROR: License not installed.	A management command (for example, a set command) was issued for a vSphere host that does not have a valid license registered for it.

Table 10 Error messages (page 2 of 3)

Message	Error
ERROR: Path is in Autostandby, force option required.	The path is in autostandby and the force option was not specified with the rpowermt set mode command line to change autostandby to active or standby mode.
ERROR: PowerPath Not Found	PowerPath is not installed on the specified host.
ERROR: User Authentication Invalid.	A rpowermt command (other than help or setup) was run, but the specified lockbox did not exist. rpowermt attempted to create a lockbox, but the supplied username and/or password was invalid.
ERROR: Host not found.	A rpowermt command (other than help or setup) was run, but the specified lockbox did not exist. rpowermt attempted to create a lockbox, but the specified host could not be reached.
ERROR: Invalid passphrase during host file creation.	The passphrase entered for a lockbox did not meet the passphrase requirements (that is, a minimum of 8 characters and must contain a lowercase, uppercase, numeric, and special character). Refer to "rpowermt server lockbox" on page 43 for more information.
ERROR: Failed to open host file.	The "rpowermt setup verify_hosts" on page 60 command verified that a lockbox exists but specified host is not present. The "rpowermt setup list_hosts" on page 59 command could not obtain a lock on the specified lockbox. Possible causes include: • User does not have the necessary privileges/permissions to access the lockbox • User does not have the necessary privileges/permissions to create a lockbox at the specified location • Another process is holding the lock • Lockbox does not exist
ERROR: License authorization failed.	The rpowermt register command is located the license file, but the file did not contain a key for the specified vSphere host.
ERROR: License file does not exist or not found.	The rpowermt register command could not locate the license file.
ERROR: License key or feature not found.	The rpowermt register command is located the license file, but the feature tag within the license file is corrupted.
ERROR: License file format error or misspelling.	The rpowermt register command is located the license file, but the license file is corrupted.
ERROR: License checkout failed.	The rpowermt register command is located the license file, but the attribute value within the license file is corrupted.
ERROR: Path latency threshold value ‹invalid value› is not an integer in the range of (0 - 3600) or (0 - 3600000ms).	An invalid path latency threshold value was specified. Refer to rpowermt set_path_latency_threshold for more information.
ERROR: Passphrase requires at least eight characters and must contain a lowercase, uppercase, numeric, and special character.	The lockbox phrase entered is not complex enough. Enter a different passphrase with at least eight characters. The passphrase must contain an uppercase and lowercase character, a numeric character, and a special character.
ERROR: Failed to find IP address for ‹unresolved hostname›.	An rpowermt command was not able to resolve the IP for the hostname supplied with the host argument.

Table 10 Error messages (page 3 of 3)

Message	Error
ERROR: Failed to store host information into host file.	An encryption-specific error occurred. Such errors, while extremely rare, may prevent rpowermt from properly managing the lockbox.
ERROR: Failed to retrieve host information from host file.	
ERROR: Failed to remove host information from host file.	
ERROR: Failed to store host list into host file.	
ERROR: Failed to retrieve host list from host file.	

Warning messages

Table 11 describes the warning messages returned by the rpowermt command.

Table 11 Warning messages

Message	Description
WARNING: Failed to reset the license server count.	License server is not responding or unreachable. It appears when you try to unregister served license and the license server is unreachable. **Workaround**: Repeat the command with the force option after communication to the license server is restored.
WARNING: PowerPath driver is already registered.	The rpowermt register command was issued for a vSphere host that already has a PowerPath/VE license registered for it.
WARNING: Device(s) not found.	A non-management command (for example, display) was issued, but no devices are being managed by PowerPath.
WARNING: License will expire in ‹number_of_days› days.	A registered PowerPath license is set to expire within the specified number of days.
WARNING: License not installed.	A non-management command (for example, display) was issued for a vSphere host that does not have a valid license registered for it.
WARNING: ‹class› device(s) not found.	A non-management command (for example, set) was issued for a vSphere host, but no devices were configured for the specified class.
Warning: Device ‹device_bus_path› is currently dead.	New dead paths were found. Refer to rpowermt restore for more information.

rpowermt error and warning messages