

Dell EMC PowerPath Family

CLI and System Messages Reference

7.x

Notes, cautions, and warnings

 **NOTE:** A NOTE indicates important information that helps you make better use of your product.

 **CAUTION:** A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

 **WARNING:** A WARNING indicates a potential for property damage, personal injury, or death.

Revision history

The following table presents the revision history of this document:

Table 1. Revision history

Revision	Date	Description
4.0	July 2020	Release of PowerPath for Windows 7.0.
3.0	May 2020	Release of PowerPath/VE for VMware vSphere 7.0 and updated with PowerStore storage systems support details.
2.0	January 2020	Release of PowerPath for Linux 7.1.
1.0	July 2019	Release of PowerPath for Linux 7.0.

Revision history.....	3
Tables.....	16
Chapter 1: PowerPath Command Line Interface Overview.....	17
Welcome to the PowerPath Command Line Interface Reference.....	17
Installing the PowerPath CLI.....	17
PowerPath CLI install directories.....	17
PowerPath messages locations.....	18
PowerPath information, warning, and error messages.....	18
List of the supported commands by platform.....	18
Saving and loading a PowerPath configuration.....	19
Automatic license registration on PowerPath/VE for VMware vSphere.....	19
Class option specification.....	19
Specifying the wide or width option.....	20
Chapter 2: emcpadm commands.....	22
emcpadm command.....	22
emcpadm check_mappings.....	22
emcpadm export_mappings.....	23
emcpadm getfreepseudos.....	23
emcpadm getusedpseudos.....	24
emcpadm help.....	24
emcpadm import_mappings.....	24
emcpadm print_mappings.....	25
emcpadm renamepseudo.....	25
emcpadm version.....	26
Chapter 3: emcpbfs_linux commands.....	27
emcpbfs_linux command.....	27
emcpbfs_linux config.....	27
emcpbfs_linux remove.....	27
Chapter 4: emcpcfg script commands.....	28
emcpcfg script.....	28
emcpcfg backup.....	28
emcpcfg cleanup.....	28
emcpcfg list cfgs.....	29
emcpcfg restore.....	29
emcpcfg set cfgs limit.....	29
Chapter 5: emcphostid commands.....	30
emcphostid command.....	30
The emcphostid command on Windows hosts.....	30

Chapter 6: emcplun commands.....	32
emcplun_linux utility.....	32
emcplun_linux help.....	32
emcplun_linux list.....	32
emcplun_linux remove.....	32
emcplun_linux scan.....	33
Chapter 7: emcpreg commands.....	35
emcpreg utility.....	35
The emcpreg command on UNIX and Linux.....	35
The emcpreg command on Windows.....	37
Chapter 8: powercf commands.....	39
powercf command.....	39
powercf.....	39
Chapter 9: powerkmd commands.....	40
powerkmd command.....	40
powerkmd display.....	40
powerkmd help.....	40
powerkmd remove.....	41
powerkmd version.....	41
Chapter 10: powermig commands.....	42
powermig command.....	42
powermig abort.....	42
powermig cleanup.....	43
powermig commit.....	44
powermig getHandle.....	44
powermig help.....	45
powermig info.....	45
powermig options.....	46
powermig pause.....	47
powermig query.....	48
powermig recover.....	49
powermig resume.....	50
powermig selectSource.....	50
powermig selectTarget.....	51
powermig setup.....	52
powermig sync.....	54
powermig throttle.....	54
powermig undoRedirect.....	55
powermig version.....	56
Chapter 11: powermt commands.....	57
powermt command.....	57
powermt check.....	57
powermt check_registration.....	58

powermt config.....	59
powermt disable hba.....	60
powermt display.....	61
powermt display latency.....	64
powermt display alua.....	65
powermt display array_serial.....	66
powermt display bus.....	67
powermt display bus-ports san-latency.....	68
powermt display dead.....	69
powermt display hba_mode.....	69
powermt display nonvirtual.....	70
powermt display options.....	71
powermt display perf bus.....	74
powermt display perf dev.....	77
powermt display port_mode.....	81
powermt display ports.....	82
powermt display unmanaged.....	84
powermt enable hba.....	84
powermt help.....	85
powermt load.....	86
powermt manage.....	86
powermt release.....	87
powermt remove.....	87
powermt restore.....	87
powermt save.....	88
powermt set all_paths_dead_retry_timeout.....	89
powermt set auto_host_registration.....	90
powermt set autostandby.....	90
powermt set autostandby agingperiod.....	92
powermt set autostandby iopflimit.....	93
powermt set device_perf_to_array_report.....	93
powermt set dev_inuse_to_array_report.....	93
powermt set mode.....	94
powermt set path_latency_monitor.....	95
powermt set path_latency_threshold.....	95
powermt set path_retry_timeout.....	96
powermt set perfmon.....	96
powermt set periodic_autorestore.....	97
powermt set persist_path_info.....	97
powermt set policy.....	98
powermt set port_disable.....	98
powermt set port_mode.....	99
powermt set reactive_autorestore.....	100
powermt set streamio_threshold.....	101
powermt set write_throttle.....	101
powermt set write_throttle_queue.....	102
powermt unmanage.....	102
powermt update host_registration.....	103
powermt update lun_names.....	103
powermt version.....	104

Chapter 12: pprootdev commands.....	105
pprootdev command.....	105
pprootdev.....	105
Chapter 13: rpowermt commands.....	106
rpowermt command.....	106
rpowermt check_registration.....	106
rpowermt check_served.....	107
rpowermt display.....	108
rpowermt display alua.....	109
rpowermt display array serial.....	110
rpowermt display bus.....	111
rpowermt display latency.....	112
rpowermt display options.....	113
rpowermt display paths.....	116
rpowermt display perf bus.....	117
rpowermt display perf bus vmname.....	119
rpowermt display perf dev.....	120
rpowermt display perf vmname.....	122
rpowermt display port_mode.....	124
rpowermt display ports.....	125
rpowermt help.....	127
rpowermt register.....	128
rpowermt restore.....	128
rpowermt set autostandby.....	129
rpowermt set autostandby iopflimit.....	131
rpowermt set autostandby agingperiod.....	131
rpowermt set auto_host_registration.....	132
rpowermt set mode.....	134
rpowermt set path_latency_monitor.....	135
rpowermt set path_latency_threshold.....	135
rpowermt set perfmon.....	136
rpowermt set vm_perfmon.....	137
rpowermt set periodic_autorestore.....	138
rpowermt set policy.....	138
rpowermt set reactive_autorestore.....	139
rpowermt set streamio_threshold.....	140
rpowermt set device_perf_to_array_report.....	141
rpowermt set dev_inuse_to_array_report.....	141
rpowermt set port_mode.....	142
rpowermt setup add_host.....	143
rpowermt setup list_hosts.....	143
rpowermt setup remove_host.....	144
rpowermt setup update_host.....	144
rpowermt setup verify_hosts.....	145
rpowermt unregister.....	146
rpowermt update host_registration.....	147
rpowermt update lun_names.....	147

rpowermt version.....	148
Chapter 14: PowerPath family environment variables.....	149
Disable autosave of configuration.....	149
Display time stamp.....	149
Display width.....	150
Set rpowermt server to location.....	150
Migration Enabler Remote SE port.....	151
Migration Enabler Remote SE server.....	151
Disable autoregistration.....	152
Disable the lockbox.....	152
Change lockbox name and location.....	152
Set a non-default license path directory.....	153
Chapter 15: PowerPath error codes.....	154
PowerPath error messages.....	154
Chapter 16: Multipathing messages.....	162
<device_id> is alive.....	165
<device_id> is dead.....	165
/etc/powermt: Permission denied.....	165
/opt/emcpower/powercfd does not exist.....	166
/opt/emcpower/powercfd: you have to be root to execute this program.....	166
Can't get Pirp_t!.....	166
Cannot install the new configuration files.....	166
Cannot mount root on <>:<>.....	166
Cannot obtain lock on <filename>.....	167
Cannot obtain driver state.....	167
Cannot open /dev/emcp, errno <number>	167
Cannot open /etc/path_to_inst. Cannot continue.....	167
Cannot open <filename>, errno <number>.....	168
Cannot open configuration file <filename>.....	168
Cannot open dev dir for reading, errno <number>.....	168
Cannot open string Initialization error.....	168
Cannot produce new files.....	169
Cannot remove device <device_id>, name in use by application.....	169
Cannot write to file.....	169
Cannot write to <filename>	169
CLARiiON device path c3t1d0s0 is currently misconfigured.....	169
Configuration database corrupted.....	170
Configuration files contains a dead path (device).....	170
Current/Active controller for volume has changed.....	170
Device is unavailable.....	170
Error number adding pnpnpath <hdisk_name> for <hdiskpower_name> to CuAt.....	170
Error number getting pnpnpath attribute for <hdiskpower_name>.....	171
Error number updating CuDv attribute for <device_id>.....	171
Error retrieving minor number(s).....	171
Failed opening adapter <device_number>.....	171
Failover mode PNR for CLARiiON LUN <WWN#> is not recommended. Please consider PAR or ALUA	171

File deleted. Use -add to create a new file.....	172
File not deleted. No licenses will be available until file is deleted and re-created.....	172
Get adapter <devno> failed number.....	172
Incomplete installation detected.....	172
Invalid base minor number.....	172
Invalid minor number.....	173
Key <key_number> in license file is bad. Use -edit option to repair.....	173
Key <key_number> is invalid.....	173
Key <key_number> was not found.....	173
Key <key_number> is already present, ignored.....	173
License file is full -- cannot add number keys.....	174
mknod of device failed.....	174
MPxIO is enabled. Verify that no arrays are managed by both MPxIO and PowerPath.....	174
Not a multiple of 8.....	174
Number of minors out of range 1-1000.....	174
One or more expired PowerPath license key(s) detected. These are typically used...use emcprep utility to review your license keys.....	175
PRO service, action 46 failed, 0x5/0x80/0x20.....	175
Pseudo Device Name Major# Minor# <ERROR: unexpected internal error>.....	175
Read of file <file> failed.....	175
There are no keys presently registered. Some functions and capabilities may be unavailable.....	176
This version of PowerPath is NOT binary compatible with your kernel.....	176
Forceload of drv/sd failed.....	176
Get hdisk path data failed.....	176
Internal DB read failed	176
Kernel symbol versioning is not enabled in your linux kernel.....	177
Killing bus <bus> to Symmetrix <id> port <port>.....	177
Failed to load extension errno = <number>.....	177
Recognize (loop) error number on device <device_id>.....	177
Storage system <storage_system_name> PowerPath mode disabled.....	177
Not all paths to <volume> are dead.....	178
SPx Not Responding on array : Trespassing volumes.....	178
ODM PdDv database contains duplicate uniquetypes. Refer to EMC Knowledge Base Article Number: 000084491 or contact EMC Customer Support personnel.....	178
Path <path_name> to <device_id> is dead.....	178
<storage_system> path <path_name> is dead. Do you want to remove it (y/n/a/q)?.....	179
Path <bus> <tgt> <lun> to <device> is dead.....	179
Open powerpath0 failed number.....	179
PRO service action 6 failed, <Sense Key>/<ASC>/<ASCQ> for Bus <bus> Tgt <target> Lun <lun>.....	179
Restored volume <clariion/vnx_volume> to default: SP<clariion/vnx_storage_processor>.....	179
Reviving bus <bus> to Symmetrix <id> port <port>.....	180
Sleeping for memory.....	180
Syscall 8133 to <device_id> failed.....	180
Trespassed volume <clariion_volume> to SP<clariion_storage_processor>.....	180
Unable to restore volume <volume> to default controller.....	180
Unable to restore volume <clariion_volume> to default SP.....	181
Unable to follow volume <volume> to controller.....	181
Encountered unexpected error 0x<code> on Dek operation command for <device id> on handle <association_handle>.....	181
Unmanage of device <device> failed.....	181

Getting unmanaged devices failed with error.....	182
Unsupported configuration (R1/R2 setup detected).....	182
Unsupported HBA.....	182
Volume <clariion_volume> followed to SP<clariion_storage_processor>.....	182
Volume <clariion_volume> is unbound.....	182
Write failed for lack of space, check for filesystem full.....	183
Cannot open PowerPath. Initialization error.....	183
Configuration files corrupted/out of date.....	183
PowerPath configuration files present - exiting.....	183
Major number is already in use.....	184
PowerPath modules are not loaded. Ensure that the modules are loaded and try again.....	184
An older EMCpower package is already installed. Please exit the installation and remove the old EMCpower package first.....	184
PowerPath pseudo bus <bus_name> cannot be configured. Uninstall and reinstall PowerPath.....	184
Unable to create kernel process.....	184
Unable to open /etc/PowerPath Extensions.....	185
Unable to start kernel process.....	185
EMC registration data is missing or corrupt.....	185
EMCP registration file does not exist. Use -add to create it.....	185
EMCP registration file is corrupt.....	185
EMCP registration file is corrupt. Delete it?.....	185
License for <storage_system_name> storage system support is missing or expired.....	186
There are no license keys now registered.....	186
There are no license keys now registered.....	186
License not installed.....	186
License not installed.....	186
Temporary upgrade key has been superseded.....	186
Key found in old registration file is invalid.....	187
Mismatch in configuration.....	187
Missing configuration files.....	187
All paths to vol <volume> are dead.....	187
Attempted action for <storage_array> storage system support is not allowed with the current license level....	188
Attempted action is not allowed for volumes with BF or NR policy.....	188
"autostandby agingperiod" must be an integer greater than or equal to 0 and less than 24855 days (596520 hours).....	188
All paths to <device_id> are dead.....	188
All paths to <volume> are dead.....	189
Array <array> Port <port> is Disabled.....	189
Array <array> Port <port> is Enabled.....	189
Assigned volume <clariion/vnx_volume> to SP<clariion/vnx_storage_processor>.....	189
<path> is in Autostandby, force option required.....	189
Bad dev value <device>, or not under Powerpath control.....	190
Bad HBA value <hba>.....	190
Bus hba is Disabled.....	190
Bus <bus_id> to <storage_system_name> port <port_name> is alive.....	190
Bus <bus_id> to <storage_system_name> port <port_name> is dead.....	190
Cannot assemble drivers for root <>:<>.....	191
Cannot manage. verify state of the class.....	191
Device cannot be managed at this time. Please verify the state of the class <class>.....	191
Cannot open filename.....	191

Cannot remove alive device <device_id>.....	192
Cannot unmanage. Verify state of the class <class>.....	192
Cannot update persistent store.....	192
<storage_system_class> device(s) not found.....	192
Complete management of class <class> failed.....	192
Configuration file created by newer version of powermt.....	193
Device will be configured on closing all the applications to the device or upon reboot.....	193
Get hdiskpower CuAt failed number.....	193
Get hdiskpower CuDv failed number	193
Device specified must be a native device.....	193
Device(s) not found.....	194
Device not found.....	194
Device(s) not found.....	194
disable/enable is not allowed on this hba type.....	194
Error occurred loading saved driver state from file /etc/powermt_custom.xml Loading continues.....	194
Error occurred loading saved driver state from file <filename>. Loading continues.....	195
Error reading configuration file / custom file IO failure.....	195
Failed to update CLARiiON user assignable LUN names.....	195
Length of the filename exceeds the maximum limit of: <number> characters.....	195
Invalid option: dev=all is required for this version.....	196
"autostandby iopflimit" must be a number greater than 1 and less than or equal to 100000000.....	196
Path Bus <bus> Tgt <tgt> Lun <lun> to <storage_system> is set to asb:iopf	196
powermt load failed.....	197
Manage of device <device> failed.....	197
Manage of device <device> failed.....	197
Please run powermt save to persist PowerPath settings in custom file.....	197
Performance measurement is not enabled & their powermt display perf command exits.....	197
Performance monitoring already enabled.....	198
Policy is not a valid policy for <storage_system_class> storage systems.....	198
bosboot: missing proto file: /usr/lib/boot/network/chrp.hdiskpower.proto.....	198
The current boot logical volume, <device>, does not exist on <device>.....	198
powermt save failed.....	199
SCSI error on device <device>.....	199
Latency Threshold exceeded by UserDev...with a new high latency of <latency_value> msec.....	199
Unable to unconfigure path Bus X Tgt Y Lun Z as firstpath policy is set.....	199
Unmanage of device <device> failed.....	200
Renaming of temporary file to file:<filename> failed.....	200
PowerPath cannot disable some of its devices because...two host reboots are required to uninstall PowerPath.....	200
An unidentified program wants to access your computer. Don't run the program unless you know where it's from or you've used it before.....	200

Chapter 17: Migration Enabler messages.....	201
Already recovered.....	202
Already in the state requested.....	202
Couldn't read disk label.....	202
Couldn't write disk label, be sure target is at least as large as source.....	202
Current node cannot be excluded	203
Daemon communication error.....	203
Daemon failed to create thread.....	203

Daemon semaphore error.....	203
Daemon socket connection failed.....	203
Device is inaccessible.....	204
Devices already involved in a migration.....	204
An error occurred while using the base daemon, please see the log file.....	204
An internal error occurred during a SYMAPI operation. Please report to EMC.....	204
Failed to fix disk label of the target device.....	204
Failed to get disk layout of device.....	205
Failed to get target size.....	205
Failed to load EFI support library.....	205
Failed to set disk layout of device.....	205
Failed to validate geometry of target device.....	206
Failed to transfer the source reservation.....	206
The feature being requested is not currently licensed.....	206
A function call exceeded allocated time.....	206
The gatekeeper device (while using the Base Daemon) has an error (Please see the log file).....	207
Handle not found.....	207
Host can still access source device.....	207
Migration stopped due to IO error - will attempt to restart.....	207
Needs recovery.....	207
No migration is in the proper state to perform the requested operation.....	208
Node <nodename> not found.....	208
PPME Resource <Cluster Disk Name> created with powermigcl.....	208
PR Keys found on Target disk during Setup.....	208
Not in proper state to perform this operation.....	208
Possible alternate cylinder overlap detected using native source geometry.....	209
PPME cannot be used with a temporary license.....	209
Remote SE environment variable(s) not allowed in setup.....	209
Source must be a cluster shared volume.....	209
Source is a CSV (use '-cluster mscsv' to setup a cluster migration).....	209
Source disk's dependent PPME resource was changed.....	210
Source device is in use (-force will override).....	210
Source faulted, must abort migration.....	210
Source faulted, must commit migration.....	210
Source and target must be from same array.....	210
Source and target are the same device.....	211
Target faulted, must abort migration.....	211
Target device is in use.....	211
Target must be at least as large as the source.....	211
Target max_transfer attribute is too small.....	211
Target is not accessible.....	211
Target device is not attached to some nodes.....	212
Target must be a pseudo.....	212
UMD database not initialized.....	212
Unable to scrub the abandoned device.....	212
Unable to scrub the abandoned device. Device is offline.....	212
Unknown device.....	213

Chapter 18: PowerPath/VE messages.....	214
Allocation of dummy <ASynclDaemon/ TestDaemon/ SynclDaemon> entry failed. Not aborting it.....	215

Bad dev value <invalid_value>, or not under PowerPath control.....	216
Bad dev value emcpower <invalid value>, or not under PowerPath control.....	216
Bad hba value <invalid_value>, or not under PowerPath control.....	216
Cannot communicate with host, version mismatch too great.....	216
Cannot communicate with host, version mismatch too great.....	217
Could not allocate SCSI device.....	217
Could not claim the path. Status: <status>.....	217
Could not create vol Daemon.....	217
Could not initialize log name: <status>.....	217
Device <device_bus_path> is currently dead.....	217
Device(s) not found.....	218
<array_type> device(s) not found.....	218
<class> device(s) not found.....	218
Device(s) not found	218
Enqueued from <ASynclDaemon/ TestDaemon/ SynclDaemon> to Abort queue Cmd <command>.....	218
Failed to Create Pirp.....	219
Failed to create Pirp	219
Failed to create Platform scsicmd	219
Failed to create scsi Pirp	219
Failed to find IP address for <unresolved hostname>.....	219
Failed to open host file.....	219
Failed to remove host information from host file.....	220
Failed to reset the license server count.....	220
Failed to retrieve host information from host file.....	220
Failed to retrieve host list from host file.....	220
Failed to store host information into host file.....	220
Failed to store host list into host file.....	221
Failed to alloc Pirp.....	221
Failed to create Char device name. Status: <status>.....	221
Global PMI disabled.....	221
Host not found.....	221
Inquiry failed. Status: <status>.....	222
Invalid option '-d' when upgrading from a previous version of installer.....	222
Invalid passphrase during host file creation.....	222
License autoregistration is disabled.....	222
License checkout failed.....	223
License data does not match host data.....	223
License feature or attribute not found.....	223
License file does not exist or not found.....	223
License file format error or misspelling.....	223
License file is invalid (inconsistent)	223
License key or feature not found.....	224
License not installed.....	224
License not installed.....	224
License type mismatch. Hyper-V enabled: PowerPath/VE for Windows license key not found.....	224
Mpx: pData <data> size <size>.....	224
MpxRecognize failed. No resources- check system maximum path limits.....	225
MpxRecognize failed.....	225
No valid licenses found.....	225
No active path to the device found.....	225

No PMI data available for this VM.....	225
No standard UID. Using the WWN as volume UID.....	225
Not a valid sense buffer	226
The path has already been claimed	226
Path is in Autostandby, force option required.....	226
Path <path> could not be disabled. It is the last working path to volume.....	226
The Path has not been claimed.....	226
Performance Monitor is not enabled.....	227
Persistence of SCSI Res fail Value is 0, generating new one.....	227
PowerPath driver is already registered.....	227
PowerPath driver is already unlicensed.....	227
PowerPath Not Found.....	227
Probe of <path> path failed. Status <status>.....	227
ProbePath Failed path	228
Path <path>: could not issue sync/async IO command: <status>.....	228
SCSI-2 reservation/release failed with SCSI status <status>.....	228
SCSI-2 Translation command failed after the maximum number of retries.....	228
Sense Info S: <sc> ASC: <asc> ASCQ: <ascq>.....	228
Status <status>. Failed to register device. User should unclaim paths.	229
Unable to send cmd (cleanup PRI) to device. SCSI status <status>.....	229
Unable to send cmd (cleanup PRO) to device. SCSI status <status>.....	229
Could not unregister the device.....	229
User Authentication Invalid.....	229
VM either does not exist or is not registered on this host.....	230
vmk_LogRegister failed: <status>.....	230
Volume is being destroyed. Cannot open.....	230
Volume not found	230
Volume to be opened not found	230
The volume is open. Cannot unclaim last path <path>.....	230
The volume is being registered or unregistered. Cannot unclaim last path <path>.....	231
Wait for n losPending count (count) to become 0.....	231
Wait for volume to be unquiesced	231

Chapter 19: PowerPath Management Appliance messages..... 232

Although served license added, there was an error contacting License Server. Check License Server status...	232
License already present.....	232
License key not valid for the domain / IP address of this Virtual Appliance.....	232
Unable to add new license. Reason: License is invalid or expired.....	233

Appendix A: PowerPath family new and changed commands..... 234

powermig new and changed commands.....	234
powermt new and changed commands.....	234
rpowermt new and changed commands.....	235

Appendix B: List of the PowerPath load balancing and failover policies..... 237

List of the PowerPath load balancing and failover policies.....	237
---	-----

Appendix C: List of fields used in the powermt and rpowermt display command output..... 239

List of all fields used in the display output.....	239
--	-----

Appendix D: Microsoft Windows IDs.....	243
Microsoft Windows application event IDs.....	243
Microsoft Windows system event IDs.....	247
Appendix E: AIX error logging facility messages.....	249
AIX error logging facility messages.....	249
Index.....	251

1	Revision history.....	3
2	CLI install directories.....	17
3	PowerPath messages locations.....	18
4	PowerPath information, warning, and error messages locations.....	18
5	Supported commands.....	18
6	Custom file install directory.....	19
7	Valid values for class option.....	20
8	Output fields.....	48
9	Attributes and description.....	76
10	Attributes and description.....	80
11	PowerPath error messages.....	154
12	New and changed commands in 6.0 and later.....	234
13	New and changed commands in 6.0 and later.....	234
14	New and changed commands in 6.0 and later.....	235
15	Application event IDs.....	243
16	System event IDs.....	247
17	Logging facility messages.....	249

PowerPath Command Line Interface Overview

This chapter contains the following topics:

Topics:

- [Welcome to the PowerPath Command Line Interface Reference](#)
- [Installing the PowerPath CLI](#)
- [PowerPath CLI install directories](#)
- [PowerPath messages locations](#)
- [PowerPath information, warning, and error messages](#)
- [List of the supported commands by platform](#)
- [Saving and loading a PowerPath configuration](#)
- [Automatic license registration on PowerPath/VE for VMware vSphere](#)
- [Class option specification](#)
- [Specifying the wide or width option](#)

Welcome to the PowerPath Command Line Interface Reference

The PowerPath Command Line Interface (CLI) allows data center personnel to use PowerPath to manage storage resources.

Installing the PowerPath CLI

On PowerPath for AIX, HP-UX, Linux, Solaris, and Windows the PowerPath CLI is installed, along with all of the necessary support files, on each PowerPath machine when PowerPath is installed.

On PowerPath/VE for VMware vSphere, you need to install the PowerPath remote multipathing rpowermt (RTOOLS) CLI package or the PowerPath Management Appliance. The *PowerPath/VE for VMware vSphere Installation and Administration Guide* provides more information.

PowerPath CLI install directories

In a default PowerPath installation, the command utilities reside in the following directories:

Table 2. CLI install directories

Platform	Directory
AIX	/usr/sbin (all commands)
Linux	/sbin (all commands); /etc/opt/emcpower/
Solaris	/etc (emcpreg, powermt); /etc/emc/bin (emcpcfg, powermig); /opt/EMCpower/bin (emcpadm)
VMware	/opt/emc/powerpath/bin (on Linux) or <i>install_dir</i> \PowerPath (on Windows)
Windows	<i>install_dir</i> \PowerPath (all commands)

PowerPath messages locations

In a default PowerPath installation, the audit, informational, warning, and error messages are logged in the following locations:

Table 3. PowerPath messages locations

Platform	Directory
AIX	Output of the <code>errpt -a</code> command (See Appendix E) <code>/var/log/syslog</code> , if enabled.
HP-UX	<code>/var/adm/syslog</code>
Linux	<code>/var/log/messages</code>
Solaris	<code>/var/adm/messages</code>
VMware	<code>/var/log/vmkernel.log</code>
Windows	Application and System Event Logs

PowerPath information, warning, and error messages

On Linux, PowerPath messages are logged to the following two files:

Table 4. PowerPath information, warning, and error messages locations

Messages	Location
install/uninstall/upgrade messages	
All other messages	Standard operating system log file - <code>/var/log/messages</code>

List of the supported commands by platform

Displays the list of the PowerPath commands by platform.

Table 5. Supported commands

Command	AIX	Linux	Solaris	VMware	Windows
emcpadm	Yes	Yes	Yes	No	No
emcpbfs_linux	No	Yes	No	No	No
emcpcfg	No	No	Yes	No	No
emcphostid	No	No	No	No	Yes
emcplun	No	Yes	No	No	No
emcpreg	Yes	Yes	Yes	No	Yes
powercf	No	No	Yes	No	No
powerkmd	No	No	No	No	Yes
powermig	Yes	Yes	Yes	No	Yes
powermt	Yes	Yes	Yes	No	Yes
pprootdev	Yes	No	No	No	No
rpowermt	No	No	No	Yes	No

Saving and loading a PowerPath configuration

You can save the PowerPath configuration details such as the PowerPath device policies and priorities, write throttle setting and write throttle queue depth, periodic autorestore setting, and path modes to a configuration file to persist your configuration details across reboot.

At boot time, PowerPath loads configuration data from the XML formatted file having the default name `powermt_custom.xml`. Both the `powermt load` command and the `powermt save` commands use this configuration file.

On HP-UX, Linux, and Solaris The configuration is saved using the `powermt save` command and when the host is booted.

On AIX and Windows The configuration is saved only when `powermt save` is run manually.

In a default PowerPath installation, the `powermt_custom` file resides in the following directory.

Table 6. Custom file install directory

Platform	Directory
AIX, HP-UX, Linux, Solaris	/etc
VMware	N/A
Windows	<code>install-dir\PowerPath</code>

Automatic license registration on PowerPath/VE for VMware vSphere

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 the automatic license registration 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.

Successful automatic license registration

```
# 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
```

Class option specification

Displays the list of supported classes by platform.

Several PowerPath commands have a `class` option that specifies a storage system type. The following table lists the valid values for the class option added in versions 6.0 and later, with the earliest supported version listed by platform. Support in later versions (including

service packs) can be assumed if no exception in the table notes otherwise. For the list of storage systems that PowerPath currently supports, see the E-Lab Interoperability Navigator.

Table 7. Valid values for class option

Command	Storage arrays	AIX	Linux	Solaris	VMware	Windows
all	All storage arrays	4.x	4.x	4.x	5.4	4.x
clariion	CLARiiON, VNXe2 (VNXe 3200 and later), Unity arrays	4.x	4.x	4.x	5.4	4.x
ess	IBM arrays	4.x	4.x	4.x	5.4	4.x
generic	Celerra, VNXe (prior to VNXe 3200) arrays	5.7	5.5	5.5	5.7	5.5
	Data Domain	N/A	6.1	N/A	6.1	N/A
hitachi	HDS arrays	4.x	4.x	4.x	5.4	4.x
hpxp	HP arrays	4.x	4.x	4.x	5.4	4.x
invista	Invista arrays	4.x	4.x	4.x	5.4	4.x
netapp	NetApp arrays	5.7	5.7 SP3	6.0	5.9	5.7 SP2
sc	Dell Compellent arrays	6.2	6.2	N/A	6.2	6.1
symm	Symmetrix arrays	4.x	4.x	4.x	5.4	4.x
unity	Unity arrays	6.1	6.1	N/A	6.1	6.1
vnx	VNX arrays	5.7	5.7	5.5	5.8	5.7
vplex	VPLEX arrays	5.7	5.7	5.5	5.8	5.7
xtremio	XtremIO arrays	5.7 SP1	5.7 SP3	5.5 P03	5.9 SP1	5.7 SP2

4.x means that this class was a valid value in PowerPath prior to version 5.0.

N/A (not applicable) means that class is not supported for that platform.

Specifying the wide or width option

You can change the width of the output by specifying the wide or width option. The wide option is not supported on Windows platforms.

The following two examples show how setting the width option changes the output display. The first example uses the default column width (85) while the second example shows the width set to 100.

```
# powermt display dev=a width=85
Pseudo name=emcpowerz
Symmetrix ID=000192601710
Logical device ID=0806
Device WWN=60000970000192601710533030383036
state=alive; policy=SymmOpt; queued-IOS=0
=====
----- Host ----- - Stor - -- I/O Path -- -- Stats ---
### HW Path      I/O Paths      Interf.  Mode      State  Q-IOS  Errors
=====
   5 qla2xxx      sdat           FA  7g:01  active  alive   0      0
   1 qla2xxx      sdaj           FA  7g:01  active  alive   0      0
   5 qla2xxx      sdy            FA  9g:01  active  alive   0      0
   1 qla2xxx      sdn            FA  9g:01  active  alive   0      0

# powermt display dev=0 width=100
Pseudo name=emcpower0a
CLARiiON ID=WRE00021400504
Logical device ID=600601760B09000054F4356E2B17D711
state=alive; policy=BasicFailover; priority=0; queued-IOS=0
Owner: default=SP B, current=SP B
=====
----- Host ----- - Stor - -- I/O Path -- -- Stats ---
```

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

emcpadm commands

This chapter contains the following topics:

Topics:

- [emcpadm command](#)

emcpadm command

The `emcpadm` command lists available PowerPath pseudo device names and used PowerPath pseudo device names or renames a PowerPath pseudo device from a used pseudo device name to an available pseudo device name.

This command is supported on AIX, Linux, and Solaris and requires root privileges. The `emcpadm` command cannot rename pseudo devices that are in use. For example, devices that are part of a VxVM disk group are considered in use, and therefore cannot be renamed.

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

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

emcpadm check_mappings

Displays a comparison of the current mappings and the device remappings in *mappings_file*.

Syntax

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

Description

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

Options

- `-v`
Specifies verbose mode.
- `-x`
Specifies XML format for the mappings file data. You must use the XML format to manually edit the mappings file. The XML data is compatible with that viewed via tools such as PowerPath `powermt` and Symmetrix `syminq`.
- `-f pseudo device/LU mappings file`
Specifies the file name and location for the mappings file.

Check mappings

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

Imported Device  Remapping Status
Mapping
-----
```

```

hdiskpower72 no change
hdiskpower73 remaps: hdiskpower83
hdiskpower74 no change
hdiskpower75 no change
hdiskpower76 remaps: hdiskpower104
hdiskpower76 (* Device name already in use)

```

For devices which have been unmanaged or removed from PowerPath control, `emcpadm` will print the `Device not found` error message:

```

Imported Device  Remapping Status
Mapping
-----
hdiskpower63     no change
hdiskpower64     no change
hdiskpower65     Error: Device not found

```

emcpadm export_mappings

Writes and exports the current mappings to the named file.

Syntax

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

Description

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

Options

- x**
Specifies XML format for the mappings file data. You must use the XML format to manually edit the mappings file. The XML data is compatible with that viewed via tools such as PowerPath `powermt` and Symmetrix `syminq`.
- f *pseudo device/LU mappings file***
Specifies the file name and location for the mappings file.

emcpadm getfreepseudos

Retrieves a list of available PowerPath pseudo devices.

Syntax

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

Description

For each device, lists the pseudo device name, instance number, major device number, and minor device number. By default, retrieves the first free pseudo device name. If both `-n # of pseudo instances` and `-b start pseudo device` are specified, retrieves information for up to the indicated number of free pseudo devices, starting at the indicated device.

Options

- f {i|d|m}**
This command is supported only on Solaris. If you do not include the `-f` option, the `emcpadm` command treats a numeric entry (for example, 2) as an instance number. The command recognizes an `emcpowerN` name without the `-f` option.
 - `-f i` indicates an instance number, for example, 1.

- `-f d` indicates a pseudo device name, for example, `emcpower1`.
- `-f m` indicates a minor number, for example, `8`.

-n # of pseudo instances

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

-b start pseudo device

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

emcpadm getusedpseudos

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

Syntax

```
emcpadm getused[pseudos]
```

Options

None.

emcpadm help

Displays usage summary for the `emcpadm` command.

Syntax

```
emcpadm help
```

Options

None.

emcpadm import_mappings


Replaces the current mappings with the mappings in *mappings file*.

Syntax

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

Description

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

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

Before importing new mappings on a node or server:

- Preview changes with `emcpadm check_mappings`.
- Shut down all applications and database systems.
- Unmount file systems.
- Deport VxVM disk groups.

Options

-v

Specifies verbose mode.

-x

Specifies XML format for the mappings file data. You must use the XML format to manually edit the mappings file. The XML data is compatible with that viewed via tools such as PowerPath `powermt` and Symmetrix `syminq`.

-f *pseudo device/LU mappings file*

Specifies the file name and location for the mappings file.

emcpadm print_mappings

Displays the current mappings.

Syntax

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

Description

Print mappings displays the mappings in the following ways:

- `print_mappings` (without options) displays the current mappings on the node or server.
- `print_mappings -f pseudo device/LU mappings file` displays the exported mappings in *pseudo device/LU mappings file*.
- `print_mappings -p pseudo device/LU mappings file` displays the persistent mappings maintained in the PowerPath device-mapping database. These mappings should be the same as the mappings on the current node or server.

Options

-p

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

-x

Specifies XML format for the mappings file data. You must use the XML format to manually edit the mappings file. The XML data is compatible with that viewed via tools such as PowerPath `powermt` and Symmetrix `syminq`.

-f *pseudo device/LU mappings file*

Specifies the file name and location for the mappings file.

Print mappings

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

(Output truncated)

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

emcpadm renamepseudo

Changes the name of the pseudo device from the pseudo device name specified by source to the pseudo device name specified by target.

Syntax

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

Description

The target pseudo device name must be available (that is, not in use).

Options

-f {i|d|m}

This command is supported only on Solaris. If you do not include the `-f` option, the `emcpadm` command treats a numeric entry (for example, 2) as an instance number. The command recognizes an `emcpowerN` name without the `-f` option.

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

-s *src pseudo instance* -t *tgt pseudo instance*

Changes the name of the pseudo device from the pseudo device name specified by `-s` (source) to the pseudo device name specified by `-t` (target). The target pseudo device name must be available (that is, not in use)

You can specify for the device name either the instance number (for example, 12) or the complete pseudo device name (for example, `emcpower12`).

emcpadm version

Displays the version of the `emcpadm` command and the PowerPath base driver.

Syntax

```
emcpadm version
```

Options

None.

emcpbfs_linux commands

This chapter contains the following topics:

Topics:

- [emcpbfs_linux command](#)

emcpbfs_linux command

The `emcpbfs_linux` utility is used to configure, remove, or upgrade PowerPath in Boot from SAN setup.

This command is supported only on Linux.

emcpbfs_linux config

Configures PowerPath in Boot from SAN setup.

Syntax

```
emcpbfs_linux config
```

Options

None.

emcpbfs_linux remove

Uninstalls PowerPath in Boot from SAN setup.

Syntax

```
emcpbfs_linux remove
```

Options

None.

emcpcfg script commands

This chapter contains the following topics:

Topics:

- [emcpcfg script](#)

emcpcfg script

The `emcpcfg` script saves and restores a PowerPath configuration if any of the configuration files have been corrupted or removed.

The `emcpcfg` script is supported only on Solaris.

The files that represent the current configuration are:

- `/kernel/drv/emcp.conf`
- `powermt_custom.xml`
- `/etc/powermt.custom`
- `/etc/emcp_devicesDB.dat`
- `/etc/emcp_devicesDB.idx`

emcpcfg backup

Saves the current configuration files for PowerPath as TAR files in the `/etc/emc/cfgs` directory.

Syntax

```
emcpcfg backup
```

Description

The `powermt save` command is automatically run after `emcpcfg backup` is run.

Options

None.

emcpcfg cleanup

Removes or deletes all of the backed up configuration files (which were backed up only using the PowerPath `emcpcfg` script).

Syntax

```
emcpcfg cleanup
```

Options

None.

emcpcfg list cfgs

Lists the available saved configurations.

Syntax

```
emcpcfg list_cfgs
```

Options

None.

emcpcfg restore

Restores the specified last good configuration that was saved.

Syntax

```
emcpcfg restore [-reboot] <saved_PowerPath_configuration>
```

Description

After running emcpcfg restore, reboot the host to return to the last saved configuration.

Options

-reboot

Performs an automatic reboot. Otherwise, a manual reboot is required.

saved_PowerPath_configuration

Last saved configuration file.

emcpcfg set cfgs limit

Sets the limits on number of configurations that can be backed up.

Syntax

```
emcpcfg set_cfgs_limit <number>
```

Description

For example, if the new limit is set to three configurations, as opposed to a previous limit of five configurations, then the two older configurations are removed to maintain only three backups. You do not need to delete the configurations manually.

Options

number

Number of configurations that can be backed up.

emcphostid commands

This chapter contains the following topics:

Topics:

- [emcphostid command](#)

emcphostid command

PowerPath requires a unique host ID to reserve shared disks in clustered systems. The `emcphostid` command sets the Host ID.

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

On HP-UX hosts, the PowerPath installation sets the host ID to the largest numerical MAC address or the hash of the Machine ID found at the time of installation. PowerPath installation sets the host ID to the largest numerical MAC address or the hash of the Machine ID found at the time of installation. On Windows hosts, run the Microsoft sysprep utility if the same operating system image is used to set up cluster nodes to ensure that each node receives a unique host ID.

The host ID is saved in `/etc/emc/emcphostid` on HP-UX hosts. On Windows hosts, the host ID is created by hashing the NetBIOS name of the computer on every boot.

The emcphostid command on Windows hosts

PowerPath requires a unique host ID to reserve shared disks in clustered systems.

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

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

emcphostid check

Checks the uniqueness of host IDs for the specified list of hosts.

Syntax

```
emcphostid check -host <HostName1> <HostName2> ...<HostNameN>
```

Description

Run `emcphostid check` to check for conflicts after installing PowerPath in cluster environment and after changing the host ID manually.

Options

`-host HostName`

Retrieves and prints the host ID for the identified host.

emcphostid get

Retrieves and prints the host ID that is currently stored in the Windows registry.

Syntax

```
emcphostid get [-host <HostName>]
```

Description

Get also checks to see whether the current host ID differs from the default host ID, and if so, prints a warning.

Options

-host *HostName*

Retrieves and prints the host ID for the identified host.

emcphostid help

Displays usage summary for the `emcphostid` command.

Syntax

```
emcphostid help
```

Options

None.

emcphostid set

Sets or resets the host ID.

Syntax

```
emcphostid set [-id <id>] [-host <HostName>] [-no_prompt]
```

Description

By default, if a host ID already exists, prompts for confirmation before resetting the host ID. Use `-no_prompt` to suppress this prompt. Without the `id` option, `emcphostid set` sets the host ID to the highest MAC address found on the host or the hash of the Machine ID at the time the command is run.

If the MAC address or Machine SID cannot be retrieved, a warning is displayed, and the host ID is set or not set or changed. You can identify the host ID as an 8-digit hexadecimal number with or without the 0x prefix.

Options

-id *id*

Host id.

-host *HostName*

Retrieves and prints the host ID for the identified host.

-no_prompt

Suppresses the default confirmation that appears when you run a command.

emcplun commands

This chapter contains the following topics:

Topics:

- [emcplun_linux utility](#)

emcplun_linux utility

The `emcplun_linux` utility is used to perform the steps necessary for dynamically adding or removing a LUN on PowerPath hosts.

emcplun_linux help

Prints usage screen.

Syntax

```
emcplun_linux help [operation]
```

Options

operation

Prints usage for: list, remove, or scan.

emcplun_linux list

Lists pseudo devices.

Syntax

```
emcplun_linux list [<pseudo-list>] [default: all]
```

Options

list [*pseudo-list*]

Lists pseudo devices and corresponding details [default: all].

```
list emcpowerX emcpowerY . .
```

emcplun_linux remove

Removes pseudo devices.

Syntax

```
emcplun_linux remove <pseudo-list>|all [-noprompt]
```

Options

pseudo-list

List of pseudo devices that are to be removed. This option removes the pseudo devices only if pseudo device is not in use.

```
emcpowerX emcpowerY ...
```

all

Remove all pseudo devices and associated native devices. This option removes the pseudo devices only if they are not in use.

-noPrompt[-no

Suppresses the default confirmation that appears when you run a command.

Remove

Example for removing emcpower<id>

```
# /etc/opt/emcpower/emcplun_linux remove emcpowera
```

If no operation is given, emcplun_linux script runs in interactive mode

emcplun_linux scan

Scans the bus and configures new devices under PowerPath.

Syntax

```
emcplun_linux scan <option> [-noprompt]
```

Options

- hba**
Performs the scanning on all available HBAs.
Supported adapter include Qlogic, Emulex, Brocade and Cisco HBAs. These included FC/iSCSI HBAs and CNAs.
- iscsi**
Performs the scanning on all available iSCSI Software initiator sessions.
- open-fcoe**
Performs the scanning on all available Open-FCoE Software initiator sessions.
- all**
Performs the scanning on all available HBAs and Software initiators.
- <# [#]>**
Performs the scanning on specific available SCSI Host numbers. This option is not supported on Software initiators.

```
2 3 4
```

aid=Array ID# [tgt=#] [hlu=#]

Limits the scanning to the specified Storage system (Array ID) and its SCSI Targets. Discovers and scans SCSI Host and Target numbers for a Storage system (Array ID).

All other options are optional [SCSI Target number(s) and HLUs]. The default hlu value is 0-20 if the hlu option is not specified.

```
aid=FCNCH090201457 [tgt=3,7-9,2] [hlu=3,7-9]
```

hba=# tgt=# [hlu=#]

Limits the scanning to the specified SCSI Host numbers, Targets, HLUs. The default hlu value is 0-20 if the hlu option is not specified.

SCSI Target numbers are mandatory and HLU numbers are optional.

```
hba=1,2,4-7 tgt=3,7-9,2 [hlu=3,7-9]
```

-noPrompt|-no

Suppresses the default confirmation that appears when you run a command.

Scan

Example for scanning hba

```
# /etc/opt/emcpower/emcplun_linux scan hba
```

Example for scanning iSCSI software initiator

```
# /etc/opt/emcpower/emcplun_linux scan iscsi
```

If no operation is given, emcplun_linux script runs in interactive mode

```
Use '/etc/opt/emcpower/emcplun_linux help' for usage of the script.  
Use '/etc/opt/emcpower/emcplun_linux help scan' for information on scan operation.
```

emcprep commands

This chapter contains the following topics:

Topics:

- [emcprep utility](#)

emcprep utility

The `emcprep` utility manages PowerPath license registration.

On some platforms, `emcprep` is invoked automatically by the PowerPath installation processes. On other platforms, you must invoke the command manually. Once installation is complete, `emcprep` can be invoked manually on any supported platform.

To run `emcprep`, you must be logged in as the root user.

NOTE: The `emcprep` utility will not accept keys or display keys with the PowerPath Encryption with RSA capability. If an encryption key is detected you will receive an error message. The multipathing messages in this book provides more information.

The emcprep command on UNIX and Linux

The `emcprep` utility stores license keys in the registration database, by default in `/etc/emcp_registration`.

emcprep add key

Adds one or more keys to the registration database.

Syntax

```
emcprep [-f <pathname>] -add <key> [<key> ...]
```

Description

The `emcprep` command only adds or removes licenses from the database. The licenses are not activated until one of the following commands are executed: `powermt config`, `powermt check_registration`, `powermt load`, or `reboot`.

Options

-f *pathname*

Provides the optional pathname of the registration database for UNIX and Linux hosts. The default is `/etc/emcp_registration`. This file should not be manually edited. It also should not be moved, as that will prevent installed products from finding valid license keys.

-add *key* [*key* ...]

Identifies the PowerPath license key or keys.

emcprep check key

Checks the validity of one or more keys and display the product information associated with each key.

Syntax

```
emcprep [-f <pathname>] -check <key> [<key> ...]
```

Options

-f *pathname*

Provides the optional pathname of the registration database for UNIX and Linux hosts. The default is `/etc/emcp_registration`. This file should not be manually edited. It also should not be moved, as that will prevent installed products from finding valid license keys.

-check key [key ...]

Checks the validity of one or more keys and display the product information associated with each key.

emcpreg edit

Edits the product information database.

Syntax

```
emcpreg [-f <pathname>] -edit
```

Description

Each key in the database is displayed and you are prompted to remove or retain it.

Options

-f pathname

Provides the optional pathname of the registration database for UNIX and Linux hosts. The default is `/etc/emcp_registration`. This file should not be manually edited. It also should not be moved, as that will prevent installed products from finding valid license keys.

emcpreg install

Installs the license key for PowerPath.

Syntax

```
emcpreg [-f <pathname>] -install
```

Options

-f pathname

Provides the optional pathname of the registration database for UNIX and Linux hosts. The default is `/etc/emcp_registration`. This file should not be manually edited. It also should not be moved, as that will prevent installed products from finding valid license keys.

emcpreg list

Lists the product information for all keys in the registration database.

Syntax

```
emcpreg [-f <pathname>] -list
```

Options

-f pathname

Provides the optional pathname of the registration database for UNIX and Linux hosts. The default is `/etc/emcp_registration`. This file should not be manually edited. It also should not be moved, as that will prevent installed products from finding valid license keys.

emcpreg remove key

Removes one or more keys from the registration database.

Syntax

```
emcpreg [-f <pathname>] -remove <key> [<key> ...]
```

Description

The keys may be added back later if needed.

The `emcprep` command only adds or removes licenses from the database. The licenses are not activated until one of the following commands are executed: `powermt config`, `powermt check_registration`, `powermt load`, or `reboot`.

Options

-f *pathname*

Provides the optional pathname of the registration database for UNIX and Linux hosts. The default is `/etc/emcp_registration`. This file should not be manually edited. It also should not be moved, as that will prevent installed products from finding valid license keys.

-remove *key* [*key* ...]

Removes one or more keys from the registration database. The keys may be added back later if needed. To change the driver state, you need to do an additional step: run the `powermt config`, `powermt check_registration`, `powermt load`, or the `reboot` command.

The emcprep command on Windows

The `emcprep` utility stores license keys in the registration database, by default in the Windows registry.

emcprep add key

Adds one or more keys to the registration database.

Syntax

```
emcprep -add <key> [<key> ...]
```

Description

The `emcprep` command only adds or removes licenses from the database. The licenses are not activated until one of the following commands are executed: `powermt config`, `powermt check_registration`, `powermt load`, or `reboot`.

Options

-add *key* [*key* ...]

Identifies the PowerPath license key or keys.

emcprep check key

Checks the validity of one or more keys and display the product information associated with each key.

Syntax

```
emcprep -check <key> [<key> ...]
```

Options

-check *key* [*key* ...]

Checks the validity of one or more keys and display the product information associated with each key.

emcprep edit

Edits the product information database.

Syntax

```
emcprep -edit
```

Description

Each key in the database is displayed and you are prompted to remove or retain it.

Options

None.

emcprep install

Installs the license key for PowerPath.

Syntax

```
emcprep -install
```

Options

None.

emcprep list

Lists the product information for all keys in the registration database.

Syntax

```
emcprep -list
```

Options

None.

emcprep remove key

Removes one or more keys from the registration database.

Syntax

```
emcprep -remove <key> [<key> ...]
```

Description

The keys may be added back later if needed.

The `emcprep` command only adds or removes licenses from the database. The licenses are not activated until one of the following commands are executed: `powermt config`, `powermt check_registration`, `powermt load`, or `reboot`.

Options

-remove key [key ...]

Removes one or more keys from the registration database. The keys may be added back later if needed. To change the driver state, you need to do an additional step: run the `powermt config`, `powermt check_registration`, `powermt load`, or the `reboot` command.

powercf commands

This chapter contains the following topics:

Topics:

- [powercf command](#)

powercf command

During system boot on Solaris hosts, the `powercf` utility configures PowerPath devices by scanning HBAs for both single-ported and multiported storage system logical devices.

A multiported logical device shows up on two or more HBAs with the same storage system subsystem/device identity. The identity comes from the serial number for the logical device. For each storage system logical device found in the scan of the HBAs, `powercf` creates a corresponding `emcpower` device entry in the `emcp.conf` file and it saves a primary path and an alternate primary path to that device.

After PowerPath is installed, you need to run `powercf` only when the physical configuration of the storage system or the host changes. Configuration changes that require you to reconfigure PowerPath devices include:

- Adding or removing HBAs
- Adding, removing, or changing storage system logical devices
- Changing the cabling routes between HBAs and storage system ports
- Adding or removing storage system interfaces

The `powercf` utility resides in the `/etc` directory. You must have superuser privileges to use the utility.

powercf

The `powercf` utility scans HBAs for single-ported and multiported storage system logical devices and compares those logical devices with PowerPath device entries in `emcp.conf`.

Syntax

```
powercf -q|-Z
```

Description

The `/kernel/drv/emcp.conf` file lists the primary and alternate path to each storage system logical device and the storage system device serial number for that logical device. The `powercf -q` command updates the existing `emcp.conf` file or creates a new one if it does not already exist.

Options

- q
Runs `powercf` in quiet mode.
Updates the `emcp.conf` file by removing PowerPath devices not found in the HBA scan and adding new PowerPath devices that were found. It saves a primary and an alternate path to each PowerPath device.
- Z
Runs automatically during system boot.
Configures an SRDF-enabled server to be bootable from an R2 mirror of a Symmetrix-based `emcpower` boot disk by a remote host.
Should be run manually whenever such a server's Symmetrix volume configuration changes due to the addition or deletion of volumes.

powerkmd remove

Removes persistent mappings for the devices.

Syntax

```
powerkmd remove [-device|-dev <device>] [-all] [-noPrompt|-no] [-force]
```

Options

-device|-dev *device*

Identifies the device ID for which persistent mapping should be removed.

-all

Removes persistent mapping for all dead devices.

-noPrompt|-no

Suppresses the default confirmation that appears when you run a command.

-force

Removes persistent mapping for the device while some migrations are active. Ensure that the devices for which persistent mapping is being removed are not part of any outstanding migration.

powerkmd version

Displays the version of powerkmd running on the host.

Syntax

```
powerkmd version
```

Options

None.

powermig commands

This chapter contains the following topics:

Topics:

- [powermig command](#)

powermig command

The `powermig` command manages migration tasks in a PowerPath environment.

powermig abort

Aborts the migration and returns the migration session to the setup state.

Syntax

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

Description

From Setup, you can either restart the synchronization with the `powermig sync` command, or clean up the migration with the `powermig cleanup` command. When you clean up the migration, you remove selected data copied from the source to the target.

You can abort a migration session at any point after running `powermig sync` and before running `powermig commit`. Once a migration has been committed, it cannot be aborted.

After aborting a migration, you can restart synchronization of the source and target or you can clean up the target to remove selected data that was copied from the source.

The `powermig abort` command fails if the handle is invalid or migration is not in a state that permits an abort.

Migration state

Required migration state to run the command: syncing, sourceSelected, or targetSelected

Migration state after the command completes: setup

Options

-handle|-hd *migrationHandle*

Number that identifies the migration session. It is assigned when you run the `powermig setup` command. You must enter the handle when running commands that reference the migration session.

NOTE: If `-all` and `-handle` are both provided then the `-all` option takes precedence and the `-handle` value is ignored.

-all

Performs abort for all migration sessions in the syncing, sourceSelected, or targetSelected state. Migrations in other states are ignored. Migrations in the needsRecovery state that cannot be recovered with the `powermig recover` command must be aborted individually by handle.

-noPrompt|-no

Suppresses the default confirmation that appears when you run a command.

powermig cleanup

Cleans up data on the source or target logical unit, depending on the state of the migration when you run the command.

Syntax

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

Description

This cleanup prevents the application or operating system from potentially seeing two logical units with identical data, which could cause confusion.

- Run this command from the committed state to clean up selected data on the source logical unit.
- Run this command from the setup state to clean up selected data copied from the source to the target logical unit. For example, you may want to destroy an aborted migration or a migration that has been setup, but is no longer desired.

Once this command completes successfully, the migration handle is no longer valid, and no record of the migration exists.

Migration state

Required migration state to run the command: committed (cleans up source) or setup (cleans up target)

Migration state after the command completes: Not applicable. No record of the migration session remains.

Options

-handle|-hd *migrationHandle*

Number that identifies the migration session. It is assigned when you run the `powermig setup` command. You must enter the handle when running commands that reference the migration session.

NOTE: If -all and -handle are both provided then the -all option takes precedence and the -handle value is ignored.

-all

Performs cleanup for all migration sessions in the committed or setup state. Migrations in other states are ignored.

NOTE: Migrations in the needsRecovery state requiring the -force option must be cleaned up individually by handle.

-format|-form

NOTE: This command option destroys all data on the source logical unit leaving it labeled and ready for use.

This option is supported only on Linux and Solaris hosts and is not supported with Virtual Encapsulation (INVE) technology type. This option will cause `powermig cleanup` to take substantially longer because it performs a full format of the disk.

Used for migrations in the committed state when a full disk format of the source logical unit is desired. This option has no effect on migrations in the setup state.

-force|-forc

Removes the migration session from Migration Enabler control. Use this option when the following are true:

- `powermig cleanup` generates an error because the logical unit requiring cleanup is permanently inaccessible.
- you are sure you want to erase all knowledge of the migration from Migration Enabler.

-noPrompt|-no

Suppresses the default confirmation that appears when you run a command.

Cleanup

```
# powermig cleanup -handle 1
Cleanup migration for handle 1 ? [Yes]/No: y
```

powermig commit

Commits the migration by permanently designating the target as the recipient of all I/O requests.

Syntax

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

Description

After this command is run, Migration Enabler no longer keeps the source and target synchronized and I/O is not sent to the source. The migration session must be in the targetSelected state for this command to succeed.

When you commit a migration, the migration enters one of two states, depending on whether the source logical unit is a pseudo- or native-named device. If the source is a pseudo-named device, the migration enters the committed state. If the source is a native-named device, the migration enters the committedAndRedirected state.

The `powermig commit` command fails if the handle is not valid or the migration is not in the targetSelected state.

Migration state

Required migration state to run the command: targetSelected

Migration state after the command completes: committed (for a pseudo source device) or committedAndRedirected (for a native source device)

Options

`-handle|-hd migrationHandle`

Number that identifies the migration session. It is assigned when you run the `powermig setup` command. You must enter the handle when running commands that reference the migration session.

 **NOTE: If `-all` and `-handle` are both provided then the `-all` option takes precedence and the `-handle` value is ignored.**

`-all`

Performs commit for all migration sessions in the targetSelected state. Migrations in other states are ignored.

`-noPrompt|-no`

Suppresses the default confirmation that appears when you run a command.

powermig getHandle

Displays the handle for a migration session in progress.

Syntax

```
powermig getHandle -dev <device>|-src <source> -tgt <target>
```

Description

The handle is assigned when you run the `powermig setup` command, and it is needed for other `powermig` commands that reference the migration session.

Migration state

Required migration state to run the command: Any state.

Migration state after the command completes: This command does not change the migration state.

If a migration with the specified device name or names does not exist, you see this message: Handle not found

Options

`-dev device`

The source or target device name for the migration whose handle you want to retrieve. The device can be a pseudo- or native-named device. The device name must match the name entered when you set up the migration session.

-src *source*

The source-device name involved in the migration whose handle you want to retrieve. The `-source` option required that you also enter the target-device name as an option in the `-target` flag. Alternatively, use the `-dev` option with either the source- or target-device name as an easier way to retrieve the handle.

-tgt *target*

The target-device name involved in the migration whose handle you want to retrieve. The `-target` option requires that you also enter the source-device name as an option in the `-source` flag. Alternatively, use the `-dev` option with either the source- or target-device name as an easier way to retrieve the handle.

getHandle

```
# powermig getHandle -dev emcpower72a
Migration Handle = 1
```

powermig help

Displays usage summary for the `powermig` command.

Syntax

```
powermig help <operation> <options>
```

Description

By default, supported commands are listed. Specify a specific command to display help for that command.

Migration state

Required migration state to run the command: Any state.

Migration state after the command completes: This command does not change the migration state.

Options

operation

Name of the `powermig` operation about which you want to view information. The possible values are: `abort`, `cleanup`, `commit`, `getHandle`, `info`, `pause`, `query`, `recover`, `resume`, `selectSource`, `selectTarget`, `setup`, `sync`, `throttle`, `undoRedirect`, `version`

options

Name of the option you want to view help. The possible values are: `-handle`, `-format`, `-force`, `-noPrompt`, `-query`, `-source`, `-target`, `-techType`, `-throttleValue`

Help

```
# powermig help setup
```

powermig info

Displays information about a specific migration session or about all active migrations.

Syntax

```
powermig info -handle <migrationHandle> [-query]|-all [-query]
```

Description

Output includes:

- source and target device names for the migration session
- underlying technology used in the migration
- state of the migration
- logical-unit fault information, where relevant

The synchronization status appears if you use the `-query` option.

Migration state

Required migration state to run the command: Any state.

Migration state after the command completes: This command does not change the migration state.

A no migrations found message displays when no migrations are present on the host.

Options

`-handle|-hd migrationHandle`

Number that identifies the migration session. It is assigned when you run the `powermig setup` command. You must enter the handle when running commands that reference the migration session.

NOTE: If `-all` and `-handle` are both provided then the `-all` option takes precedence and the `-handle` value is ignored.

`-query|-q`

Queries all migrations for their synchronization status. This is a resource-intensive operation if a large number of migrations are in progress and in the syncing state.

`-all`

Specifies that information be displayed for all Migration Enabler migrations in progress.

Info

```
# powermig info -handle 14
=====
Hnd   Source          Target          Tech          State
===   =====          =====          =====          =====
14    emcpower48a      emcpower45a    HostCopy      syncing(12%) [PAUSED]
```

The following output is applicable only for migrations in an MSCS cluster environment. The `powermig info` command here displays cluster and non-cluster migrations initiated from the migration node. It does not display cluster migrations initiated from other nodes. You will be able to differentiate between cluster and non-cluster migration by (cl) flag near the techType value. The handle value for the cluster migration starts at c1 and continues for the number of cluster migrations.

```
=====
Hnd   Source          Target          Tech          State
===   =====          =====          =====          =====
1    harddisk18      harddisk19      HostCopy (cl)  setup
2    harddisk16      harddisk17      HostCopy       setup
```

powermig options

Sets PowerPath Migration Enabler options which apply to all migrations, or all migrations within a specific technology type.

Syntax

```
powermig options [-<option> <value>]
```

Description

Running `powermig options` with no option merely displays all available options and their values.

NOTE: This command changes the Host Copy ceiling value, and also switches between the traditional throttle based behavior and the new Host Copy ceiling behavior. This command affects all the current and new Host Copy migrations.

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

Migration state

Required migration state to run the command: Any state.

Migration state after the command completes: This command does not change the migration state.

Options

option

Name of the option for which you want to view. The possible value is: `hostcopy_ceiling`

value

Name of the value for which you want to view. The possible values are: `CeilingValue`, `none`.

ceilingValue

Specifies Host Copy ceiling value in megabytes (MB) per second. This integer value can range from 1 to 1000000 (1,000,000) MB per second.

none

Disables Host Copy ceiling. When set to `none`, per migration throttling and the active sessions limit are enabled.

Options

```
# powermig options -hostcopy_ceiling 10
hostcopy_ceiling: 10 Megabytes per second
```

powermig pause

Pauses a migration session in the syncing state and pauses the synchronization of the source and target logical units.

Syntax

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

Description

 **NOTE: This command is supported with all technologies except Virtual Encapsulation (INVE).**

In the case of a SYMCLONE migration, `powermig pause` is supported only if the source and target are the same size. Pausing a migration allows you to free host resources for other operations. When you resume the migration, the synchronization process picks up where it left off.

Migration state

Required migration state to run the command: `syncing`

Migration state after the command completes: `syncing [PAUSED]`

Options

`-handle|-hd migrationHandle`

Number that identifies the migration session. It is assigned when you run the `powermig setup` command. You must enter the handle when running commands that reference the migration session.

 **NOTE: If `-all` and `-handle` are both provided then the `-all` option takes precedence and the `-handle` value is ignored.**

`-all`

Performs pause for all migration sessions in the syncing state. Migrations in other states are ignored.

`-noPrompt|-no`

Suppresses the default confirmation that appears when you run a command.

powermig query

Displays the state of a migration.

Syntax

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

Migration state

Required migration state to run the command: Any state.

Migration state after the command completes: This command does not change the migration state.

Options

-handle|-hd *migrationHandle*

Number that identifies the migration session. It is assigned when you run the `powermig setup` command. You must enter the handle when running commands that reference the migration session.

NOTE: If `-all` and `-handle` are both provided then the `-all` option takes precedence and the `-handle` value is ignored.

-all

Performs query for all migration sessions.

-noPrompt|-no

Suppresses the default confirmation that appears when you run a command.

Query

```
# powermig query -hd 92
Handle: 92
Source: emcpower1a (512.01 MB)
Target: emcpower2a (4.21 GB, thin)
Technology: HostCopy
Migration state: syncing
Throttle Value: 2
Percent InSync: 46% (237.50 MB copied)
Sync Start Time: Mon Aug 1 11:09:21 2011
Recent Throughput: 14.84 MB/s (118.75 MB in 8 seconds)
Estimated Time to Completion (using recent throughput): 18 seconds
Estimated Completion Time: Mon Aug 1 11:09:53 2011
```

where:

NOTE: The output fields that display vary based on the migration state.

Table 8. Output fields

Field	Description
Handle	Number that identifies the migration session.
Source	Source logical unit. Includes the LUN size for technologies that sync and thin if the source is virtually provisioned. LUN size is not displayed for Virtual Encapsulation (INVE) migrations. Once committed, if the devices have been swapped, the size information will also be swapped.
Target	Target logical unit. Includes the LUN size for technologies that sync and thin if the target is virtually provisioned. LUN size is not displayed for Virtual Encapsulation (INVE) migrations. Once committed, if the devices have been swapped, the size information will also be swapped.
Technology	Underlying technology used in the migration.

Table 8. Output fields (continued)

Field	Description
Migration state	State of the migration.
Percent InSync	Percentage the source and target logical units are in-sync and displays the amount of data copied.
Throttle Value	Value of the throttle, which slows down or speeds up the synchronization of the source and target logical units.
Sync Start Time	Date and time that the migration begins syncing. This begins after running <code>powermig sync</code> .
Sync End Time	Date and time that the migration completes syncing. This occurs after the last <code>powermig query</code> runs and moves the migration to the <code>sourceSelected</code> state. For Host Copy only, the last <code>powermig query</code> occurs automatically at the final checkpoint when the copy is complete.
Total Sync Time	Elapsed time from the start of syncing until the last <code>powermig query</code> during the syncing state runs. This includes time that the host was down, but not time the migration was paused.
Recent Throughput	Throughput and the amount of data copied since <code>powermig query</code> was last run.
Overall Average Throughput	Determined by dividing the source logical unit size by the total sync time. This field displays after the migration has finished syncing.
Estimated Time to Completion	Estimated sync time remaining using the most recent throughput calculation.
Estimated Completion Time	Estimated date and time when sync should complete.

powermig recover

Recovers the migration command that was in progress when an interruption occurred due to a migration error or process crash.

Syntax

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

Description

If a migration error occurs, the recovery may fail until the cause of the error is identified and resolved.

Migration state

Required migration state to run the command: `needsRecovery`

Migration state after the command completes: The next state in the migration process. For example, if the transition from syncing to `sourceSelected` is interrupted, after running `powermig recover` the migration session enters the `sourceSelected` state.

Options

-handle|-hd *migrationHandle*

Number that identifies the migration session. It is assigned when you run the `powermig setup` command. You must enter the handle when running commands that reference the migration session.

 **NOTE: If `-all` and `-handle` are both provided then the `-all` option takes precedence and the `-handle` value is ignored.**

-all

Performs recover for all migration sessions in the `needsRecovery` state. Migrations in other states are ignored.

`-noPrompt|-no`

Suppresses the default confirmation that appears when you run a command.

powermig resume

Restarts a migration that was paused by the powermig pause command.

Syntax

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

Description

 **NOTE: This command is supported with all technology types except Virtual Encapsulation (INVE).**

In the case of a TimeFinder/Clone migration, the `powermig resume` command works only if the source and target are the same size. Use the `powermig info -all -query` command to determine which migrations are paused. A migration state of PAUSED indicates a paused migration session.

Migration state

Required migration state to run the command: paused

Migration state after the command completes: syncing

Options

`-handle|-hd` *migrationHandle*

Number that identifies the migration session. It is assigned when you run the `powermig setup` command. You must enter the handle when running commands that reference the migration session.

 **NOTE: If `-all` and `-handle` are both provided then the `-all` option takes precedence and the `-handle` value is ignored.**

`-all`

Performs resume for all migration sessions in the paused state. Migrations in other states are ignored.

`-noPrompt|-no`

Suppresses the default confirmation that appears when you run a command.

powermig selectSource

Designates the source logical unit as the recipient of all I/O requests.

Syntax

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

Description

When this command completes successfully, the migration is in the source and target logical units continue to be synchronized. In the `sourceSelected` state, the This command allows you to back out of the `targetSelected` state if necessary. Synchronization of the source and target continues after running this command.

The `powermig selectSource` command fails if the handle is invalid or if the migration is not in the `targetSelected` state or the `sourceSelected` state. You cannot commit a migration from this step. To commit a migration, select the target first and then commit the migration.

Migration state

Required migration state to run the command: `targetSelected`

Migration state after the command completes: `sourceSelected`

Options

-handle|-hd *migrationHandle*

Number that identifies the migration session. It is assigned when you run the `powermig setup` command. You must enter the handle when running commands that reference the migration session.

 **NOTE: If `-all` and `-handle` are both provided then the `-all` option takes precedence and the `-handle` value is ignored.**

-all

Performs `selectSource` for all migration sessions in the `targetSelected` state. Migrations in other states are ignored.

-noPrompt|-no

Suppresses the default confirmation that appears when you run a command.

SelectSource

```
# powermig selectSource -handle 1  
Transition to sourceSelected state for handle x? [Yes]/No:
```

powermig selectTarget

Designates the target logical unit as the recipient of all I/O requests.

Syntax

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

Description

When this command completes successfully, the migration transitions to the `TargetSelected` state. In the `targetSelected` state the source and target logical units continue to be synchronized.

The `powermig selectTarget` command fails if the handle is invalid or the device is not in the `sourceSelected` state or the `targetSelected` state. When the target logical unit is larger than the source, additional space on the target is unusable until the migration is committed.

 **NOTE: This command is not supported on MSCS cluster migrations.**

Migration state

Required migration state to run the command: `sourceSelected`

Migration state after the command completes: `targetSelected`

Options

-handle|-hd *migrationHandle*

Number that identifies the migration session. It is assigned when you run the `powermig setup` command. You must enter the handle when running commands that reference the migration session.

 **NOTE: If `-all` and `-handle` are both provided then the `-all` option takes precedence and the `-handle` value is ignored.**

-all

Performs `selectTarget` for all migration sessions in the `sourceSelected` state. Migrations in other states are ignored.

-noPrompt|-no

Suppresses the default confirmation that appears when you run a command.

SelectTarget

```
# powermig selectTarget -handle 1
Transition to targetSelected state for handle x? [Yes]/No:
```

powermig setup

Defines the source and target device name involved in the migration and the underlying technology used in conjunction with Migration Enabler.

Syntax

```
powermig setup -techType <techType> -src <sourceName> -tgt <targetName>|-file <fileName> [-throttleValue <throttleValue>] [-cluster mscd|mscsv] [-techHost <hostname> [-techPort <port#>]] [-exclude_nodes <nodename1> <nodename2>] [-force] [-noprompt|-no]
```

Description

The `powermig setup` command output displays a migration handle, which identifies the migration session. You need the migration handle in subsequent `powermig` commands that reference the migration session.

The migration handle is a number between 1 and 30,000. The first migration handle is 1, and each subsequent migration handle increases by 1, until the maximum of 30,000 is reached. Once the maximum is reached, the handle number restarts at 1, or at the lowest number not in use by another migration session.

After the `powermig setup` command completes, the migration transitions to the setup state. You can then continue the migration by synchronizing the source and target logical units. Or you can abandon the migration by running the `powermig cleanup` command to remove selected data from the source or target logical unit.

This command fails if:

- The source or target is inaccessible.
- The source or target is part of another migration session.
- The source and target devices have mismatched labels or the target is unlabeled.

After this command completes, Migration Enabler protects the source and target from removal or remapping. Once the migration enters the setup state, the target-device name and any aliases become inaccessible, and attempts to access the target fail.

NOTE: If environment variables are present, they will be used to replace the values stored in the Migration Enabler database (UMD) at setup, causing `powermig setup` to fail.

On AIX and Windows, when the `max_transfer` attribute of the target is smaller than that of the source, `powermig setup` may fail. Modify the target's `max_transfer` size to be at least equal to that of the source.

Migration state

Required migration state to run the command: Not applicable as the migration session is not set up until after running the command.

Migration state after the command completes: setup

Options

-techType|-tt *techType*

The technology type used in conjunction with Migration Enabler. Where `techType` is OR (Open Replicator), INVE (Virtual encapsulation), HOSTCOPY (HostCopy), and SYMCLONE (TimeFinder/Clone). Migrations to an encrypted thin target, or from an encrypted thin source, do not preserve the source's allocation on the target.

-src *sourceName*

Source device from which you want to migrate data. Specify the device name as a base name (for example, `emcpowerc`), or as a fully qualified path name (for example, `/dev/rdisk/emcpowerc`). For the `-src` option, you must specify the source Name used by the application.

NOTE: The *sourceName* must be the same as the name applications are configured to use.

-tgt *targetName*

Target device to which you want to migrate data. Specify the device name as a base name (for example, `emcpowerc`), or as a fully qualified path name (for example, `/dev/rdisk/emcpowerc`). On Windows, use the name as displayed by `powermt display dev=all`.

-file *fileName*

File name that contains migration pairs. For example, `-file /tmp/pairs`. The migration pair consists of a source and target device. Create the text file with one migration pair per line. Migration pairs with handles that exist from previous migrations will be ignored.

If an error displays for a migration pair, setup will continue with the next pair. After you resolve the cause of the error, run **powermig setup** with the same file again to complete.

-throttleValue|-tv *throttleValue*

Value of the throttle, which slows down or speeds up the synchronization of the source and target logical units. Values range from 0 to 9, where 0 is the fastest speed and 9 is the slowest. The default value is 5 for Open Replicator, 0 for SYMCLONE, and 2 for HostCopy migrations. The throttle allows you to change (while a synchronization is in progress) the speed at which data is synchronized. The default values are: 5 for Open Replicator; 2 for HostCopy; and 0 for SYMCLONE.

This option is valid only with Open Replicator, SYMCLONE, or HostCopy migrations. With Open Replicator, the throttle value determines resource consumption on the target array. With HostCopy, throttle value affects resource consumption on the host.

If you need a faster copy, lower it, keeping in mind this will utilize more host resources. If you need to free up your host's resources increase the throttle value to slow down HostCopy. If you need to postpone utilizing host resources use `powermig pause`, then `powermig resume` when host resources are more readily available for copying.

-techHost|-th *hostname*

Host name or IP address of the remote Solutions Enabler server used to optionally specify remote server name (or IP address) where the Solutions Enabler server daemon is running. This is only applicable when using Solutions Enabler remotely with Migration Enabler technology types Open Replicator (OR) and SYMCLONE (TimeFinder/Clone).

-techPort|-tp *port#*

Used with `-techHost` to optionally specify the port number used by the remote Solutions Enabler server daemon. The default is 2707. Use the default value.

-exclude_nodes *nodename1 nodename2 ...*

Excludes nodes that do not have access to the source and target device when in a Geo Cluster with SQL AlwaysOn replication and mirroring.

-force|-forc

Overrides error when the target is in use. On Solaris, `-force` also overrides label validation error after source label is copied to target.

-cluster [*mscd|mscsv*]

The `mscd` option is used to setup a cluster migration for cluster disk. The `mscsv` option is used to setup migration for Cluster Shared Volume. Running `-cluster` without `mscd` or `mscsv` is equivalent to `-cluster mscd`.

-noPrompt|-no

Suppresses the default confirmation that appears when you run a command.

Setup

```
# powermig setup -src emcpowerc -tgt emcpowerh -tt hostcopy
Setup migration? [Yes]/No: yes
Migration Handle = 1
```

The following example uses the `-file` option to move all migration pairs listed in the `/tmp/pairs` file to the setup state.

```
# powermig setup -techType HOSTCOPY -file /tmp/pairs -noprompt
Source = emcpower21c, Target = emcpower4c
Migration Handle = 244
Source = emcpower20c, Target = emcpower19c
Migration Handle = 245
Source = emcpower7c, Target = emcpower5c
Migration Handle = 246
Source = emcpower6c, Target = emcpower26c
```

```
Migration Handle = 247
...
```

powermig sync

Starts the synchronization of the target logical unit with the source logical unit.

Syntax

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

Description

Depending on the size of the source, the synchronization can take some time to complete.

Once the source and target logical units are synchronized the migration transitions to the sourceSelected state. You can check the status of a synchronization using the `powermig query` command. For Open Replicator, HOSTCOPY, or SYMCLONE migrations, you can change the speed at which the source-target data is synchronized using the `powermig throttle` command.

The migration is in the syncing state while data is being copied. When the synchronization completes, the migration transitions to the sourceSelected state.

The `powermig sync` command fails if the handle is invalid or if the migration is not in the setup state.

Migration state

Required migration state to run the command: setup

Migration state after the command completes: sourceSelected

Options

-handle|-hd *migrationHandle*

Number that identifies the migration session. It is assigned when you run the `powermig setup` command. You must enter the handle when running commands that reference the migration session.



NOTE: If -all and -handle are both provided then the -all option takes precedence and the -handle value is ignored.

-all

Performs sync for all migration sessions in the setup state. Migrations in other states are ignored.

-noPrompt|-no

Suppresses the default confirmation that appears when you run a command.

Sync

```
# powermig sync -handle 1
Start sync for handle 1? [Yes]/No: y
```

powermig throttle

Sets the speed at which data is synchronized between a source-and-target pair.

Syntax

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

Description

This command is supported with all technology types except INVE (Virtual encapsulation). If you run this command when using INVE (Virtual encapsulation) as the underlying technology, while the migration transitions is in the sourceSelected state, a message alerts you that the migration is in the wrong state to run this command.

Run this command while a synchronization is in progress. This command affects resource consumption differently depending on the underlying technology used.

The `powermig throttle` command can only be changed in an OR (Open Replicator) migration while state is syncing. It can also be set when running `powermig setup`.

Migration state


Required migration state to run the command: syncing, setup

Migration state after the command completes: This command does not change the migration state.

Options

`-handle|hd migrationHandle`

Number that identifies the migration session. It is assigned when you run the `powermig setup` command. You must enter the handle when running commands that reference the migration session.

 **NOTE: If `-all` and `-handle` are both provided then the `-all` option takes precedence and the `-handle` value is ignored.**

`-all`

Performs throttle for all migration sessions in the syncing state. Migrations in other states are ignored.

`-throttleValue|tv throttleValue`

Value of the throttle, which slows down or speeds up the synchronization of the source and target logical units. Values range from 0 to 9, where 0 is the fastest speed and 9 is the slowest. The default value is 5 for Open Replicator, 0 for SYMCLONE, and 2 for HostCopy migrations. The throttle allows you to change (while a synchronization is in progress) the speed at which data is synchronized. The default values are: 5 for Open Replicator; 2 for HostCopy; and 0 for SYMCLONE.

This option is valid only with Open Replicator, SYMCLONE, or HostCopy migrations. With Open Replicator, the throttle value determines resource consumption on the target array. With HostCopy, throttle value affects resource consumption on the host.

If you need a faster copy, lower it, keeping in mind this will utilize more host resources. If you need to free up your host's resources increase the throttle value to slow down HostCopy. If you need to postpone utilizing host resources use `powermig pause`, then `powermig resume` when host resources are more readily available for copying.

`-noPrompt|no`

Suppresses the default confirmation that appears when you run a command.

powermig undoRedirect

Stops the redirection of I/O from the source to the target.

Syntax

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

Description

Use this command only when the source logical unit is a native-named device.

Before using this command, bring down any application sending I/O to the target. While the application is down, run `powermig undoRedirect`, and then reconfigure the application to use the target-device name. This prevents application I/O errors.

Migration state

Required migration state to run the command: committedAndRedirected

Migration state after the command completes: committed

Options

-handle|-hd *migrationHandle*

Number that identifies the migration session. It is assigned when you run the `powermig setup` command. You must enter the handle when running commands that reference the migration session.

NOTE: If `-all` and `-handle` are both provided then the `-all` option takes precedence and the `-handle` value is ignored.

-all

Performs `undoRedirect` for all migration sessions in the `committedAndRedirected` state. Migrations in other states are ignored.

-force

Forces the redirection of I/O from the source to the target to stop even if the source device is in use. If you use `-force` while I/O is running to the source device, then I/O errors or disk corruption may occur.

-noPrompt|-no

Suppresses the default confirmation that appears when you run a command.

undoRedirect

```
# powermig undoRedirect -handle 1
Undo IO redirection for handle x? [Yes]/No:
```

powermig version

Displays the version of PowerPath Migration Enabler running on the host.

Syntax

```
powermig version
```

Description

PowerPath Migration Enabler has the same version as the PowerPath software package with which it is distributed.

Migration state

Required migration state to run the command: Any state.

Migration state after the command completes: This command does not change the migration state.

Options

None.

powermt commands

This chapter contains the following topics:

Topics:

- [powermt command](#)

powermt command

The `powermt` command monitors and manages the PowerPath environment.

powermt check

Checks specified paths and, if desired, removes from the PowerPath configuration any paths marked dead.

Syntax

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

Description

If the PowerPath license is invalid, the `powermt check` command issues a warning and continues its processing. Use the `powermt check_registration` command to determine the problem with the PowerPath license.

Options

hba=*hba#*|all

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

dev=*path*|*device*|all

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

class={*all*|*symm*|*vx*|*clariion*|*vplex*|*invista*|*netapp*|*hitachi*|*hpxp*|*ess*|*xtremio*|*unity*|*sc*|*powermax*|*generic*|*powerstore*}

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

force

Suppresses the confirmation prompt and silently removes all paths marked dead (or, if *reconfig* is specified, silently reconfigures all mismatched paths). Do not use the *force* option in production environments.

reconfig

Supported only on Solaris and Windows. Reconfigures mismatched or misconfigured paths. Path mismatching occurs more frequently on iSCSI networks, where path login order can cause the condition. This option does not remove paths that are truly dead (use `powermt check` for this).

Reconfiguring mismatched paths (Solaris and Windows only)

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

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

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

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

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

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

powermt check_registration

Checks the state of the PowerPath license on a host.

Syntax

```
powermt check_registration
```

Description

Check registration determines whether the host has:

- A valid permanent license or evaluation license
- An expired evaluation license. Evaluation licenses provide the same functionality as permanent licenses, but an evaluation license expires on the last day of a specified month.
- An invalid or corrupt license
- No license

Apply one license per server. When different devices use different licenses, the `powermt set` command with the `dev=all` option may apply the most restrictive license to all devices.

You cannot run certain `powermt` commands if the license is a Base license or is expired, or if there is no license (except on a host connected exclusively to CLARiiON AX-series arrays). With an unlicensed version of PowerPath or with a PowerPath Base license you cannot run these `powermt set` commands: `powermt set policy` (with third-party arrays) or `powermt set mode`.

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

Certain versions of PowerPath provide full functionality with or without a license when the host is connected exclusively to CLARiiON AX-series storage systems. Note that models earlier than the CLARiiON AX4-5 (for example, AX100 and AX150) are not supported with PowerPath for AIX or HP-UX.

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

Options

None.

Check registration

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

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

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

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

For an evaluation license, output similar to the following appears:

```
Expires: <month>, <year>
```

For an expired license, output similar to the following appears:

```
Expired: <month>, <year>
```

powermt config

Configure paths to logical devices.

Syntax

```
powermt config
```

Description

The config command performs the following tasks:

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

Config does not remove previously configured paths when they become dead paths.

After config completes, you can change the configuration with the following `powermt` commands:

- `powermt load`
- `powermt set autostandby`
- `powermt set autostandby iopflimit`
- `powermt set autostandby agingperiod`
- `powermt set auto_host_registration`
- `powermt set mode`
- `powermt set path_latency_monitor`
- `powermt set path_latency_threshold`
- `powermt set path_retry_timeout`
- `powermt set perfmon`
- `powermt set periodic_autorestore`
- `powermt set policy`
- `powermt set port_disable`
- `powermt set reactive_autorestore`
- `powermt set streamio_threshold`
- `powermt set dev_inuse_to_array_report`
- `powermt set device_perf_to_array_report`
- `powermt set port_mode`

If the PowerPath license is invalid, `powermt config` issues a warning and continues its processing. Use `powermt check_registration` to determine the problem with the PowerPath license.

On AIX:

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

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

On Linux, before running `powermt config`, ensure that the required special native SCSI disk block device files exist in the `/dev` directory.

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

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

Options

None.

powermt disable hba

Disables all paths configured through the specified HBA.

Syntax

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

Description

All I/O on the port is stopped, and devices configured through the port are closed. This command is supported with the following platform configurations:

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

Run `powermt disable` when, for example, you want to stop I/O to an HBA that needs to be replaced. When you disable an HBA port, all paths configured through the HBA are labeled `Dead` in the `powermt display dev=device` output. To check whether an HBA is enabled or disabled, use the `powermt display hba_mode` command. In the `powermt display hba_mode` output, `Disabled` appears in the `Mode` column.

Note the following limitations:

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

Options

hba=hba#

Disables the specified HBA port using the HBA ID on the host, for example, 1280.

powermt display

Displays information about HBAs or devices configured for and managed by PowerPath.

Syntax

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

Description

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

The `powermt display dev` command includes an Array failover mode field, which displays the failover mode setting for Unity/VNX/CLARiiON arrays. This setting determines the failover mode for a server's HBA ports, and is configured through the Unity/VNX/CLARiiON Navisphere software.

Options

dev=device|all

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

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

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

every=#seconds

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

wide|width=#columns

NOTE: The wide option is not supported on Windows platforms.

Wide sets the line length to the width of the window and width sets a range within which the `powermt display` output appears. The default is 80. The acceptable range is 80 to 160 columns. The output widens or narrows within the range depending on the width of the window. This value applies only to the specific output; it is not persistent. You can set a persistent value for the `width` option by using the environment variable `PP_DISPLAY_WIDTH`. If you set the `every` option along with the `width` option, the `powermt display` output widens or narrows to accommodate changes you make to the window width.

Display

```
# powermt display dev=all verbose
Pseudo name=emcpowera
XtremIO ID=JWXED130500271 [PPXtremIO4]
Logical device ID=514F0C59BF0005E4 [LND198060_1]
state=alive; policy=LeastBlocks; queued-I/Os=0
=====
----- Host ----- - Stor - -- I/O Path -- -- Stats ---
### HW Path I/O Paths Interf. Mode State Q-I/Os Errors
=====
  1 lpfc      sdq      X1-C1-f1 active  alive    0      0
  0 lpfc      sde      X1-C1-f2 active  alive    0      0
  1 lpfc      sdu      X1-C2-f2 active  alive    0      0
  0 lpfc      sda      X1-C2-f1 active  alive    0      0
```

Display output for NVMe

```
# powermt display dev=all
Pseudo name=emcpowera
PowerMax ID=000197600236
Logical device ID=000019c0
Device WWN=600009700BC723EC36ED0036000006A8
```

```
state=alive; policy=ADaptive; queued-IOS=0
=====
----- Host ----- - Stor - -- I/O Path -- -- Stats ---
### HW Path          I/O Paths  Interf.  Mode   State  Q-IOS Errors
=====
  4 qla2xxx           nvme7n2  FN  4g:25 active  alive  0    0
  4 qla2xxx           nvme3n2  FN  2g:25 active  alive  0    0
  4 qla2xxx           nvme6n2  FN  2g:24 active  alive  0    0
  4 qla2xxx           nvme2n2  FN  4g:24 active  alive  0    0
  0 qla2xxx           nvme4n2  FN  4g:24 active  alive  0    0
  0 qla2xxx           nvme5n2  FN  2g:25 active  alive  0    0
  0 qla2xxx           nvme0n2  FN  2g:24 active  alive  0    0
  0 qla2xxx           nvme1n2  FN  4g:25 active  alive  0    0
```

```
Pseudo name=emcpowere0
PowerStore ID=FNM00185000725
Logical device ID=68CCF098000936E5FD3407C240D6E3D9
state=alive; policy=ADaptive; queued-IOS=0
```

```
=====
----- Host ----- - Stor - -- I/O Path -- -- Stats ---
### HW Path          I/O Paths  Interf.  Mode   State  Q-IOS Errors
=====
 18 qla2xxx           sdyc      A0-NB-f3 active  alive  0    0
 15 qla2xxx           sdcu      A0-NA-f2 active  alive  0    0
 15 qla2xxx           sdlm      A0-NA-f3 active  alive  0    0
 18 qla2xxx           sdll      A0-NA-f2 active  alive  0    0
 15 qla2xxx           sdgc      A0-NB-f2 active  alive  0    0
 15 qla2xxx           sdhs      A0-NB-f3 active  alive  0    0
```

The following example is from a PowerPath for Windows host.

```
# powermt display dev=all verbose
Pseudo name=harddisk15
Symmetrix ID=000195700879
Logical device ID=1D9A
Device WWN=60000970000197200863533033334345
state=alive; policy=SymmOpt; queued-IOS=1
=====
----- Host ----- - Stor - -- I/O Path -- -- Stats ---
### HW Path          I/O Paths  Interf.  Mode   State  Q-IOS Errors
=====
 5 port5\path0\tgt3\lun8 c5t3d8 FA  4d:37 active  alive  0    1
 5 port5\path0\tgt0\lun8 c5t0d8 FA  7d:37 active  alive  0    1
 6 port6\path0\tgt1\lun8 c6t1d8 FA  4d:37 active  alive  0    1
 6 port6\path0\tgt0\lun8 c6t0d8 FA  7d:37 active  alive  0    1
 5 port5\path0\tgt3\lun8 c5t3d8 FA  8e:00 active  dead   0    1
 5 port5\path0\tgt0\lun8 c5t0d8 FA  7e:00 active  dead   0    1
 6 port6\path0\tgt1\lun8 c6t1d8 FA  8e:00 active  dead   0    1
 6 port6\path0\tgt0\lun8 c6t0d8 FA  7e:00 active  dead   0    1
```

In case of NDM, after successful NDM creates operation, the new paths to VMAX3 array gets created.

Perform NDM from VMAX2 to VMAX3 by referring to migration steps. Ensure that `powermt display dev=devicename` is as follows during various stages of NDM:

1. Before NDM with only VMAX2 devices present in host:

```
# powermt display dev=emcpowerbb
Pseudo name=emcpowerbb
Symmetrix ID=000195700879
Logical device ID=1D9B
Device WWN=60000970000195700879533031443942
state=alive; policy=SymmOpt; queued-IOS=0
=====
----- Host ----- - Stor - -- I/O Path -- -- Stats ---
### HW Path          I/O Paths  Interf.  Mode   State  Q-IOS Errors
=====
 3 qla2xxx           sdge      FA  8e:00 active  alive  0    0
 3 qla2xxx           sdbz      FA  7e:00 active  alive  0    0
 0 qla2xxx           sdac      FA  8e:00 active  alive  0    0
 0 qla2xxx           sdh       FA  7e:00 active  alive  0    0
```

2. After establishing NDM session:

```
# powermt display dev=emcpowerbb
Pseudo name=emcpowerbb
Symmetrix ID=000197200863, 000195700879
Logical device ID=00004266, 1D9B
Device WWN=60000970000195700879533031443942
state=alive; policy=SymmOpt; queued-IOs=0
=====
----- Host ----- - Stor - -- I/O Path -- -- Stats ---
### HW Path I/O Paths Interf. Mode State Q-IOs Errors
=====
  3 qla2xxx sdcn FA 1e:08 active alive 0 0
  3 qla2xxx sdbq FA 5e:08 active alive 0 0
  0 qla2xxx sdbb FA 1e:08 active alive 0 0
  0 qla2xxx sda FA 5e:08 active alive 0 0
  3 qla2xxx sdge FA 8e:00 active alive 0 0
  3 qla2xxx sdbz FA 7e:00 active alive 0 0
  0 qla2xxx sdac FA 8e:00 active alive 0 0
  0 qla2xxx sdh FA 7e:00 active alive 0 0
```

3. After NDM CUTOVER:

```
# powermt display dev=emcpowerbb
Pseudo name=emcpowerbb
Symmetrix ID=000197200863
Logical device ID=00004266
Device WWN=60000970000195700879533031443942
state=alive; policy=SymmOpt; queued-IOs=0
=====
----- Host ----- - Stor - -- I/O Path -- -- Stats ---
### HW Path I/O Paths Interf. Mode State Q-IOs Errors
=====
  3 qla2xxx sdcn FA 1e:08 active alive 0 0
  3 qla2xxx sdbq FA 5e:08 active alive 0 0
  0 qla2xxx sdbb FA 1e:08 active alive 0 0
  0 qla2xxx sda FA 5e:08 active alive 0 0
```

i **NOTE:** Ensure that for NDM device, device ID and serial number of VMAX3 are displayed, and not the VMAX2 details.

The following example includes an SRDF Metro device:

```
# powermt display dev=emcpower4
Pseudo name=emcpower4
Symmetrix ID=000197200863, 000195700879
Logical device ID=00004213, 1A0A
Device WWN=60000970000195700879533031413041
Standard UID=naa.60000970000195700879533031413041
type=Conventional; state=alive; policy=SymmOpt; queued-IOs=0
=====
----- Host ----- - Stor - -- I/O Path -- -- Stats ---
### HW Path I/O Paths Interf. Mode State Q-IOs Errors
=====
  1 vmhba3 C0:T2:L4 FA 1e:08 active alive 0 0
  1 vmhba3 C0:T3:L4 FA 5e:08 active alive 0 0
  2 vmhba2 C0:T1:L4 FA 1e:08 active alive 0 0
  2 vmhba2 C0:T3:L4 FA 5e:08 active alive 0 0
  2 vmhba2 C0:T0:L4 FA 7e:00 active alive 0 0
  1 vmhba3 C0:T0:L4 FA 8e:00 active alive 0 0
  2 vmhba2 C0:T2:L4 FA 8e:00 active alive 0 0
  1 vmhba3 C0:T1:L4 FA 7e:00 active alive 0 0
```

If autostandby proximity is on for Symmetrix, `asb:prox` paths are displayed for SRDF Metro devices.

powermt display latency

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

Syntax

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

Description

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

- Current is the time, in microseconds, it took the last I/O on the path to complete.
- Max is the 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 `powermt set path_latency_threshold` command. To turn on path latency monitoring, use the `powermt set path_latency_monitor` command.

The `powermt display options` command shows if the functionality is enabled or disabled

Options

dev=*device***|all**

Displays latency information for the specified device or for all devices.


class={all|symm|vnx|clarion|vplex|invista|netapp|hitachi|hpxp|ess|xtremio|unity|sc|powermax|generic|powerstore}

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

every=*#seconds*

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

wide|width=*#columns*

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

Wide sets the line length to the width of the window and width sets a range within which the `powermt display` output appears. The default is 80. The acceptable range is 80 to 160 columns. The output widens or narrows within the range depending on the width of the window. This value applies only to the specific output; it is not persistent. You can set a persistent value for the `width` option by using the environment variable `PP_DISPLAY_WIDTH`. If you set the `every` option along with the `width` option, the `powermt display` output widens or narrows to accommodate changes you make to the window width.

Display latency

```
powermt display latency dev=harddisk5
Pseudo name=harddisk5
CLARiiON ID=FNM00022300718 [Storage Group 1]
Logical device ID=6006016038E02400BCCBAF9428AADE1 [LUN 8]
state=alive; policy=CLAROpt; ; queued-IOs=0
Owner: default=SP A, current=SP A                      Array failover mode: 4
=====
----- Host ----- -Stor - I/O Path - - Latency (us)
### HW Path          I/O Paths          Interf.          Mode
State          Current Max
=====
5 port5\path0\tgt0\lun2      c5t0d2      SP A0      active      alive 0 0
5 port5\path0\tgt1\lun2      c5t1d2      SP B1      active      alive 0 0
6 port6\path0\tgt0\lun2      c6t0d2      SP A1      active      alive 0 0
6 port6\path0\tgt1\lun2      c6t1d2      SP B0      active      alive 0 0
7 port7\path0\tgt0\lun2      c7t0d2      SP A0      active      alive 0 0
7 port7\path0\tgt1\lun2      c7t1d2      SP B1      active      alive 0 0
8 port8\path0\tgt0\lun2      c8t0d2      SP A1      active      alive 0 0
8 port8\path0\tgt1\lun2      c8t1d2      SP B0      active      alive 0 0
```


powermt display alua

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

Syntax

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

Description

The valid asymmetric access states are: Active/optimized, Active/non-optimized, Standby, Unavailable, Offline, Transitioning

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

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

Options

dev=device|all

Specifies a valid ALUA device or all ALUA devices.


class={all|symm|vnx|clarion|vplex|invista|netapp|hitachi|hpxp|ess|xtremio|unity|sc|powermax|generic|powerstore}

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

every=#seconds

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

wide|width=#columns

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

Wide sets the line length to the width of the window and width sets a range within which the powermt display output appears. The default is 80. The acceptable range is 80 to 160 columns. The output widens or narrows within the range depending on the width of the window. This value applies only to the specific output; it is not persistent. You can set a persistent value for the width option by using the environment variable PP_DISPLAY_WIDTH. If you set the every option along with the width option, the powermt display output widens or narrows to accommodate changes you make to the window width.

Display ALUA

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

# powermt display alua dev=8 class=clarion every=2
Pseudo name=emcpower8a
CLARiion ID=APM00090700078 [lclb039]
Logical device ID=600601601C102400AC335B99B796E011 [Thin LUN 166]
state=alive; policy=CLAROpt; queued-IOS=0
```

```

Owner: default=SP A, current=SP A          Array failover mode: 4
=====
---- Host ---- - Stor - ----- I/O Path ----- -- Stats ---
### I/O Paths  Interf.  ALUA State          Mode      State      Errors
=====
3072 c2t5006016041E0A22Ed9s0 SP A0      Active/optimized      active  alive  0
3072 c2t5006016841E0A22Ed9s0 SP B0      Active/non-optimized  active  alive  0
3074 c3t5006016141E0A22Ed9s0 SP A1      Active/optimized      active  alive  0
3074 c3t5006016941E0A22Ed9s0 SP B1      Active/non-optimized  active  alive  0

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

```

<repeated every two seconds>

powermt display array_serial

Displays array serial numbers for all the paths to a volume.

Syntax

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

Description

The `powermt display array_serial` command helps to identify paths of distributed volumes from multi-frame appliances such as, VPLEX Metro clusters and VMAX SRDF Active Active configurations.

With the array serial numbers displayed in all the paths, the paths going to the local and the remote array are easily identifiable.

Options

dev=*device***|all**

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

class={*all|symm|vnx|clarion|vplex|invista|netapp|hitachi|hpxp|ess|xtremio|unity|sc|powermax|generic|powerstore*}

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

Display array_serial

```

Pseudo name=emcpowera
Symmetrix ID=000197600771, 000197600545
Logical device ID=00000043, 0000176B
Device WWN=60000970000197600545533031373642
state=alive; policy=SymmOpt; queued-IOs=0
=====
---- Host ---- - Stor - -- I/O Path -- --- Array Serial ---
### HW Path      I/O Paths  Interf.  Mode      State      Number
=====
  3 lpfc          sdav      FA  2d:08  active  alive  000197600771
  3 lpfc          sdap      FA  4d:08  active  alive  000197600545
  3 lpfc          sdab      FA  1d:08  active  alive  000197600545
  3 lpfc          sdaj      FA  1d:08  active  alive  000197600771
  2 lpfc          sdt       FA  1d:08  active  alive  000197600771
  2 lpfc          sdn       FA  2d:08  active  alive  000197600771
  2 lpfc          sdh       FA  4d:08  active  alive  000197600545

```

powermt display bus

Displays all paths between the HBAs on a host and the storage-system ports that the HBAs can access.

Syntax

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

Description

The output has one line per path. It displays information one line per bus. Alternatively, you can type `display paths` to display path information.

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

Options


class={all|symm|vnx|clarion|vplex|invista|netapp|hitachi|hpxp|ess|xtremio|unity|sc|powermax|generic|powerstore}

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

every=#seconds

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

wide|width=#columns

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

Wide sets the line length to the width of the window and width sets a range within which the `powermt display` output appears. The default is 80. The acceptable range is 80 to 160 columns. The output widens or narrows within the range depending on the width of the window. This value applies only to the specific output; it is not persistent. You can set a persistent value for the `width` option by using the environment variable `PP_DISPLAY_WIDTH`. If you set the `every` option along with the `width` option, the `powermt display` output widens or narrows to accommodate changes you make to the window width.

Display bus

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

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

Ess logical device count=10
=====
```

```

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

```

Hitachi logical device count=9

```

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

```

```

~]# powermt display bus
PowerStore logical device count=5

```

```

----- Host Bus Adapters ----- Storage System ----- - I/O Paths -
### HW Path                               ID                               Interface                               Total                               Dead
=====
 15 qla2xxx                               FNM00185000725                   A0-NB-f2                               5                               0
 15 qla2xxx                               FNM00185000725                   A0-NA-f2                               5                               0
 15 qla2xxx                               FNM00185000725                   A0-NA-f3                               5                               0
 15 qla2xxx                               FNM00185000725                   A0-NB-f3                               5                               0
 18 qla2xxx                               FNM00185000725                   A0-NB-f3                               5                               0
 18 qla2xxx                               FNM00185000725                   A0-NA-f2                               5                               0

```

PowerMax logical device count=158

```

----- Host Bus Adapters ----- Storage System ----- - I/O Paths -
### HW Path                               ID                               Interface                               Total                               Dead
=====
 16 lpfc                                   000197600236                       FN 4g:25                               158                              0
 16 lpfc                                   000197600236                       FN 4g:24                               158                              0
 16 lpfc                                   000197600236                       FN 2g:25                               158                              0
 16 lpfc                                   000197600236                       FN 2g:24                               158                              0
 17 lpfc                                   000197600236                       FN 4g:25                               158                              0
 17 lpfc                                   000197600236                       FN 4g:24                               158                              0
 17 lpfc                                   000197600236                       FN 2g:25                               158                              0
 17 lpfc                                   000197600236                       FN 2g:24                               158                              0

```

powermt display bus-ports san-latency

PowerPath reports SAN latency values for paths to PowerMax arrays running microcode 5978.444.444 and above.

Syntax

```
powermt display bus-ports <san-latency|speed|wn> [class=<class>|all]
```

Description

As the vendor-specific commands used to calculate these values are not queued on the array port, their response is not dependent on how loaded the array is and hence accurately reflects the actual latency between the host initiator and array ports. SAN administrators can analyze these values to determine issues in the network that may be causing higher latencies and fix them.

Following latency values are determined automatically every 30 minutes for each host initiator to array target pair and are logged:

- Previous Latency – Latency value reported from the previous iteration.
- Current Latency – Most recently reported Latency value.
- Max Latency – Max latency reported in the last 24 hours. This value gets reset every 24 hours.

Options

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

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

Output for Windows platform

```
powermt display bus-ports san-latency

Storage class = Symmetrix
=====
-- Host Bus Adapters --- Storage System ---      SAN Latency (ns)
### HW Path      ID      Interface  Previous  Current  Max
=====
 4 port4\path0    000197600771  FA  2d:08    157200    185200    185200
 4 port4\path0    000197600771  FA  1d:08    172200    184900    184900
 4 port4\path0    000197600236  FA  2d:06    202100    224300    224300
 4 port4\path0    000197600236  FA  1d:06    159600    214500    214500
 5 port5\path0    000197600771  FA  2d:08    417200    434600    434600
 5 port5\path0    000197600771  FA  1d:08    391900    317
```

powermt display dead

Lists the volumes which have dead paths and also the total number of dead paths. Supported only on Linux.

Syntax

```
powermt_display_dead
```

Description

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

Options

None.

powermt display hba_mode

Displays information about all HBAs configured on the host.

Syntax

```
powermt display hba_mode
```

Description

The displayed information includes the HBA mode, which is either enabled or disabled. When the mode is enabled, the HBA is active and working as it normally would within a PowerPath environment. When the mode is disabled, it means that the HBA was disabled previously using the `powermt disable` command.

Options

None.

Display hba_mode

```
# powermt display hba_mode
Symmetrix logical device count=65
VPLEX logical device count=20
VNX logical device count=10
XtremIO logical device count=19
Unity logical device count=51
PowerStore logical device count=5
PowerMax logical device count=158
```

```

=====
----- Host Bus Adapters -----      ----- I/O Paths -----      Stats
###  HW Path                          Summary    Total    Dead    Q-I/Os  Mode
=====
 15  qla2xxx                            optimal    436     0       0       Enabled
 16  lpfc                                optimal    667     0       0       Enabled
 17  lpfc                                optimal    667     0       0       Enabled
 18  qla2xxx                            optimal    426     0       0       Enabled
 19  192.168.11.74                       optimal    10      0       0       Enabled
 20  192.168.10.78                       optimal    10      0       0       Enabled
 21  192.168.11.74                       optimal    10      0       0       Enabled
 22  192.168.10.78                       optimal    10      0       0       Enabled
 23  192.168.11.74                       optimal    10      0       0       Enabled
 24  192.168.11.74                       optimal    10      0       0       Enabled
 25  192.168.10.78                       optimal    10      0       0       Enabled
 26  192.168.11.74                       optimal    10      0       0       Enabled
 27  192.168.11.74                       optimal    10      0       0       Enabled
 28  192.168.10.78                       optimal    10      0       0       Enabled
 29  192.168.10.78                       optimal    10      0       0       Enabled
 30  192.168.10.78                       optimal    10      0       0       Enabled
#

```

powermt display nonvirtual

Displays, along with the standard powermt display output, information on the physical devices associated with Virtual Volumes.

Syntax

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

Description

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

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

Refer to the EMC Invista documentation for more information on setting up Invista devices for Migration Enabler.

Options

dev=<device>|all

Displays information about the specified device or all devices.

class={all|symm|vnx|clarion|vplex|invista|netapp|hitachi|hp|ess|xtremio|unity|sc|powermax|generic|powerstore}

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

every=<#seconds>

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

wide|width=<#columns>

NOTE: The wide option is not supported on Windows platforms.

Wide sets the line length to the width of the window and width sets a range within which the powermt display output appears. The default is 80. The acceptable range is 80 to 160 columns. The output widens or narrows within the range depending on the width of the window. This value applies only to the specific output; it is not persistent. You can set a persistent value for the `width` option by using the environment variable `PP_DISPLAY_WIDTH`. If you set the `every` option along with the `width` option, the `powermt display` output widens or narrows to accommodate changes you make to the window width.

Display nonvirtual

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

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

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

<output truncated>

powermt display options

Displays PowerPath options for each storage system class.

Syntax

```
powermt display options
```

Description

The display options command displays the following settings for each storage system class:

- The user-assignable LUN names of Unity, VNX, and CLARiiON.
- The status of path latency monitoring and the path latency threshold, if one is set.
- The status of periodic autorestore and reactive autorestore
- The status of autostandby (array class based) and the threshold, if one is set. The possible values for proximity-based autostandby are on, off, and N/A. For classes where proximity-based autostandby are on, the threshold value is default or a user-defined value.
- Performance data collection settings
- Array specific feature settings
- Status of the storage system class (managed, unmanaged, manage_incomplete, or unmanage_incomplete)

To change the settings, use the appropriate `powermt set` or `powermt load` commands to change the default storage system class. Unity, VNX, and CLARiiON user-assignable LUN names are displayed by default.

Options

None.

Display options

```
# powermt display options

Show CLARiiON LUN names:      true

Path Latency Monitor: Off

Performance Monitor: disabled

Autostandby:  IOs per Failure (iopf): enabled
               iopf aging period      : 1 d
               iopf limit              : 6000

Storage
System Class  Attributes
-----
Symmetrix     periodic autorestore = on
               reactive autorestore = on
               status = managed
               proximity based autostandby = off
               path retry timeout = 5
               all paths dead retry timeout = 45
               auto host registration = enabled
               app finger printing = disabled
               device to array performance report = enabled
               device in use to array report = enabled

CLARiiON     periodic autorestore = on
               reactive autorestore = on
               status = managed
               proximity based autostandby = N/A
               path retry timeout = 5
               all paths dead retry timeout = 45
               auto host registration = disabled
               app finger printing = N/A
               device to array performance report = N/A
               device in use to array report = N/A

Invista      periodic autorestore = on
               reactive autorestore = on
               status = managed
               proximity based autostandby = N/A
               path retry timeout = 5
               all paths dead retry timeout = 45
               auto host registration = N/A
               app finger printing = N/A
               device to array performance report = N/A
               device in use to array report = N/A

Hitachi      periodic autorestore = on
               reactive autorestore = on
               status = managed
               proximity based autostandby = N/A
               path retry timeout = 5
               all paths dead retry timeout = 45
               auto host registration = N/A
               app finger printing = N/A
               device to array performance report = N/A
               device in use to array report = N/A

HP xp        periodic autorestore = on
               reactive autorestore = on
               status = managed
               proximity based autostandby = N/A
               path retry timeout = 5
               all paths dead retry timeout = 45
               auto host registration = N/A
               app finger printing = N/A
               device to array performance report = N/A
               device in use to array report = N/A
```



```

Ess      periodic autorestore = on
        reactive autorestore = on
        status = managed
        proximity based autostandby = N/A
        path retry timeout = 5
        all paths dead retry timeout = 45
        auto host registration = N/A
        app finger printing = N/A
        device to array performance report = N/A
        device in use to array report = N/A

VPLEX   periodic autorestore = on
        reactive autorestore = on
        status = managed
        proximity based autostandby = on
        proximity based autostandby threshold = default
        path retry timeout = 5
        all paths dead retry timeout = 45
        auto host registration = N/A
        app finger printing = N/A
        device to array performance report = N/A
        device in use to array report = N/A

VNX     periodic autorestore = on
        reactive autorestore = on
        status = managed
        proximity based autostandby = N/A
        path retry timeout = 5
        all paths dead retry timeout = 45
        auto host registration = disabled
        app finger printing = N/A
        device to array performance report = N/A
        device in use to array report = N/A

        NetApp      periodic autorestore = on
        reactive autorestore = on
        status = managed
        proximity based autostandby = N/A
        path retry timeout = 5
        all paths dead retry timeout = 45
        auto host registration = N/A
        app finger printing = N/A
        device to array performance report = N/A
        device in use to array report = N/A

XtremIO periodic autorestore = on
        reactive autorestore = on
        status = managed
        proximity based autostandby = N/A
        path retry timeout = 5
        all paths dead retry timeout = 45
        auto host registration = N/A
        app finger printing = N/A
        device to array performance report = N/A
        device in use to array report = N/A

Unity   periodic autorestore = on
        reactive autorestore = on
        status = managed
        proximity based autostandby = N/A
        path retry timeout = 5
        all paths dead retry timeout = 45
        auto host registration = disabled
        app finger printing = N/A
        device to array performance report = N/A
        device in use to array report = N/A

SC      periodic autorestore = on
        reactive autorestore = on
        status = managed
        proximity based autostandby = N/A
        path retry timeout = 5
        all paths dead retry timeout = 45

```

```

auto host registration = N/A
app finger printing = N/A
device to array performance report = N/A
device in use to array report = N/A

PowerStore
periodic autorestore = on
reactive autorestore = on
status = managed
proximity based autostandby = N/A
path retry timeout = 5
all paths dead retry timeout = 45
auto host registration = N/A
app finger printing = N/A
device to array performance report = N/A
device in use to array report = N/A

PowerMax
periodic autorestore = on
reactive autorestore = on
status = managed
proximity based autostandby = N/A
path retry timeout = 5
all paths dead retry timeout = 45
auto host registration = N/A
app finger printing = N/A
device to array performance report = N/A
device in use to array report = N/A

#

```

In case of NDM, after successful NDM create operation, the new paths to VMAX3 array gets created.

powermt display perf bus

Displays path performance metrics for each bus.

Syntax

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

Description

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

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

Options

continuous

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

verbose

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

xml

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

nowait

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

Display perf dev bus

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

Symmetrix logical device count=418

```
=====
----- Host Bus Adapters ----- Storage System ----- Retry Error
### HW Path ID Interface delta delta
=====
```

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

PowerStore logical device count=5

```
=====
----- Host Bus Adapters ----- Storage System ----- Retry Error
### HW Path ID Interface delta delta
=====
```

###	HW Path	ID	Storage Interface	System	Retry delta	Error delta
15	qla2xxx	FNM00185000725	A0-NA-f2		0	0
				Metrics---Reads--Writes		
				Low Latency (ms) - -		
				High Latency (ms) - -		
15	qla2xxx	FNM00185000725	A0-NA-f3		0	0
				Metrics---Reads--Writes		
				Low Latency (ms) - -		
				High Latency (ms) - -		
15	qla2xxx	FNM00185000725	A0-NB-f2		0	0
				Metrics---Reads--Writes		
				Low Latency (ms) - -		
				High Latency (ms) - -		
15	qla2xxx	FNM00185000725	A0-NB-f3		0	0
				Metrics---Reads--Writes		
				Low Latency (ms) - -		
				High Latency (ms) - -		
18	qla2xxx	FNM00185000725	A0-NB-f3		0	0
				Metrics---Reads--Writes		
				Low Latency (ms) - -		

```

High Latency (ms) - -
18 qla2xxx
      FNM00185000725 A0-NA-f2 0 0
      Metrics----Reads--Writes
      Low Latency (ms) - -
      High Latency (ms) - -
+-----+
PowerMax logical device count=158
=====
----- Host Bus Adapters ----- Storage System ----- Retry Error
### HW Path ID Interface delta delta
=====
16 lpfc
      000197600236 FN 4g:25 0 0
      Metrics----Reads--Writes
      Low Latency (ms) - -
      High Latency (ms) - -
16 lpfc
      000197600236 FN 4g:24 0 0
      Metrics----Reads--Writes
      Low Latency (ms) - -
      High Latency (ms) - -
16 lpfc
      000197600236 FN 2g:25 0 0
      Metrics----Reads--Writes
      Low Latency (ms) - -
      High Latency (ms) - -
16 lpfc
      000197600236 FN 2g:24 0 0
      Metrics----Reads--Writes
      Low Latency (ms) - -
      High Latency (ms) - -
17 lpfc
      000197600236 FN 4g:25 0 0
      Metrics----Reads--Writes
      Low Latency (ms) - -
      High Latency (ms) - -
17 lpfc
      000197600236 FN 4g:24 0 0
      Metrics----Reads--Writes
      Low Latency (ms) - -
      High Latency (ms) - -
17 lpfc
      000197600236 FN 2g:25 0 0
      Metrics----Reads--Writes
      Low Latency (ms) - -
      High Latency (ms) - -
17 lpfc
      000197600236 FN 2g:24 0 0
      Metrics----Reads--Writes
      Low Latency (ms) - -
      High Latency (ms) - -

```

Table 9. Attributes and description

Attributes	Description
Latency - Reads	Latency value (ms) for Read IOs.
Latency - Writes	Latency value (ms) for Write IOs.
Latency - Low	Lowest latency value (ms) for IOs.

Table 9. Attributes and description (continued)

Attributes	Description
Latency - High	Highest latency value (ms) for IOs.
Lat <= 1ms	Number of IOs with Latency value less than or equal to 1 millisecond.
1ms < Lat <= 10ms	Number of IOs with Latency value greater than 1 millisecond but less than or equal to 10 milliseconds.
10ms < Lat <= 100ms	Number of IOs with Latency value greater than 10 milliseconds but less than or equal to 100 milliseconds.
Lat > 100ms	Number of IOs with Latency value greater than 100 milliseconds.

powermt display perf dev

Display storage device performance metrics for all logical devices.

Syntax

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

Description

The output has one line per path. There is a table for each type of storage device. The storage device count above each table includes the type of storage device to which that table applies.

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

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

Options

dev=*device*|all

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

continuous

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

verbose

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

xml

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

nowait

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

Display perf dev

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

state=alive; policy=SymmOpt; queued-IOs=10
Read bytes/s      KB<=4K  4K<KB<=8K  8K<KB<=128K  KB>128K  All
                  4.50M    3.25M      13.1M         -         20.8M
```

```

Write bytes/s      4.50M    3.26M    13.1M    -    20.8M
Total bytes/s
Read Avg Response ms 1.03    0.985    1.01    -    1.02
Write Avg Response ms 1.05    0.989    1.02    -    1.04
All Avg Response ms
Read IOs/s        2.80K    396      796     -    3.99K
Write IOs/s       2.79K    397      798     -    3.99K
Total IOs/s

```

```

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

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

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

```

```

# powermt display perf dev=es verbose
Timestamp = 07:19:46 UTC, 14 Apr 2020
Sample Interval = 60
Pseudo name=emcpoweres
PowerStore ID=FNM00185000725
Logical device ID=68CCF09800B1D8C093C6A05CF54BBAFB
state=alive; policy=ADaptive; 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
=====
 15 qla2xxx          sdcx
      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 - -

 15 qla2xxx          sdgf
      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	-	-		
15 qla2xxx	sdhv			0	0
	Low Latency (ms)	-	-		
	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	-	-		
15 qla2xxx	sdlr			0	0
	Low Latency (ms)	-	-		
	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	-	-		
18 qla2xxx	sdlr			0	0
	Low Latency (ms)	-	-		
	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	-	-		
18 qla2xxx	sdvf			0	0
	Low Latency (ms)	-	-		
	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	-	-		

Pseudo name=emcpowerla
PowerMax ID=000197600236
Logical device ID=600009700BC723EC36ED002800000B05
state=alive; policy=ADaptive; 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
=====
16 lpfc              nvme0n33

```

	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	-	-		
16 lpfc	nvme1n33				
	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	-	-		
16 lpfc	nvme2n33				
	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	-	-		
16 lpfc	nvme3n33				
	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	-	-		
17 lpfc	nvme4n33				
	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	-	-		
17 lpfc	nvme5n33				
	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	-	-		
17 lpfc	nvme6n33				
	Low Latency (ms)	-	-	0	0
	High Latency (ms)	-	-		
	Lat <= 1ms (#)	0	0		
	1ms < Lat <= 10ms (#)	0	0		
	10ms < Lat <= 100ms (#)	0	0		

Table 10. Attributes and description

Attributes	Description
KB<=4K	IOs less than or equal to 4KB in size.
4K<KB<=8K	IOs greater than 4KB but less than or equal to 8KB in size.

Table 10. Attributes and description (continued)

Attributes	Description
8K<KB<=128K	IOs greater than 8KB but less than or equal to 128KB in size.
KB>128K	IOs greater than 128KB in size.
All	IOs of all sizes.
Read bytes/s	IO throughput (bytes/s) for Read IOs.
Write bytes/s	IO throughput (bytes/s) for Write IOs.
Total bytes/s	IO throughput (bytes/s) for All IOs.
Read Avg Response	Average response time (ms) for Read IOs.
Write Avg Response	Average response time (ms) for Write IOs.
All Avg Response	Average response time (ms) for All IOs.
Read IOs/s	IO count (IOs /s) for Read IOs.
Write IOs/s	IO count (IOs /s) for Write IOs.
Total IOs/s	IO count (IOs /s) for All IOs.
Latency - Reads	Latency value (ms) for Read IOs.
Latency - Writes	Latency value (ms) for Write IOs.
Latency - Low	Lowest latency value (ms) for IOs.
Latency - High	Highest latency value (ms) for IOs.
Lat <= 1ms	Number of IOs with Latency value less than or equal to 1 millisecond.
1ms < Lat <= 10ms	Number of IOs with Latency value greater than 1 millisecond but less than or equal to 10 milliseconds.
10ms < Lat <= 100ms	Number of IOs with Latency value greater than 10 milliseconds but less than or equal to 100 milliseconds.
Lat > 100ms	Number of IOs with Latency value greater than 100 milliseconds.

powermt display port_mode

Displays the mode of all ports on all storage systems configured on the host.

Syntax

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

Description

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

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

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

Options

dev=device

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

class={all|symm|vnx|clarion|vplex|invista|netapp|hitachi|hpxp|ess|xtremio|unity|sc|powermax|generic|powerstore}

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

Display port mode

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

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

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

powermt display ports

Displays attached array port specific information, one line per port.

Syntax

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

Description

The `display ports` command displays the state of the storage-system ports that are visible to the HBAs under PowerPath control. The output has one line per port and includes the Array failover mode field, which displays the failover mode setting for Unity/VNX/CLARiiON arrays. This setting determines the failover mode for a server's HBA ports, and is configured through the Unity/VNX/CLARiiON Navisphere software.

Options

dev=*device*|all

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

class={all|symm|vnx|clarion|vplex|invista|netapp|hitachi|hpxp|ess|xtremio|unity|sc|powermax|generic|powerstore}

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

every=#seconds

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

wide|width=#columns

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

Wide sets the line length to the width of the window and width sets a range within which the powermt display output appears. The default is 80. The acceptable range is 80 to 160 columns. The output widens or narrows within the range depending on the width of the window. This value applies only to the specific output; it is not persistent. You can set a persistent value for the `width` option by using the environment variable `PP_DISPLAY_WIDTH`. If you set the `every` option along with the `width` option, the `powermt display` output widens or narrows to accommodate changes you make to the window width.

Display ports

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

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

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

```
Storage class = PowerStore
=====
----- Storage System ----- -- I/O Paths -- --- Stats ---
ID          Interface  Wt_Q  Total  Dead  Q-IOs  Errors
=====
FNM00185000725 A0-NB-f3   256   10    0     0     0
FNM00185000725 A0-NB-f2   256    5    0     0     0
FNM00185000725 A0-NA-f2   256   10    0     0     0
FNM00185000725 A0-NA-f3   256    5    0     0     0

Storage class = PowerMax
=====
----- Storage System ----- -- I/O Paths -- --- Stats ---
ID          Interface  Wt_Q  Total  Dead  Q-IOs  Errors
=====
000197600236 FN  4g:25   256   316    0     0     0
000197600236 FN  4g:24   256   316    0     0     0
000197600236 FN  2g:25   256   316    0     0     0
000197600236 FN  2g:24   256   316    0     0     0
```

```
# powermt display ports dev=emcpowera
Pseudo name=emcpowera
Symmetrix ID=000195700363
```

```

Logical device ID=7743
state=alive; policy=SymmOpt; queued-IOs=0; write_throttle=off
=====
----- Host ----- - Storage Sys - - Path - -- Stats --
### HW Path I/O Paths Interf. Wt_Q Status Q-IOs Errors
=====
 4 lpfc sdat FA 8gB 256 a,v 0 0
 4 lpfc sdbg FA 9gB 256 a,v 0 0
 5 lpfc sdcz FA 8gB 256 a,v 0 0
 5 lpfc sddm FA 9gB 256 a,v 0 0

```

In case of NDM, after successful NDM create operation, the new paths to VMAX3 array gets created.

powermt display unmanaged

Displays information about logical devices that have been excluded from PowerPath management through the `powermt unmanage` command.

Syntax

```
powermt display unmanaged
```

Description

You can use the native device as the `dev` option to the `powermt manage dev` command.

Options

None.

Display unmanaged

```

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

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

```

Devices unmanaged with `powermt unmanage class=class` will not be listed.

powermt enable hba

Enables an HBA port that was previously disabled with the `powermt disable` command.

Syntax

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

Description

Running `powermt enable` resumes I/O to the specified HBA port and its devices. This command is supported with the following platform configurations:

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

This command fails if the PowerPath load balancing and failover policy is `bf` or `nr`.

Options

`hba=hba#`

Enables the specified HBA port, which resumes I/O to the HBA and its devices.

powermt help

Displays usage summary for the `powermt` command.

Syntax

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

Description

The usage syntax is also displayed when invalid commands or options are entered.

Options

`command`

If a specific command is provided, displays the usage only of the command specified.

Help

```
powermt help display

display - Display multipathing information.

Usage:
  powermt <command>
           [class={all|symm|vnx|clariion|vplex|invista|netapp|hitachi|hpxp|ess|xtremio|unity|
sc|powermax|generic|powerstore}
           ]

powermt display [dev=<device>|all [verbose]] [class=<class>|all] [every=<#seconds>]
               [width=<#col>]
powermt display alua dev=<device>|all [class=<class>|all] [every=<#seconds>]
               [width=<#col>]
powermt display latency [dev=<device>|all] [class=<class>|all]
               [every=<#seconds>] [width=<#col>]
powermt display nonvirtual dev=<device>|all
               [every=<#seconds>] [class=<class>|all] [width=<#col>]
powermt display options
powermt display bus [class=<class>|all] [every=<#seconds>]
               [width=<#col>]
powermt display perf dev=<device>|all [continuous] [verbose] [xml] [nowait]

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

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

powermt load

Loads the current configuration the following settings from a saved configuration file.

Syntax

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

Description

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

If the PowerPath license is invalid, `powermt load` issues a warning and continues its processing. Use `powermt check_registration` to determine the problem with the license.

On HP-UX, Solaris, and Windows

The `powermt load` command also causes devices to be configured:

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

On Linux

After you run `powermt load`, you must run `powermt config` to configure pseudo devices; otherwise, only native devices are configured.

On Solaris

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

Options

file=filename

This is an optional option. The system will save or load the default configuration file if one is not specified. The default configuration file name and location must be used during boot. If you specify a filename and it is not in your current directory, then you must include the full file path as well as the filename in the file specification. For example, `powermt load file=/etc/emc/powermt_custom.xml`

powermt manage

Places a specified logical device or storage system class under PowerPath management.

Syntax

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

Description

By default, all storage system classes are under PowerPath management. You can use `powermt unmanage` to exclude a storage system class from PowerPath management. Use `powermt manage class` to return that storage system class to PowerPath control.

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

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

Options

dev=path|device

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

NOTE: `powermt manage` cannot accept a pseudo device as an option, since pseudo devices are not assigned to unmanaged devices/classes. Note, however, that this restriction does not apply to the `powermt unmanage` command as pseudo device names are assigned to managed devices/classes.

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

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

powermt release

Releases PowerPath pseudo device (`emcpower`) names for later reuse.

Syntax

```
powermt release
```

Description

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

Options

None.

powermt remove

Removes the specified path (or paths) from the PowerPath list of configured paths.

Syntax

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

Options

`hba=hba#|all`

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

`dev=path|device|all`

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

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

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

`force`

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

powermt restore

Tests and restores the specified paths.

Syntax

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

Description

Restore issues test I/Os and responds to the test results as follows:

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

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

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

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

If the PowerPath license is invalid, `powermt restore` issues a warning and continues its processing. To determine the problem with the PowerPath license, use `powermt check_registration`.

Options

`hba=hba#|all`

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

`dev=path|device|all`

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

`class={all|symm|vnx|clarion|vplex|invista|netapp|hitachi|hpxp|ess|xtremio|unity|sc|powermax|generic|powerstore}`

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

powermt save

Saves the PowerPath settings in the PowerPath configuration file.

Syntax

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

Description

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

Run `powermt save` to update your PowerPath configuration whenever you make changes to hardware on the host. If you do not capture these changes in the `powermt` configuration file and then need to restore the system, the file that is restored will be incorrect. These settings save important PowerPath driver information required for the PowerPath state persistence across host reboots.

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

The configuration file is saved when `powermt save` is run manually and, on some platforms, at boot time. A saved configuration can be reloaded with `powermt load`.

Do not manually modify the configuration file.

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

On Windows

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

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

The Windows PowerPath Administrator GUI prompts you to save the file after configuration changes are made in the GUI.

Options

`file=filename`

This is an optional option. The system will save or load the default configuration file if one is not specified. The default configuration file name and location must be used during boot. If you specify a filename and it is not in your current directory, then you must include the full file path as well as the filename in the file specification. For example, `powermt load file=/etc/emc/powermt_custom.xml`

powermt set all_paths_dead_retry_timeout

Sets the amount of time an I/O will be retried after PowerPath has determined that all paths have become dead.

Syntax

```
powermt set all_paths_dead_retry_timeout=<#seconds>[class=<class>|all]
```

Description

This parameter specifies the amount of time an I/O will be retried after PowerPath has determined that all paths have become dead. The I/O will hang till this timeout expires after all the paths have become dead. If at least one of the paths to the LUN comes back alive within this period, the I/O will be dispatched on an alive path and can succeed. If none of the paths restore back, then an I/O failure will be returned to the application after this timeout period expires.

- During Symmetrix microcode upgrade, all paths to a Symmetrix array may go down simultaneously for a brief amount of time before they come back alive. Hence, to prevent I/O errors, ensure that the all paths dead retry timeout value is set to at least 30 seconds before a NDU is done on a Symmetrix array. After the NDU, the value can be changed again.
- This parameter can also be used to fasten the failover from one cluster node to another when all its paths go down. However, if the nodes are connected to a Symmetrix array and if NDU is going to be performed, then the timeout value needs to be reset to at least 30 seconds for the duration of the NDU to avoid application failures.

The `powermt display options` command shows the all paths dead retry timeout value. When you change all paths dead retry timeout and then run `powermt save`, this setting persists across reboot.

Options

`all_paths_dead_retry_timeout=#seconds`

Seconds is a value between 1-120 seconds that determines the amount of time that I/O will be retried for after all paths have become dead. By default, on a fresh installation, PowerPath will set this timeout value to 45 seconds.

`class={all|symm|vnx|clarion|vplex|invista|netapp|hitachi|hpxp|ess|xtremio|unity|sc|powermax|generic|powerstore}`

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

All paths dead retry timeout

The following command sets the retry timeout to 120 seconds on all devices:

```
# powermt set all_paths_dead_retry_timeout=120 class=all
```

powermt set auto_host_registration

Enables or disables auto host registration.

Syntax

```
powermt set auto_host_registration={on|off} class=<class>
```

Description

PowerPath Auto Host Registration is a feature where host sends special registration I/O commands down each path to the array to register the paths. Usually the registration is performed by an array agent like the Navisphere or Unisphere agent running on a host. The information sent in the registration command includes the initiator type, hostname, IP address, etc. By enabling PowerPath to send the registration information across the SAN to arrays, host information can be bound to the array automatically as soon as the host is visible. It simplifies the configuration process.

Use `powermt display options` to verify if the `auto_host_registration` is on before trying `powermt update host_registration`.

NOTE: If `powermt update host_registration` is tried when the feature is disabled, error message is not displayed. Virtual Machine details are not captured in Windows during Auto Host Registration.

Auto Host registration for PowerMaxOS

This feature is available for PowerMaxOS 5978 and above. PowerPath sends host details such as hostname, OS version, server hardware vendor, PowerPath license WWN details, host connectivity type, and PowerPath version as part of host registration payload to PowerMax. The host registration information is sent to PowerMax as soon as Host reboots or when user triggers `powermt update host_registration` command manually. Host registration information is sent automatically every 24 hours. If any change in host configuration, host registration is sent within 1 hour by PowerPath . PowerPath also sends Cluster details in Windows environment to array.

NOTE: Auto host registration is enabled by default for symm arrays.

Auto host registration for Unity, VNX, and CLARiiON

This feature is carried forward from PowerPath previous releases for VNX, Unity, and CLARiiON. Host sends the information that includes initiator type, hostname, and IP address as part of host registration payload. This registration is sent only once on host start or PowerPath starts.

NOTE: Auto host registration is disabled by default for CLARiiON, VNX and Unity.

Auto host registration is persistent across reboot. The `powermt load` and `powermt save` commands can be used to load and save the configuration.

The `powermt display options` command shows if the functionality is enabled or disabled.

Options

`auto_host_registration=on`

Enables auto host registration.

`auto_host_registration=off`

Disables auto host registration. By default, auto host registration is disabled.

powermt set autostandby

Enables or disables autostandby.

Syntax

```
powermt set autostandby={on|off|reinitialize} [trigger={prox [class=<class> [threshold=<value>]]|iopf}]
```

Description

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

- Proximity-based autostandby (asb:prox) applies to path modes for multi-frame storage appliances, such as VMAX SRDF/Metro Arrays, and VPLEX cluster configuration where a host is connected to and manages a distributed volume across multiple VPLEX clusters and frames within a given VPLEX system. The proximity-based autostandby for VMAX SRDF/Metro is disabled by default to avoid issues with low-latency remote sites.

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

The default for the prox trigger is off. While autostandby is enabled or if autostandby is reinitialized with the prox trigger, PowerPath performs a test on all paths, sets the appropriate paths to active or autostandby, and determines the preferred paths to a volume from federated storage.

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

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

When you turn autostandby off, and then run `rpowermt save`, then this setting persists across reboot. The autostandby modes for each path do not persist across reboot. For example, if path set A is in autostandby mode and path set B is in active mode, then on reboot path set A can change to active mode and path set B can change to autostandby mode. On each reboot (assuming the global autostandby setting is on), the system performs latency measurements and re-runs the algorithms for selecting which path or path set to put in autostandby mode and which path or path set to put in active mode.

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

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

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

Options

autostandby=on

Enables autostandby. If you do not select a trigger, both prox and iopf are enabled. For example, `rpowermt set autostandby=on` globally turns on autostandby with both the proximity-based and the IOsPerFailure-based autostandby features.

autostandby=off

Disables autostandby. PowerPath clears the global setting for enabling autostandby (or the selected type of autostandby), and persists that setting if you subsequently run the `rpowermt save` command. PowerPath also changes all paths in autostandby (or the selected type of autostandby) to active mode. If you do not select a trigger, then both prox and iopf are enabled. For example, `rpowermt set autostandby=off` globally turns off autostandby with both the proximity-based and the IOsPerFailure-based autostandby features.

autostandby=reinitialize

If the proximity trigger is enabled, this option recalibrates the trigger and determines anew the preferred paths to VPLEX distributed volumes — this can be useful after a cluster has been physically moved, for example. If autostandby is not enabled, then an error message displays. If you do not select a trigger, then prox is enabled. For example, `rpowermt set autostandby=reinitialize` reinitializes autostandby with the proximity-based autostandby feature. The minimum latencies for each cluster and the mode changes for each path are logged in audit messages.

trigger=prox [class=class [threshold=value]]

Selects only the proximity-based autostandby feature. This determines and selects the preferred paths to a volume from federated storage and places the non-preferred 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.

The `threshold` value, in microseconds, can range from 0 to 5000 (rounded to multiples of 50 after reboot). The threshold can be set only on a per-class basis.

Threshold value is used to decide paths of which Frame belonging to a multi-frame storage appliance. If the latency difference between two frames is greater than the user-defined value, then the frame with higher latency is marked as non-preferred and other as preferred. If the latency difference is lesser than the user-defined threshold value then the paths to both the frames are marked as active.

i **NOTE: This setting is applicable only for VPLEX and Symmetrix devices. For VPLEX devices, this setting is enabled by default with a threshold value of 0. For Symmetrix devices, this setting is not enabled by default and needs to be set manually.**

`trigger=iopf`

Selects only the IOsPerFailure-based autostandby feature. This determines whether a path that has intermittent I/O failures should be placed into autostandby mode (asb:iopf).

powermt set autostandby agingperiod

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

Syntax

```
powermt set autostandby agingperiod={<#days>|<#hours>h|<#hours>H}
```

Description

When a path is placed into autostandby mode through the IOsPerFailure trigger, PowerPath adds the aging amount to the current time. The aging period timeout allows PowerPath to return the paths to full use without user intervention. Error monitoring is resumed when the paths are returned to an active state. Each path ages separately. This may result in each path having different times at which they return to active mode.

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

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

Options

`agingperiod={#days|#hoursh|#hoursH}`

Specifies the number of days or hours after which a path in autostandby mode due to intermittent I/O failures automatically resets to active mode. The default number of days is 1. The range of acceptable values is 0 to 24854 days or 596520 hours. Use H or h to specify hours (for example, 6H or 6h for 6 hours). If you pass a value without H or h then it is considered as number of days.

Specifying 0 or 0h turns off the periodic reset of a path in autostandby due to intermittent I/O failures. These paths in autostandby (asb:iopf) never return to active mode without manual intervention.

Autostandby agingperiod

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

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

The following command resets flaky paths (asb:iopf paths) to active after 6 hours:

```
powermt set autostandby agingperiod=6H
```

powermt set autostandby iopflimit

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

Syntax

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

Description

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

Options

iopflimit=*value*

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

powermt set device_perf_to_array_report

Enables device performance statistics reporting to the array.

Syntax

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

Description

User has to specify class option to enable or disable device performance reporting for an array.

This feature is available for PowerMaxOS 5978 microcode and above. PowerPath collects statistics for number of reads/writes, IO bandwidth of reads/writes, and latency of reads/writes present in Unisphere. For more information see Unisphere documentation.

PowerPath sends performance statistics to PowerMax once in a minute for a maximum of 400 devices. This feature is persistent across host restart.

Options

device_perf_to_array_report={on|off}

Enables (on) or disables (off) device performance statistics reporting to the array.

class={all|netapp|symm|vnx|clariion|vplex|invista|hitachi|hp|x|ess|xtremio|unity|sc|generic|powerstore}

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

powermt set dev_inuse_to_array_report

Enables device report to the array.

Syntax

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

Description

`powermt set dev_inuse_to_array_report` enables or disables device report to the array. User has to specify class option to enable or disable this feature for an array. This feature is available for PowerMaxOS 5978 microcode and above. PowerPath sends the

application name which is using the device. PowerPath also sends if the device is mounted on the host and the last time the device was in use. This information is sent to array once in a day. This feature is persistent across host restart.

Options

dev_inuse_to_array_report={on|off}

`on` enables device report to the array.

`off` disables device report to the array.

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

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

powermt set mode

Sets paths to active or standby mode.

Syntax

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

Description

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

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

The `powermt set mode=<active|standby> force` command will log only an audit log message. There will not be any message transition from `iopf` to active or from active to `iopf`.

Options

mode=active

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

mode=standby

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

hba=hba#|all

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

dev=path|device|all

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

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

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

force

Manually overrides autostandby. If a path is in autostandby mode, you must use the `force` option to change it to active or standby. If the path is in autostandby and the `force` option is not specified, then an error is returned.

powermt set path_latency_monitor

Enables or disables path latency monitoring.

Syntax

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

Description

When you enable or disable path latency monitoring, the setting applies globally to all paths.

NOTE: Before enabling latency monitoring for the first time, you should not only read this section but also `powermt set path_latency_threshold` for setting a threshold for path latency.

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

The `powermt display options` command shows if the functionality is enabled or disabled

To avoid extraneous log messages, set a threshold value before turning on latency monitoring.

Options

`path_latency_monitor=on`

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

`path_latency_monitor=off`

Disables I/O latency measurement for all paths, and resets all paths' Current and Max latencies to zero. The default setting is off. Disabling latency monitoring does not change a previously set threshold. On re-enabling, the last threshold entered goes back into effect.

powermt set path_latency_threshold

Sets a time interval in seconds within which I/Os should complete.

Syntax

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

Description

When a threshold has been set, PowerPath generates system log messages indicating each threshold crossing that results in a new Max latency (high watermark) for a path. This threshold value applies to all paths in your environment.

For example, if

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

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

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

When changing the threshold, use the following three steps:

1. Disable path latency monitoring so that high watermarks are zeroed.

2. Set the new threshold.
3. Re-enable path latency monitoring.

Options

`path_latency_threshold=#seconds|#milliseconds`

Sets a time interval in seconds within which I/Os should complete. The value applies to all paths. The default value is 0. The range of acceptable values is 0 to 3600 seconds or 0 to 3600000 milliseconds. Seconds is the default.

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

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

powermt set path_retry_timeout

Sets the amount of time I/O will be held in the lower layers on a path failure.

Syntax

```
powermt set path_retry_timeout={default|<#seconds>} [class=<class>|all]
```

Description

By default, on a fresh installation, PowerPath sets this value to 5 seconds.

Path retry timeout reduces I/O completion time during path failures by changing the timeout values in the SCSI layer.

The `powermt display options` command shows the path retry timeout value. When you set path retry timeout and then run `powermt save`, this setting persists across reboot.

Options

`path_retry_timeout={default|#seconds}`

`default` does not change the SCSI timeout values and hence I/O hang times are directly dependent on the protocol and the type of HBA used. `Seconds` is a value between 1-25 seconds that determines the amount of time that I/O hangs before I/O is retried on a path failure. On PowerPath installation, the default value is 5 seconds. If this timeout value is changed, it may take up to five minutes for the change to take effect.

`class={all|symm|vnx|clarion|vplex|invista|netapp|hitachi|hpxp|ess|xtremio|unity|sc|powermax|generic|powerstore}`

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

Path retry limit

The following command sets the retry timeout to 20 seconds on all devices:

```
# powermt set path_retry_timeout=20 class=all
```

powermt set perfmon

Enables or disables performance monitoring for all devices.

Syntax

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


Description

Performance monitoring helps characterize I/O patterns and possibly aide in diagnosing I/O problems.

When performance monitoring is enabled, `powermt display perf dev=all` and `powermt display perf bus` displays performance metrics for all paths to all devices and bus ports.

The `powermt display options` command shows if the functionality is enabled or disabled

Options

`perfmon=on`

Enables performance monitoring for all devices. This command initializes the counters, clears all performance measurements including high and low watermarks for latency, and accepts a sampling interval for data collection in seconds with a default interval of 900 seconds (15 minutes). The legacy path latency monitoring threshold measurements are not affected. There is no output.

`perfmon=off`

Disables performance monitoring and clears all measurements.

`interval=#seconds`

The sampling interval in seconds for data collection. The sampling interval range is from a minimum of 60 seconds (1 minute) to a maximum of 86,400 seconds (1,440 minutes or 1 day) with a default of 900 seconds (15 minutes).

To change the performance monitoring sampling interval, you must first set `perfmon=off` and then set `perfmon=on`.

`powermt set periodic_autorestore`

Enables or disables the periodic autorestore feature.

Syntax

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

Description

Periodic autorestore can be used along with reactive autorestore to define when paths are restored after path failure. To determine the periodic autorestore setting, use `powermt display options`.

Options

`periodic_autorestore=on`

Enables periodic autorestore. PowerPath periodically tests dead paths and, if they pass the test, restores them to service. The default is `on`.

`periodic_autorestore=off`

Disables periodic autorestore. Even when periodic autorestore is off, path testing continues to be done under certain conditions as described in the *PowerPath Family Product Guide*, and automatic path restoration continues to occur based on the results of that testing.

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

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

`powermt set persist_path_info`

Controls the `powermt save` behavior with respect to path information.

Syntax

```
powermt set persist_path_info={on|off}
```

Description

After you set the value of `persist_path_info`, run `powermt save` to persist this value in the custom file.

This command is available only to the PowerPath Windows platform. PowerPath UNIX platforms need Devices path information to be persisted always as Ghost volume is an important feature to these platforms.

The `powermt display options` command shows if the functionality is enabled or disabled. This shows the current value saved in Kernel and not the one in `powermt` custom file.

Options

`persist_path_info={on|off}`

Enables (on) or disables (off) path persistence information. The default is on.

NOTE: If the HBA SCSI target persistent binding is disabled, it is recommended to set `persist_path_info` to off. Without HBA persistent binding, the SCSI ID of the path persisted in the custom file may change after actions like a HBA enable/disable or a server reboot. This may cause `powermt` to display misleading information. To enable SCSI target persistent binding, refer to the HBA management software documentation.

powermt set policy

Sets the load-balancing and failover policy for devices configured in PowerPath.

Syntax

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

Description

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

Options

`policy={ad|co|lb|li|re|rr|si|so}`

Sets the load-balancing and failover policy to the specified value.

`dev=device|all`

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

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

`class={all|symm|vnx|clarion|vplex|invista|netapp|hitachi|hpxp|ess|xtremio|unity|sc|powermax|generic|powerstore}`

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

powermt set port_disable

Enables or disables a storage system port, and closes all devices configured through the port.

Syntax

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

Description

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

Note the following limitations:

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

Options

`port_disable=on`

Disables a port on the specified device. Once disabled, all paths configured through a port are marked dead and no I/O is sent to the port.

`port_disable=off`

Enables a port on the specified device. The default is off.

`dev=device`

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

Port disable

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

```
# powermt set port_disable=on dev=c5t50060482CAFD7A5Dd1s0
# powermt display port_mode dev=c5t50060482CAFD7A5Dd1s0
Pseudo name=emcpower27a
Symmetrix ID=000187430377
Logical device ID=0001
state=alive; policy=SymmOpt; queued-IOs=0
=====
----- Host ----- Stor-- I/O Path-- Stats-
### HW Path          I/O Paths          Interf.  Mode  State Q-IOs  Errors
=====
3072 pci@1c,600000/SUNW,q1c@1/fp@0,0 c3t5006048ACAFD7A7Cd1s0 FA 13dB active alive 0 0
3074 pci@1d,700000/SUNW,q1c@1/fp@0,0 c5t50060482CAFD7A5Dd1s0 FA 14bA active dead 0 0
```

powermt set port_mode

Sets the mode of all paths of a port on any storage system

Syntax

```
powermt set port_mode={active|standby} dev=<path> [force]
```

Description

All paths on an active port will function normally. When all paths on a port are moved to standby using this command, the paths to the port are only used when there are no other viable paths for IO. There is no output.

Run the `powermt display` command to see the status of the paths.

Options

`port_mode=active`

User I/O is delivered to a path. If you designate a path as active, it is continuously scheduled for I/O according to the load-balancing and failover policy in use. An active port operates as it normally would within a PowerPath environment.

`port_mode=standby`

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

dev=path

The device specified must be a native device.

force

Sets the paths irrespective of their current mode. Otherwise, it sets or generates warning message based on their current mode.

Port mode

The command sets all paths belonging to the port corresponding to a device 'sdaa' as standby.

```
bash-4.0# powermt display dev=hdisk604
Pseudo name=hdiskpower166
Symmetrix ID=000196800207
Logical device ID=09A9
Device WWN=60000970000196800207533030394139
state=alive; policy=SymmOpt; queued-I/Os=0
=====
----- Host ----- - Stor - -- I/O Path -- -- Stats ---
### HW Path I/O Paths Interf. Mode State Q-I/Os Errors
=====
  1 fcs11 hdisk1301 FA 2e:04 active alive 0 0
  1 fcs11 hdisk1274 FA 4e:04 active alive 0 0
  0 fcs10 hdisk631 FA 2e:04 active alive 0 0
  0 fcs10 hdisk604 FA 4e:04 active alive 0 0
bash-4.0# powermt set port mode=standby dev=hdisk604
bash-4.0# powermt display dev=hdisk604
Pseudo name=hdiskpower166
Symmetrix ID=000196800207
Logical device ID=09A9
Device WWN=60000970000196800207533030394139
state=alive; policy=SymmOpt; queuedI/Os=0
=====
----- Host ----- - Stor - -- I/O Path -- -- Stats ---
### HW Path I/O Paths Interf. Mode State Q-I/Os Errors
=====
  1 fcs11 hdisk1301 FA 2e:04 active alive 0 0
  1 fcs11 hdisk1274 FA 4e:04 standby alive 0 0
  0 fcs10 hdisk631 FA 2e:04 active alive 0 0
  0 fcs10 hdisk604 FA 4e:04 standby alive 0 0
```

powermt set reactive_autorestore

Enables or disables the PowerPath reactive autorestore facility.

Syntax

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

Description

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. `powermt set reactive_autorestore` is not supported on Windows.

The `powermt display options` command shows if the functionality is enabled or disabled

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. Do not set this 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 `powermt restore` at the command line.

Options

`reactive_autorestore=on`

Enables reactive autorestore, which reactively tests dead paths and, if they pass the test, restores them to service. For non-disruptive upgrades (NDU), reactive autorestore should be set to `on`. The default is `on`.

`reactive_autorestore=off`

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

`class={all|symm|vnx|clarion|vplex|invista|netapp|hitachi|hpxp|ess|xtremio|unity|sc|powermax|generic|powerstore}`

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

powermt set streamio_threshold

Sets the I/O threshold values for when to switch paths while the Stream I/O (si) load-balancing and failover policy is in effect.

Syntax

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

Description

The `streamio` value is accepted only if the load-balancing policy in effect on the device is `si`. If the PowerPath license is invalid, setting the `streamio_threshold` will generate an error message. You can check whether Stream I/O threshold is set by viewing the `powermt display dev` output.

Options

`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 logical device. `all` specifies all logical devices. The default is `all`.

`class={all|symm|vnx|clarion|vplex|invista|netapp|hitachi|hpxp|ess|xtremio|unity|sc|powermax|generic|powerstore}`

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

powermt set write_throttle

Enables or disables I/O write throttling to the specified logical devices.

Syntax

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

Description

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

This command is available only to the PowerPath AIX, HP-UX, and Solaris platforms.

Options

write_throttle={on|off}

Enables (on) or disables (off) write throttling to the specified logical device. For new logical devices, write throttling is disabled by default.

class={symm|all}

Limits the command to the specified type of storage system. `all` specifies all storage-system types eligible for write throttling. The default is `all`. If `class` is specified along with `dev`, the command is applied to logical devices that meet all specified constraints.

dev=device|all

Limits the change to the specified logical device. `all` specifies all logical devices eligible for write throttling. The default is `all`.

powermt set write_throttle_queue

Sets the write throttling queue depths for a storage-system port connected to a specified device.

Syntax

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

Description

The queue-depth setting limits the number of writes to all devices enabled for write throttling that can be outstanding (from PowerPath perspective) on the storage-system port. The queues are allocated within PowerPath, one per storage-system port.

This command is available only to the PowerPath AIX and Solaris platforms.

Options

write_throttle_queue=queue_depth

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

class={symm|all}

Limits the command to the specified type of storage system. `all` specifies all storage-system types eligible for write throttling. The default is `all`. If `class` is specified along with `dev`, the command is applied to logical devices that meet all specified constraints.

dev=path|device|all

Limits the change to the storage-system port associated with the specified path, or the port connected to the specified device. `all` specifies all storage-system ports to all devices eligible for write throttling. The default is `all`.

powermt unmanage

Excludes devices or storage system classes from PowerPath management.

Syntax

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

Description

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

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

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

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

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

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

Options

`dev=path|device`

Specifies the pathname of a logical device to be excluded from PowerPath management.

`class={symm|vnx|clarion|vplex|invista|netapp|hitachi|hpxp|less|xtremio|unity|sc|powermax|generic|powerstore}`

Specifies the storage system class to be excluded from PowerPath management. If `class` is specified along with `dev`, the command is applied to HBAs or devices that meet all specified constraints.

powermt update host_registration

Updates host registration.

Description

`powermt update host_registration` command triggers a host registration to all connected arrays that support host registration. This command can be used immediately to update an array whenever there is any change in host details.

This command always returns a success even if there are no arrays connected, if `auto_host_registration` is disabled on the host or if `auto_host_registration` is not supported/enabled on the connected arrays.

This command does not return any status when executed. Ensure that `auto_host_registration` feature is enabled before executing this command.

powermt update lun_names

Retrieves the latest XtremIO, Unity/VNX/CLARiON, Dell SC, and IBM XIV user-assignable LUN names.

Syntax

```
powermt update lun_names
```

Description

A user-assignable LUN name is a character string that a user or system manager associates with a LUN and assigns through Unisphere or Navisphere. CLARiON LUN nice names are updated automatically every hour. Use this command to manually update nice names.

Options

None.

powermt version

Displays the version of PowerPath running on the host.

Syntax

```
powermt version
```

Description

PowerPath version names use the term service pack (SP) for 6.x.x level releases (for example, 6.0 SP1 instead of 6.0.1) and patch (P) for 6.x.x.x level releases (for example, 6.0 SP1 P01). Installed hot fixes are displayed by default.

Options

None.

Version

```
# powermt version  
Dell EMC powermt for PowerPath (c) Version 6.0 build (xxx)
```


pprootdev commands

This chapter contains the following topics:

Topics:

- [pprootdev command](#)

pprootdev command

The `pprootdev` command enables and disables PowerPath functionality on a boot device.

The `pprootdev` tool:

- Enables and disables PowerPath load balancing and failover functionality on a storage system boot device.
- Enables the `bosboot` tool to function correctly with PowerPath.

pprootdev

Enables and disables PowerPath load balancing and failover functionality on a storage system boot device.

Syntax

```
pprootdev {on | off | fix | fixback | query}
```

Description

The `pprootdev` tool cannot change the state of `rootvg` on a running system. It does, however, modify ODM data that other tools use to determine what devices `rootvg` is using. For this reason, some commands report information that may appear to be incorrect if they are run after `pprootdev` and before a restart.

NOTE: If the `/tmp` directory is full and the disk containing the OS does not have free partitions available, the `pprootdev` command will fail. Ensure that there is enough space available before running the command.

Options

<code>on</code>	Enables PowerPath pseudo boot. This command changes AIX configuration rules and updates the boot image so that AIX uses <code>hdiskpower</code> devices to vary on the <code>rootvg</code> in PHASE1 the next time the system boots.
<code>off</code>	Disables PowerPath pseudo boot.
<code>fix</code>	Enables <code>bosboot</code> to function correctly after booting from a PowerPath device (until the next system boot).
<code>fixback</code>	Enables <code>varyonvg</code> of <code>rootvg</code> in a mirrored <code>rootvg</code> environment as this will transfer the PVIDs onto the <code>hdiskpower</code> device in the ODM.
<code>query</code>	Displays if PowerPath load balancing and failover functionality is enabled or disabled on the system boot device.

rpowermt commands

This chapter contains the following topics:

Topics:

- [rpowermt command](#)

rpowermt command

The `rpowermt` command monitors and manages the PowerPath/VE environment.

rpowermt check_registration

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>]
```

Options

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.

Check registration

If PowerPath/VE is already licensed with Served or Unserved, it continues to use as Served or Unserved license until it gets unregistered.

For PowerPath/VE to get licensed using FX, host should have at least one data device from VMAX3 array, which has FX software bundle is available.

NOTE: The Type field is displayed as FX in the following output.

```
PowerPath License Information:
-----
Host ID       : 5e784be0-8472-b26e-d117-b0262818f13c
Type         : FX
State        : licensed
License search path: /etc/emc:/etc/emc/licenses:/opt/EMCpower:/opt/EMCpower/licenses

License file(s): /etc/emc/*.lic
                  /etc/emc/licenses/*.lic
                  /opt/EMCpower
                  /opt/EMCpower/licenses
Warning: License file does not exist or not found.
```

If the host is licensed with served license the system generates the following output:

```
# rpowermt host=<host> check_registration
PowerPath License Information:
-----

Host ID       : 4dc16302-1265-a5c4-0394-002219287b68
Type          : served (counted)
State         : licensed
Registered To : EMC
Issue Date    : 13-feb-2015
Feature       : PowerPathMP
Feature Version: 5.4
Registering Svr: rpowermt@<rpowermt_server>-<licensing_server>
License Server : <licensing_server>
Days until expiration: 317
License search path: /root
License file(s): /root/POWERMT.LIC
```

 **NOTE:** The Type field is displayed as served in the output that is shown.

rpowermt check_served

Displays the license configuration information on the electronic license management server. This command works only if valid Server license is available for RTOOLS.

Syntax

```
rpowermt check_served
```

Options

None.

Check served

```
# rpowermt check_served

License details:
  Server:      localhost.localdomain
  Issuer:      EMC
  Feature:     PowerPathMP
  Version:     5.4
  Issued:      24-mar-2014
  Expiration:  31-dec-2015
  Notice:      ACTIVATED BY <user> for Testing
  Num_lic:     5
  Overdraft:   1
-----

License details:
  Server:      localhost.localdomain
  Issuer:      EMC
  Feature:     PowerPathMP
  Version:     5.4
  Issued:      26-jan-2014
  Expiration:  31-dec-2014
  Notice:      ACTIVATED BY <user> for Testing
  Num_lic:     10
  Overdraft:   0
-----

ELM server info:
  Licenses:    15
  Overdraft:   1
  Total count: 16
  In use:      0 (float 0)
  Available:   16
```

rpowermt display

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

Syntax

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

Description

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

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

Options

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:

- PowerPath/VE pseudo device name
- Native path identifier - Format of the native device path is, `hwpath:Cx:Ty:Lz` value

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

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

verbose

This option is only supported with `rpowermt display dev` command. This option displays the long nice name for XtremIO devices.

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=*#columns*

Width sets a range within which the `rpowermt display` output appears. The default is 80. The acceptable range is 80 to 160 columns. The output widens or narrows within the range depending on the width of the window. This value applies only to the specific output; it is not persistent. You can set a persistent value for the `width` option by using the environment variable `PP_DISPLAY_WIDTH`. If you set the `every` option along with the `width` 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.

Display

```
# rpowermt display dev=emcpower140 host=x.x.x.x
Pseudo name=emcpower140
XtremIO ID=JWXED130500271
```

```
Logical device ID=514F0C59BF0007F7
Standard UID=naa.514f0c59bf0007f7 [vmq216018_34]
type=Conventional; state=alive; policy=ADaptive; queued-IOS=0
=====
----- Host ----- - Stor - -- I/O Path -- -- Stats ---
### HW Path          I/O Paths   Interf.  Mode   State   Q-IOS Errors
=====
  3 vmhba4            C0:T6:L33 X1-C2-f1 active  alive   0     0
  3 vmhba4            C0:T5:L33 X1-C1-f2 active  alive   0     0
  2 vmhba5            C0:T6:L33 X1-C2-f2 active  alive   0     0
  2 vmhba5            C0:T5:L33 X1-C1-f1 active  alive   0     0
```

Display output for NVMe Device

```
# rpowermt display dev=emcpower31 class=symm host=x.x.x.x
Pseudo name=emcpower31
Symmetrix ID=000197600236
Logical device ID=00001962
Device WWN=600009700BC723EC36ED00360000065B
Standard UID=eui.600009700bc723ec36ed00360000065b
type=NVMe; state=alive; policy=SymmOpt; queued-IOS=0
=====
----- Host ----- - Stor - -- I/O Path -- -- Stats ---
### HW Path          I/O Paths   Interf.  Mode   State   Q-IOS Errors
=====
  4 vmhba64           C0:T0:L6497 SA 2g:24 active  alive   0     0
  3 vmhba65           C0:T0:L6497 SA 2g:24 active  alive   0     0
  4 vmhba64           C0:T1:L6497 SA 2g:25 active  alive   0     0
  3 vmhba65           C0:T1:L6497 SA 2g:25 active  alive   0     0
  4 vmhba64           C0:T2:L6497 SA 4g:24 active  alive   0     0
  3 vmhba65           C0:T2:L6497 SA 4g:24 active  alive   0     0
  4 vmhba64           C0:T3:L6497 SA 4g:25 active  alive   0     0
  3 vmhba65           C0:T3:L6497 SA 4g:25 active  alive   0     0
```

rpowermt display alua

Displays ALUA state for all paths to an ALUA device, that is, storage systems that support the SCSI-standard ALUA.

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>]
```

Description

The valid asymmetric access states are: 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.

Options

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:

- PowerPath/VE pseudo device name
- Native path identifier - Format of the native device path is, *hwpath:Cx:Ty:Lz* value

class={all|symm|vnx|clarion|vplex|invista|netapp|hitachi|hpxp|ess|xtremio|unity|sc|powermax|generic|powerstore}

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

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=#columns

Width sets a range within which the `rpowermt` display output appears. The default is 80. The acceptable range is 80 to 160 columns. The output widens or narrows within the range depending on the width of the window. This value applies only to the specific output; it is not persistent. You can set a persistent value for the `width` option by using the environment variable `PP_DISPLAY_WIDTH`. If you set the `every` option along with the `width` 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.

ALUA output for NVMe Devices

```
rpowermt display alua dev=eui.600009700bc723ec36ed00360000065b class=symm host=x.x.x.x
Pseudo name=emcpower31
Symmetrix ID=000197600236
Logical device ID=00001962
Device WWN=600009700BC723EC36ED00360000065B
Standard UID=eui.600009700bc723ec36ed00360000065b
type=NVMe; state=alive; policy=SymmOpt; queued-IOS=0
=====
---- Host ---- - Stor - ----- I/O Path ----- -- Stats ---
### I/O Paths Interf. ALUA State Mode State Errors
=====
  4 C0:T0:L6497 SA 2g:24 Active/optimized active alive 0
  3 C0:T0:L6497 SA 2g:24 Active/optimized active alive 0
  4 C0:T1:L6497 SA 2g:25 Active/optimized active alive 0
  3 C0:T1:L6497 SA 2g:25 Active/optimized active alive 0
  4 C0:T2:L6497 SA 4g:24 Active/optimized active alive 0
  3 C0:T2:L6497 SA 4g:24 Active/optimized active alive 0
  4 C0:T3:L6497 SA 4g:25 Active/optimized active alive 0
  3 C0:T3:L6497 SA 4g:25 Active/optimized active alive 0
```

rpowermt display array serial

Displays array serial numbers for all the paths to a volume.

Syntax

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

Description

The `rpowermt display array_serial` command helps to identify paths of distributed volumes from multi-frame appliances such as, VPLEX Metro clusters and VMAX SRDF Active Active configurations.

With the array serial numbers displayed in all the paths, the paths going to the local and the remote array are easily identifiable.

Options

dev=device|all

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

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

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

Display array_serial

```
Pseudo name=emcpowera
Symmetrix ID=000197600771, 000197600545
Logical device ID=00000043, 0000176B
Device WWN=6000097000019760054533031373642
state=alive; policy=SymmOpt; queued-IOS=0
=====
----- Host ----- - Stor - -- I/O Path -- --- Array Serial ---
### HW Path          I/O Paths  Interf.  Mode   State   Number
=====
  3 lpfc              sdav      FA  2d:08 active  alive  000197600771
  3 lpfc              sdap      FA  4d:08 active  alive  000197600545
  3 lpfc              sdab      FA  1d:08 active  alive  000197600545
  3 lpfc              sda      FA  1d:08 active  alive  000197600771
  2 lpfc              sdt      FA  1d:08 active  alive  000197600771
  2 lpfc              sdn      FA  2d:08 active  alive  000197600771
  2 lpfc              sdh      FA  4d:08 active  alive  000197600545
  2 lpfc              sdb      FA  1d:08 active  alive  000197600545
```

rpowermt display bus

Displays all paths between the HBAs on a host and the storage-system ports that the HBAs can access.

Syntax

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

Description

The output has one line per path. It displays information one line per bus. Alternatively, you can type display paths to display path information.

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

Options

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

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

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=#columns

Width sets a range within which the rpowermt display output appears. The default is 80. The acceptable range is 80 to 160 columns. The output widens or narrows within the range depending on the width of the window. This value applies only to the specific output; it is not persistent. You can set a persistent value for the width option by using the environment variable PP_DISPLAY_WIDTH. If you set the every option along with the width 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.

Output for VMAX 236 (SCSI and NVMe) and PowerStore array

```
# rpowermt display bus host=x.x.x.x
Symmetrix logical device count=18
=====
----- Host Bus Adapters ----- Storage System ----- - I/O Paths -
### HW Path ID Interface Total Dead
=====
  1 vmhba4          000197600236   FA 1d:06          1    0
  2 vmhba5          000197600236   FA 1d:06          1    0
  3 vmhba65         000197600236   SA 2g:24         17    0
  3 vmhba65         000197600236   SA 2g:25         17    0
  3 vmhba65         000197600236   SA 4g:24         17    0
  3 vmhba65         000197600236   SA 4g:25         17    0
  4 vmhba64         000197600236   SA 2g:24         17    0
  4 vmhba64         000197600236   SA 2g:25         17    0
  4 vmhba64         000197600236   SA 4g:24         17    0
  4 vmhba64         000197600236   SA 4g:25         17    0

PowerStore logical device count=14
=====
----- Host Bus Adapters ----- Storage System ----- - I/O Paths -
### HW Path ID Interface Total Dead
=====
  1 vmhba4          FNM00185100383  A1-NB-f3          7    0
  1 vmhba4          FNM00185100383  A1-NA-f2          7    0
  1 vmhba4          FNM00185100383  A1-NA-f3          7    0
  1 vmhba4          FNM00185100383  A1-NB-f2          7    0
  1 vmhba4          FNM00185000725  A0-NB-f2          7    0
  1 vmhba4          FNM00185000725  A0-NA-f2          7    0
  1 vmhba4          FNM00185000725  A0-NA-f3          7    0
  1 vmhba4          FNM00185000725  A0-NB-f3          7    0
  2 vmhba5          FNM00185100383  A1-NA-f3          7    0
  2 vmhba5          FNM00185100383  A1-NA-f2          7    0
  2 vmhba5          FNM00185100383  A1-NB-f2          7    0
  2 vmhba5          FNM00185100383  A1-NB-f3          7    0
  2 vmhba5          FNM00185000725  A0-NB-f3          7    0
  2 vmhba5          FNM00185000725  A0-NB-f2          7    0
  2 vmhba5          FNM00185000725  A0-NA-f3          7    0
  2 vmhba5          FNM00185000725  A0-NA-f2          7    0
```

rpowermt display latency

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

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>]
```

Description

The rpowermt display latency output includes the following two columns:

- Current is the time, in microseconds, it took the last I/O on the path to complete.
- Max is the 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 `powermt set path_latency_threshold` command. To turn on path latency monitoring, use the `powermt set path_latency_monitor` command.

The `powermt display options` command shows if the functionality is enabled or disabled

Options

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 `powermt display dev` are:

- PowerPath/VE pseudo device name
- Native path identifier - Format of the native device path is, `hwpath:Cx:Ty:Lz` value

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

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

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=#columns

Width sets a range within which the `powermt display` output appears. The default is 80. The acceptable range is 80 to 160 columns. The output widens or narrows within the range depending on the width of the window. This value applies only to the specific output; it is not persistent. You can set a persistent value for the `width` option by using the environment variable `PP_DISPLAY_WIDTH`. If you set the `every` option along with the `width` option, the `powermt 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.

powermt display options

Displays PowerPath/VE options for each storage system class.

Syntax

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

Description

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

- Whether Unity, VNX, and CLARiiON user-assignable LUN names are displayed
- Whether path latency monitoring is enabled and the path latency threshold, if one is set
- Whether periodic autorestore and reactive autorestore is enabled

- Whether autostandby is enabled (array class based) and the threshold, if one is set. The possible values for proximity-based autostandby are `on`, `off`, and `N/A`. For classes where proximity-based autostandby are `on`, the threshold value is `default` or a user-defined value.
- Performance data collection settings
- Status of the storage system class (managed, unmanaged, manage_incomplete, or unmanage_incomplete)

To change the settings, use `rpowermt set` commands to change the default storage system class. Unity, VNX, and CLARiiON user-assignable LUN names are displayed by default.

Options

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.

Display options

```
# rpowermt display options host=x.x.x.x

Show CLARiiON LUN names:      true

Path Latency Monitor: Off

Performance Monitor: disabled

Autostandby:  IOs per Failure (iopf): enabled
              iopf aging period      : 1 d
              iopf limit             : 6000

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

              proximity based autostandby = off
              auto host registration = enabled
              app finger printing = disabled
              device to array performance report = enabled
              device in use to array report = enabled

CLARiiON      periodic autorestore = on
              reactive autorestore = on

              proximity based autostandby = N/A
              auto host registration = disabled
              app finger printing = N/A
              device to array performance report = N/A
              device in use to array report = N/A

Invista       periodic autorestore = on
              reactive autorestore = on

              proximity based autostandby = N/A
              auto host registration = N/A
              app finger printing = N/A
              device to array performance report = N/A
              device in use to array report = N/A

Hitachi       periodic autorestore = on
```

```

reactive autorestore = on

proximity based autostandby = N/A
auto host registration = N/A
app finger printing = N/A
device to array performance report = N/A
device in use to array report = N/A

HP xp    periodic autorestore = on
         reactive autorestore = on

         proximity based autostandby = N/A
         auto host registration = N/A
         app finger printing = N/A
         device to array performance report = N/A
         device in use to array report = N/A

Ess      periodic autorestore = on
         reactive autorestore = on

         proximity based autostandby = N/A
         auto host registration = N/A
         app finger printing = N/A
         device to array performance report = N/A
         device in use to array report = N/A

Generic  periodic autorestore = on
         reactive autorestore = on

         proximity based autostandby = N/A
         auto host registration = N/A
         app finger printing = N/A
         device to array performance report = N/A
         device in use to array report = N/A

VPLEX    periodic autorestore = on
         reactive autorestore = on

         proximity based autostandby = on
         proximity based autostandby threshold = default
         auto host registration = N/A
         app finger printing = N/A
         device to array performance report = N/A
         device in use to array report = N/A

VNX      periodic autorestore = on
         reactive autorestore = on

         proximity based autostandby = N/A
         auto host registration = enabled
         app finger printing = N/A
         device to array performance report = N/A
         device in use to array report = N/A

NetApp   periodic autorestore = on
         reactive autorestore = on

         proximity based autostandby = N/A
         auto host registration = N/A
         app finger printing = N/A
         device to array performance report = N/A
         device in use to array report = N/A

XtremIO  periodic autorestore = on
         reactive autorestore = on

         proximity based autostandby = N/A
         auto host registration = N/A
         app finger printing = N/A
         device to array performance report = N/A
         device in use to array report = N/A

Unity    periodic autorestore = on

```

```

reactive autorestore = on

proximity based autostandby = N/A
auto host registration = disabled
app finger printing = N/A
device to array performance report = N/A
device in use to array report = N/A

PowerStore    periodic autorestore = on
               reactive autorestore = on

proximity based autostandby = N/A
auto host registration = N/A
app finger printing = N/A
device to array performance report = N/A
device in use to array report = N/A

```

rpowermt display paths

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>]
```

Options

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

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

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=#columns

Width sets a range within which the `rpowermt display` output appears. The default is 80. The acceptable range is 80 to 160 columns. The output widens or narrows within the range depending on the width of the window. This value applies only to the specific output; it is not persistent. You can set a persistent value for the `width` option by using the environment variable `PP_DISPLAY_WIDTH`. If you set the `every` option along with the `width` 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.

Display paths

```

# rpowermt display paths host=x.x.x.x
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

```

Display paths with NVMe devices

```

# rpowermt display paths host=x.x.x.x
Symmetrix logical device count=18
=====
----- Host Bus Adapters ----- ----- Storage System ----- - I/O Paths -
### HW Path          ID          Interface      Total      Dead
=====
1 vmhba4          000197600236      FA  1d:06          1          0
2 vmhba5          000197600236      FA  1d:06          1          0
3 vmhba65         000197600236      SA  2g:24          17         0
3 vmhba65         000197600236      SA  2g:25          17         0
3 vmhba65         000197600236      SA  4g:24          17         0
3 vmhba65         000197600236      SA  4g:25          17         0
4 vmhba64         000197600236      SA  2g:24          17         0
4 vmhba64         000197600236      SA  2g:25          17         0
4 vmhba64         000197600236      SA  4g:24          17         0
4 vmhba64         000197600236      SA  4g:25          17         0

PowerStore logical device count=14
=====
----- Host Bus Adapters ----- ----- Storage System ----- - I/O Paths -
### HW Path          ID          Interface      Total      Dead
=====
1 vmhba4          FNM00185100383    A1-NB-f3          7          0
1 vmhba4          FNM00185100383    A1-NA-f2          7          0
1 vmhba4          FNM00185100383    A1-NA-f3          7          0
1 vmhba4          FNM00185100383    A1-NB-f2          7          0
1 vmhba4          FNM00185000725    A0-NB-f2          7          0
1 vmhba4          FNM00185000725    A0-NA-f2          7          0
1 vmhba4          FNM00185000725    A0-NA-f3          7          0
1 vmhba4          FNM00185000725    A0-NB-f3          7          0
2 vmhba5          FNM00185100383    A1-NA-f3          7          0
2 vmhba5          FNM00185100383    A1-NA-f2          7          0
2 vmhba5          FNM00185100383    A1-NB-f2          7          0
2 vmhba5          FNM00185100383    A1-NB-f3          7          0
2 vmhba5          FNM00185000725    A0-NB-f3          7          0
2 vmhba5          FNM00185000725    A0-NB-f2          7          0
2 vmhba5          FNM00185000725    A0-NA-f3          7          0
2 vmhba5          FNM00185000725    A0-NA-f2          7          0

```

rpowermt display perf bus

Displays path performance metrics for each bus.

Syntax

```

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

```

Description

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

Options

continuous

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

verbose

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

xml

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

nowait

If this option is specified, `powermt` will not wait for the interval time to display the next updated PMI statistics. 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.

Display perf dev bus

```
# rpowermt display perf bus host=<host>
Timestamp = 12:03:35 UTC, 19 Feb 2015
Sample Interval = 60

Symmetrix logical device count=77
=====
---- Host Bus Adapters ---- --- Storage System ---   Retry   Error
###  HW Path                ID           Interface   delta   delta
=====
1  vmhba2
           000195700879  FA  9f:01           0       0
           Metrics----Reads--Writes
           Low Latency (ms)    -       -
           High Latency (ms)  -       -

1  vmhba2
           000195700879  FA 10f:01           0       0
           Metrics----Reads--Writes
           Low Latency (ms)    -       -
           High Latency (ms)  -       -

4  vmhba3
           000195700879  FA  9f:01           0       0
           Metrics----Reads--Writes
           Low Latency (ms)    -       -
           High Latency (ms)  -       -

4  vmhba3
           000195700879  FA 10f:01           0       0
           Metrics----Reads--Writes
           Low Latency (ms)    -       -
           High Latency (ms)  -       -
```

rpowermt display perf bus vmname

Display bus performance metrics for a specific VM or for all VMs for which VMPMI (performance monitoring) is enabled.

Syntax

```
rpowermt display perf bus vmname=<VM Name>|all [continuous] [nowait] host=<hostname>
[username=<username> [password=<password>|no_password]] [cim_sessionid=<cim sessionid>]
```

Options

vmname=VM Name|all

Displays performance metrics for VMPMI enabled VMs. Gathers metrics for either a specific VM or all VMPMI enabled VMs.

continuous

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

nowait

If this option is specified, `rpowermt` will not wait for the interval time to display the next updated VMPMI statistics. It returns immediately to report the remaining time interval for the next updated VMPMI 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.

Display perf bus vmname

```
# rpowermt display perf bus vmname=RHEL673VM host=<host>

VM Name = RHEL673VM
VM Identifier = 50 1d 43 03 1a 2f b0 09-08 ea 88 ca bf b9 c0 23
Timestamp = 21:58:12 UTC, 30 Sep 2015
Sample Interval = 60

+-----+

VNX logical device count=3
=====
----- Host Bus Adapters ----- Storage System -----
### HW Path ID Interface
=====
  1 vmhba2
                                APM00141314672 SP A4
                                Metrics----Reads--Writes
                                IOs/Sec 254 0
                                Bytes/Sec 12.5M 0.0
                                Avg latency 13.1 -

  1 vmhba2
                                APM00141314672 SP B4
                                Metrics----Reads--Writes
                                IOs/Sec 59 0
                                Bytes/Sec 3.00M 0.0
```

```

                Avg latency  0.000    0.000

2 vmhba3

                APM00141314672  SP A5
                Metrics----Reads--Writes
                IOs/Sec  1.42K          0
                Bytes/Sec 70.6M        0.0
                Avg latency  17.7      -

2 vmhba3

                APM00141314672  SP B5
                Metrics----Reads--Writes
                IOs/Sec  1.53K          0
                Bytes/Sec 76.0M        0.0
                Avg latency  19.2

```

rpowermt display perf dev

Display storage device performance metrics for all logical devices.

Syntax

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

Description

The output has one line per path. There is a table for each type of storage device. The storage device count above each table includes the type of storage device to which that table applies.

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

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

Options

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:

- PowerPath/VE 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, displays the single most recently calculated metrics for all storage devices.

verbose

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

xml

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

nowait

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

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.

Display perf dev

```
# rpowermt display perf dev=emcpower140 verbose host=<host>
Timestamp = 12:08:35 UTC, 19 Feb 2015
Sample Interval = 60
Pseudo name=emcpower140
XtremIO ID=JWXED130500271
Logical device ID=514F0C59BF0007F7
Standard UID=naa.514f0c59bf0007f7
type=Conventional state=alive; policy=ADaptive; queued-I/Os=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 I/Os/s	-	-	-	-	-
Write I/Os/s	-	-	-	-	-
Total I/Os/s	-	-	-	-	-

```
=====
```

###	Host	I/O Paths	Metrics	Reads	Writes	Retry delta	Error delta
2	vmhba5	C0:T5:L33	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		
			I/Os/sec (#)	-	-		
			Bytes/sec	-	-		
2	vmhba5	C0:T6:L33	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		
			I/Os/sec (#)	-	-		
			Bytes/sec	-	-		
3	vmhba4	C0:T5:L33	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		
			I/Os/sec (#)	-	-		
			Bytes/sec	-	-		
3	vmhba4	C0:T6:L33	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              -      -

```

rpowersmt display perf vmname

Provides a snapshot of the performance metrics for a VM in a sample period.

Syntax

```
rpowersmt display perf vmname=<VM Name|all> [continuous] [verbose] [nowait] host=<hostname>
[username=<username> [password=<password>|no_password]] [cim_sessionid=<cim sessionid>]
```

Description

The output has one line per path. There is a table for each type of storage device. The storage device count above each table includes the type of storage device to which that table applies.

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

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

NOTE: Two sample intervals are required for this command to display performance metrics.

Options

vmname=VM Name|all

Displays performance metrics for VMPMI enabled VMs. Gathers metrics for either a specific VM or all VMPMI enabled VMs.

continuous

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

verbose

Displays both the VM's overall performance and the performance metrics of the devices attached to the VM.

nowait

If this option is specified, `rpowersmt` will not wait for the interval time to display the next updated VMPMI statistics. It returns immediately to report the remaining time interval for the next updated VMPMI 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.

Display perf VM

```
# rpowersmt display perf vmname="RHELH12VM1" host=<>

VM Identifier = 50 2e 78 1a c4 af c9 b8-dc 52 f6 8b d3 a4 c8 78
Avg Response Time (ms) =      8.00
Total Throughput (bytes/sec) =  30.7M
Total IOPS = 7.49K
=====
                        Read      Write
=====
```

```

IOs/Sec                2          7.49K
Throughput (Bytes/Sec) 11.3K     30.7M
Avg Response time(ms)  8.04      8.00

```

The following example uses the verbose option.

```

# rpowermt display perf vmname="RHELH12VM1" host=<> verbose

VM Identifier = 50 2e 78 1a c4 af c9 b8-dc 52 f6 8b d3 a4 c8 78
Avg Response Time (ms) = 8.08
Total Throughput (bytes/sec) = 30.4M
Total IOPS = 7.42K
=====
                               Read      Write
=====
IOs/Sec                       -      7.42K
Throughput (Bytes/Sec)         -      30.4M
Avg Response time(ms)         -      8.08

Timestamp = 22:23:00 UTC, 31 Aug 2015
Sample Interval = 60
Pseudo name=emcpower5
Symmetrix ID=000195700879
Logical device ID=1771
Device WWN=60000970000195700879533031373731
Standard UID=naa.60000970000195700879533031373731
type=Conventional; state=alive; policy=SymmOpt

Read bytes/s                  -
Write bytes/s                  15.2M
Total bytes/s                  15.2M

Read IO/s                      -
Write IO/s                     3.71K
Total IO/s                     3.71K

Read Avg Response ms          -
Write Avg Response ms         8.08
All Avg Response ms           8.08

=====
----- Host ----- Metrics -----
### HW Path          I/O Paths          Reads---Writes
=====
  3 vmhba5           C0:T0:L2
                        Latency (ms)      -    8.06
                        IOs/sec (#)         0   1.85K
                        Bytes/sec          0   7.56M
  2 vmhba2           C0:T5:L2
                        Latency (ms)      -    8.10
                        IOs/sec (#)         0   1.86K
                        Bytes/sec          0   7.63M
+-----+

Timestamp = 22:23:00 UTC, 31 Aug 2015
Sample Interval = 60
Pseudo name=emcpower7
Symmetrix ID=000195700879
Logical device ID=1773
Device WWN=60000970000195700879533031373733
Standard UID=naa.60000970000195700879533031373733
type=Conventional; state=alive; policy=SymmOpt

Read bytes/s                  -
Write bytes/s                  15.2M
Total bytes/s                  15.2M

Read IO/s                      -
Write IO/s                     3.71K
Total IO/s                     3.71K

```

```

Read   Avg Response ms      -
Write  Avg Response ms      8.08
All    Avg Response ms      8.08

=====
----- Host ----- Metrics -----
### HW Path          I/O Paths          Reads---Writes
=====
   3 vmhba5          C0:T0:L4
                        Latency (ms)         -      8.10
                        IOs/sec (#)           0    1.84K
                        Bytes/sec            0    7.52M
   2 vmhba2          C0:T5:L4
                        Latency (ms)         -      8.05
                        IOs/sec (#)           0    1.87K
                        Bytes/sec            0    7.67M
+-----+

```

rpowermt display port_mode

Displays the mode of all ports on all storage systems configured on the host.

Syntax

```
rpowermt display port_mode [class=<class>|all] host=<hostname> [username=<username>
[password=<password>|no_password]] [cim_sessionid=<CIM sessionid>]
```

Description

The port mode is either active or standby. An active port is active and operates as it normally would within a PowerPath environment. When a port is standby, the path is held in reserve.

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

Options

class={all|symm|vnx|clariion|vplex|invista|netapp|hitachi|hpx|less|xtremio|unity|sc|powermax|generic|powerstore}

Limits the command to the specified type of storage system. `all` specifies all storage-system types. 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.

Display port mode

```

# rpowermt display port_mode host=<host IP>

Storage class = CLARiiON
=====
----- Storage System ----- I/O Paths ----- Stats -----
ID          Interface      Wt_Q    Total    Dead    Q-IOs    Mode
=====
APM00092800238  SP B2        256     6        0        0    Active
APM00092800238  SP A1        256     6        0        0    Active
APM00092800238  SP A0        256     6        0        0    Active
APM00092800238  SP B1        256     6        0        0    Active

```

```
Storage class = VPLEX
=====
----- Storage System ----- -- I/O Paths -- --- Stats ---
ID          Interface      Wt_Q    Total    Dead    Q-IOs  Mode
=====
FCNCH090201457  CL2-0C      256     10      0      0  Active
FCNCH092905441  CL1-04      256     10      0      0  Active
FCNCH090201457  CL2-04      256     10      0      0  Active
FCNCH092905441  CL1-0C      256     10      0      0  Active

Storage class = VNX
=====
----- Storage System ----- -- I/O Paths -- --- Stats ---
ID          Interface      Wt_Q    Total    Dead    Q-IOs  Mode
=====
APM00141314672  SP A5        256     5       0       0  Active
APM00141314672  SP B5        256     5       0       0  Active
APM00141314672  SP A4        256     5       0       0  Active
```

rpowermt display ports

Displays attached array port specific information, one line per port.

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>]
```

Description

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

Options

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:

- PowerPath/VE pseudo device name
- Native path identifier - Format of the native device path is, `hwpath:Cx:Ty:Lz` value

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

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

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=<#columns>

Width sets a range within which the `rpowermt display` output appears. The default is 80. The acceptable range is 80 to 160 columns. The output widens or narrows within the range depending on the width of the window. This value applies only to the specific output; it is not persistent. You can set a persistent value for the `width` option by using the environment variable `PP_DISPLAY_WIDTH`. If you set the `every` option along with the `width` 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.

Display ports

Storage class = Symmetrix

```
=====
----- Storage System -----
----- ID          Interface          Wt_Q          -- I/O Paths --   --- Stats ---
----- Total          Dead          Q-IOs      Errors
=====
```

ID	Storage Interface	System	Wt_Q	Total	Dead	Q-IOs	Errors
000197801413	FA 2d:04		256	54	0	1	0
000197000760	FA 1d:04		256	6	0	0	0
000197801413	FA 1d:04		256	54	0	1	0
000197801413	FA 2d:05		256	52	0	1	0
000197000760	FA 1d:05		256	6	0	0	0
000197801413	FA 1d:05		256	52	0	1	0
000197000760	FA 3d:05		256	6	0	0	0
000197000760	FA 3d:04		256	6	0	0	0

Storage class = VPLEX

```
=====
----- Storage System -----
----- ID          Interface          Wt_Q          -- I/O Paths --   --- Stats ---
----- Total          Dead          Q-IOs      Errors
=====
```

ID	Storage Interface	System	Wt_Q	Total	Dead	Q-IOs	Errors
PMTPPA000000003	CL1-08		256	24	0	0	0
PMTPPA000000003	CL1-00		256	24	0	0	0
PMTPPA000000004	CL2-08		256	24	0	0	0
PMTPPA000000004	CL2-0C		256	24	0	0	0

Storage class = XtremIO

```
=====
----- Storage System -----
----- ID          Interface          Wt_Q          -- I/O Paths --   --- Stats ---
----- Total          Dead          Q-IOs      Errors
=====
```

ID	Storage Interface	System	Wt_Q	Total	Dead	Q-IOs	Errors
APM00141702677	X1-C2-f2		256	6	0	0	0
FNMO0164200011	X1-C2-f2		256	23	23	0	23
APM00141702677	X2-C2-f2		256	6	0	0	0
FNMO0164200011	X2-C2-f2		256	23	23	0	23
APM00141702677	X1-C1-f2		256	6	0	0	0
FNMO0164200011	X1-C1-f2		256	23	23	0	23
APM00141702677	X2-C1-f2		256	6	0	0	0
FNMO0164200011	X2-C1-f2		256	23	23	0	23
APM00141702677	X2-C1-f1		256	6	0	0	0
FNMO0164200011	X2-C1-f1		256	23	23	0	23
APM00141702677	X2-C2-f1		256	6	0	0	0
FNMO0164200011	X2-C2-f1		256	23	23	0	23
APM00141702677	X1-C2-f1		256	6	0	0	0
FNMO0164200011	X1-C2-f1		256	23	23	0	23
APM00141702677	X1-C1-f1		256	6	0	0	0
FNMO0164200011	X1-C1-f1		256	23	23	0	23

Storage class = Unity

```
=====
----- Storage System -----
----- ID          Interface          Wt_Q          -- I/O Paths --   --- Stats ---
----- Total          Dead          Q-IOs      Errors
=====
```

ID	Storage Interface	System	Wt_Q	Total	Dead	Q-IOs	Errors
APM00171906729	SP A4		256	21	0	0	0
APM00171906729	SP B4		256	21	0	0	0
APM00171906729	SP B5		256	21	0	0	0
APM00171906729	SP A5		256	21	0	0	0

Storage class = SC

```
=====
----- Storage System -----
----- ID          Interface          Wt_Q          -- I/O Paths --   --- Stats ---
----- Total          Dead          Q-IOs      Errors
=====
```

ID	Interface	Wt_Q	Total	Dead	Q-IOs	Errors
5000D3100346C400	14724-20	256	12	0	0	0
5000D3100346C400	14724-1f	256	12	0	0	0
5000D3100346C400	14725-1d	256	14	0	0	0
5000D3100346C400	14725-1e	256	14	0	0	0
5000D3100346C400	-1d	256	2	0	0	0
5000D3100346C400	-1e	256	2	0	0	0
5000D3100346C400	-1f	256	2	0	0	0
5000D3100346C400	-20	256	2	0	0	0

Storage class = PowerStore

ID	Storage System	Interface	Wt_Q	Total	Dead	Q-IOs	Errors
FNM00185100383	A1-NA-f3		256	14	0	0	0
FNM00185100383	A1-NB-f3		256	14	0	0	0
FNM00185100383	A1-NA-f2		256	14	0	0	0
FNM00185100383	A1-NB-f2		256	14	0	0	0
FNM00185000725	A0-NB-f3		256	14	0	0	0
FNM00185000725	A0-NB-f2		256	14	0	0	0
FNM00185000725	A0-NA-f2		256	14	0	0	0
FNM00185000725	A0-NA-f3		256	14	0	0	0

Display ports with NVMe devices

Storage class = Symmetrix

ID	Storage System	Interface	Wt_Q	Total	Dead	Q-IOs	Errors
000197600236	SA	2g:24	256	34	0	0	0
000197600236	SA	2g:25	256	34	0	0	0
000197600236	SA	4g:24	256	34	0	0	0
000197600236	SA	4g:25	256	34	0	0	0

rpowermt help

Displays usage summary for the rpowermt command.

Syntax

rpowermt help [<command>]

Description

The usage syntax is also displayed when invalid commands or options are entered.

Options

command

If a specific command is provided, displays the usage only of the command specified.

Help

```
# rpowermt help display host=<host>

display - Display multipathing information.

Usage:
  rpowermt <command>
    [class={all|symm|vnx|clariion|vplex|invista|netapp|hitachi|hpxp|ess|xtremio|unity|
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 alua 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 bus [class=<class>|all] [every=<#seconds>]
[width=<#col>]
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>]
rpowermt display port_mode [class=<class>|all]
rpowermt display ports [dev=<device>|all] [class=<class>|all]
[every=<#seconds>] [width=<#col>]

```

rpowermt register

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>]
```

Options

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.

rpowermt restore

Tests and restores the specified paths.

Syntax

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

Description

Restore issues test I/Os and responds to the test results as follows:

- 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 use it. In addition, other dead paths on the same HBA or storage system port may be tested.

- If a live path fails the test, `rpowermt restore` marks it dead and prints a warning (every time the path fails the test). In addition, other paths that share the same HBA and port may be marked dead, and other paths that share only the HBA or only the port may be tested.
- If a dead path fails the test, `rpowermt restore` prints a warning (every time the path fails the test).

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

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

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

Options

hba=*hba#*|all

Specifies the IP address or hostname of the vSphere host.

dev=*path*|*device*|all

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

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

Limits the command to the specified type of storage system. `all` specifies all storage-system types. 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.

rpowermt set autostandby

Enables or disables autostandby. Autostandby is disabled by default.

Syntax

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

Description

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

- Proximity-based autostandby (`asb:prox`) applies to path modes for multi-frame storage appliances, such as VMAX SRDF/Metro Arrays, and VPLEX cluster configuration where a host is connected to and manages a distributed volume across multiple VPLEX clusters and frames within a given VPLEX system. The proximity-based autostandby for VMAX SRDF/Metro is disabled by default to avoid issues with low-latency remote sites.

In mixed environments, such as VPLEX clusters with paths that are configured to cross-connected distributed volumes as well as paths configured to non-distributed volumes, the paths to the non-distributed volume inherit the cluster mode allocation that the proximity algorithm applies to its VPLEX cluster. For example, the PowerPath proximity algorithm designates a VPLEX cluster as local and preferred for a cross-connected, distributed volume on the host. After the distributed volume is designated as local, all paths to

that cluster will be set as active. In addition, even if there is a non-distributed volume in the cluster, paths to that non-distributed volume are also set as active. This setting is inherited from the global setting set on the VPLEX cluster.

The default for the prox trigger is off. While autostandby is enabled or if autostandby is reinitialized with the prox trigger, PowerPath performs a test on all paths, sets the appropriate paths to active or autostandby, and determines the preferred paths to a volume from federated storage.

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

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

The autostandby modes for each path do not persist across reboot. For example, if path set A is in autostandby mode and path set B is in active mode, then on reboot path set A can change to active mode and path set B can change to autostandby mode. On each reboot (assuming the global autostandby setting is on), the system performs latency measurements and re-runs the algorithms for selecting which path or path set to put in autostandby mode and which path or path set to put in active mode.

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

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

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

Options

autostandby=on

Enables autostandby. If you do not select a trigger, both prox and iopf are enabled. For example, `rpowermt set autostandby=on` globally turns on autostandby with both the proximity-based and the IOsPerFailure-based autostandby features.

autostandby=off

Disables autostandby. PowerPath clears the global setting for enabling autostandby (or the selected type of autostandby), and persists that setting. PowerPath also changes all paths in autostandby (or the selected type of autostandby) to active mode. If you do not select a trigger, then both prox and iopf are enabled. For example, `rpowermt set autostandby=off` globally turns off autostandby with both the proximity-based and the IOsPerFailure-based autostandby features.

autostandby=reinitialize

 NOTE: Applies to the proximity trigger only. Reinitialize has no effect on an enabled iopf trigger or when autostandby is turned off.

If the proximity trigger is enabled, this option recalibrates the trigger and determines anew the preferred paths to volumes from federated storage — this can be useful after a cluster has been physically moved, for example. If autostandby is not enabled, then an error message displays. If you do not select a trigger, then prox is enabled. For example, `rpowermt set autostandby=reinitialize` reinitializes autostandby with the proximity-based autostandby feature. The minimum latencies for each cluster and the mode changes for each path are logged in audit messages.

trigger=prox [class=class [threshold=value]]

Selects only the proximity-based autostandby feature. This determines and selects the preferred paths to a volume from federated storage and places the non-preferred paths into autostandby mode (asb:prox).

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

The `threshold` value, in microseconds, can range from 0 to 5000 (rounded to multiples of 50 after reboot). The threshold can be set only on a per-class basis.

Threshold value is used to decide paths of which Frame belonging to a multi-frame storage appliance. If the latency difference between two frames is greater than the user-defined value, then the frame with higher latency is marked as non-preferred and other as preferred. If the latency difference is lesser than the user-defined threshold value then the paths to both the frames are marked as active.

trigger=iopf

Selects only the IOsPerFailure-based autostandby feature. This determines whether a path that has intermittent I/O failures should be placed into autostandby 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 iopflimit

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

Syntax

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

Description

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

Options

iopflimit=value

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

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

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

Syntax

```
rpowermt set autostandby agingperiod={<#days>|<#hours>h|<#hours>H} host=<hostname>
[username=<username> [password=<password>|no_password]] [cim_sessionid=<CIM sessionid>]
```

Description

When a path is placed into autostandby mode through the IOsPerFailure trigger, PowerPath adds the aging amount to the current time. The aging period timeout allows PowerPath to return the paths to full use without user intervention. Error monitoring is resumed when the paths are returned to an active state. Each path ages separately. This may result in each path having different times at which they return to active mode.

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

This command sets the aging period for autostandby iopf globally within PowerPath and does not pertain to a specific storage array or volume. When you set aging period, this setting persists across reboot. The `rpowermt display options` command shows the aging period value.

Options

`agingperiod={#days|#hoursh|#hoursH}`

Specifies the number of days or hours after which a path in autostandby mode due to intermittent I/O failures automatically resets to active mode. The default number of days is 1. The range of acceptable values is 0 to 24854 days or 596520 hours. Use H or h to specify hours (for example, 6H or 6h for 6 hours). If you pass a value without H or h then it is considered as number of days.

Specifying 0 or 0h turns off the periodic reset of a path in autostandby due to intermittent I/O failures. 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 auto_host_registration

Enables or disables auto host registration.

Syntax

```
rpowermt set auto_host_registration={on|off} class=<class> host=<hostname>[username=<username>
[password=<password>|no_password]][cim_sessionid=<cim sessionid>]
```

Description

`rpowermt set auto_host_registration` enables or disables PowerPath to send host registration details to the array. User has to specify the class option for any array class to enable or disable the auto host registration.

Use `rpowermt display options` to verify if the `auto_host_registration` is on before trying `rpowermt update host_registration`.

NOTE: If `rpowermt update host_registration` is tried when the feature is disabled, error message is not displayed. Cluster details, OS patch level, and hardware vendor name are not captured in PPVE during Auto Host Registration.

Auto host registration for PowerMaxOS

This feature is available for PowerMaxOS 5978 and above. PowerPath sends host details such as, hostname, OS version, server hardware vendor, PowerPath license, WWN details, host connectivity type and VM information, and PowerPath version as part of host registration payload to VMAX. The host registration information is sent to VMAX as soon as Host reboots or when user triggers `rpowermt update host_registration` command manually. Host registration information is sent automatically every 24 hours. If any change in host configuration, host registration is sent within 1 hour by PowerPath

PowerPath also sends Virtual Machines details in ESX environment to array.

NOTE: Auto host registration is enabled by default for symm arrays.

Auto host registration for Unity, VNX, and CLARiiON

This feature is carried forward from PowerPath previous releases for VNX, Unity, and CLARiiON. Host sends the information that includes initiator type, hostname, and IP address as part of host registration payload. This registration is sent only once on host start or PowerPath starts.

NOTE: Auto host registration is disabled by default for CLARiiON, VNX, and Unity.

Auto host registration setting is persistent across restart.

The `rpowermt display options` command shows if the functionality is enabled or disabled.

Options

auto_host_registration=on

Enables auto host registration.

auto_host_registration=off

Disables auto host registration.

class={all|symm|vnx|clarion|vplex|invista|netapp|hitachi|hpxp|ess|xtremio|unity|sc|powermax|generic|powerstore}

Limits the command to the specified type of storage system. `all` specifies all storage-system types. 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.

Auto host registration

The `rpowermt display options` command shows that the functionality is enabled in the following snippet.

```
# rpowermt display options host=x.x.x.x
    Show CLARiiON LUN names:      true
    Path Latency Monitor: Off
    Performance Monitor: disabled
Autostandby:  IOs per Failure (iopf): enabled
              iopf aging period   : 1 h
              iopf limit          : 6000
Storage
System Class Attributes
-----
Symmetrix    periodic autorestore = on
              reactive autorestore = on
              proximity based autostandby = off
              auto host registration = enabled
              app finger printing = disabled
              device to array performance report = enabled
              device in use to array report = enabled
```

rpowermt set mode

Sets paths to active or standby mode.

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>]
```

Description

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

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

Options

mode=active

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

mode=standby

The path is held in reserve. Being set to standby does not mean a path will not be used. It only means the weight of the path is heavily adjusted to preclude its use in normal operations. A standby path still can be selected if it is the best path for a request. For example, if the active paths are so heavily loaded that the weighting factor in favor of the active path is overcome, the I/O will be sent to the standby path. Standby mode persists across reboots.

hba=hba#|all

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

dev=path|device|all

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

class={all|symm|vnx|clarion|vplex|invista|netapp|hitachi|hpxp|ess|xtremio|unity|sc|powermax|generic|powerstore}

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

force

Manually overrides autostandby. If a path is in autostandby mode, you must use the force argument to change it to active or standby. If the path is in autostandby and the force option is not specified, then an error is returned.

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

Enables or disables path latency monitoring.

Syntax

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

Description

When you enable or disable path latency monitoring, the setting applies globally to all paths.

NOTE: Before enabling latency monitoring for the first time, you should not only read this section but also `rpowermt set path_latency_threshold` for setting a threshold for path latency.

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. To set a path latency threshold, use the `rpowermt set path_latency_threshold` command.

The `rpowermt display options` command shows if the functionality is enabled or disabled

To avoid extraneous log messages, set a threshold value before turning on latency monitoring.

Options

`path_latency_monitor=on`

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

`path_latency_monitor=off`

Disables I/O latency measurement for all paths, and resets all paths' Current and Max latencies to zero. The default setting is off. Disabling latency monitoring does not change a previously set threshold. On re-enabling, the last threshold entered goes back into effect.

`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

Sets a time interval in seconds within which I/Os should complete.

Syntax

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

Description

When a threshold has been set, PowerPath generates system log messages indicating each threshold crossing that results in a new Max latency (high watermark) for a path. This threshold value applies to all paths in your environment.

For example, if

- the global `path_latency_threshold` is set to 2, and
- 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, periodically disable, then re-enable latency monitoring so that high watermarks are zero-ed and threshold crossings are captured in the system log with regularity.

When changing the threshold, use the following three steps:

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

Options

`path_latency_threshold=#seconds|#milliseconds`

Sets a time interval in seconds within which I/Os should complete. The value applies to all paths. The default value is 0. The range of acceptable values is 0 to 3600 seconds or 0 to 3600000 milliseconds. Seconds is the default.

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

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

Enables or disables performance monitoring for all devices.

Syntax

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


Description

Performance monitoring helps characterize I/O patterns and possibly aide in diagnosing I/O problems.

When performance monitoring is enabled, `rpowermt display perf dev=all` and `rpowermt display perf bus` displays performance metrics for all paths to all devices and bus ports.

Options

perfmon=on

Enables performance monitoring for all devices. This command initializes the counters, clears all performance measurements including high and low watermarks for latency, and accepts a sampling interval for data collection in seconds with a default interval of 900 seconds (15 minutes). The legacy path latency monitoring threshold measurements are not affected. There is no output.

perfmon=off

Disables performance monitoring and clears all measurements.

interval=#seconds

The sampling interval in seconds for data collection. The sampling interval range is from a minimum of 60 seconds (1 minute) to a maximum of 86,400 seconds (1,440 minutes or 1 day) with a default of 900 seconds (15 minutes).

To change the performance monitoring sampling interval, you must first set `perfmon=off` and then set `perfmon=on`.

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 vm_perfmon

Enables or disables performance monitoring for the specified VM.

Syntax

```
rpowermt set vm_perfmon={on|off} vmname=<VM Name> host=<hostname> [username=<username>
[password=<password>|no_password]] [cim_sessionid=<cim sessionid>]
```

Description

Performance monitoring helps characterize I/O patterns and possibly aide in diagnosing I/O problems.

You must enable `rpowermt set perfmon` before enabling `vm_perfmon`. The `vm_perfmon` sampling interval is the same as the `rpowermt set perfmon` sampling interval.

When performance monitoring is enabled, `rpowermt display perf vmname` displays performance metrics for all paths in the VM.

Options

vm_perfmon=on

Enables performance monitoring for the specified VM.

vm_perfmon=off

Disables performance monitoring for the specified VM and clears all measurements. The default is `off`.

vmname=VM Name

Displays performance metrics for VMPMI enabled VMs.

Set VM perfmon

```
# rpowermt set vm_perfmon=on vmname="RHELH12VM1" host=<>
```

rpowermt set periodic_autorestore

Enables or disables the periodic autorestore feature.

Syntax

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

Description

Periodic autorestore can be used along with reactive autorestore to define when paths are restored after path failure. To determine the periodic autorestore setting, use `rpowermt display options`.

Options

`periodic_autorestore=on`

Enables periodic autorestore. PowerPath periodically tests dead paths and, if they pass the test, restores them to service. The default is `on`.

`periodic_autorestore=off`

Disables periodic autorestore. Even when periodic autorestore is off, path testing continues to be done under certain conditions as described in the *PowerPath Family Product Guide*, and automatic path restoration continues to occur based on the results of that testing.

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

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

`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 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

Sets the load-balancing and failover policy for devices configured in PowerPath.

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>]
```

Description

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 Unity, VNX, and CLARiiON devices, and ad (Adaptive) for Invista, VPLEX, VNXe, PowerStore, 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. You cannot manually set the policy for a device to bf or nr.

Options

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

Sets the load-balancing and failover policy to the specified value.

dev=device|all

Limits the policy change to the specified devices.

class={all|symm|vnx|clarion|vplex|invista|netapp|hitachi|hp|ess|xtremio|unity|sc|powermax|generic|powerstore}

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

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 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 reactive_automorestore

Enables or disables PowerPath's reactive automorestore facility.

Syntax

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

Description

Reactive automorestore is the automatic restoration in the context of an I/O of physically alive paths that are still in the dead (unusable) state.

The `rpowermt display options` command shows if the functionality is enabled or disabled

Reactive automorestore can be used with periodic automorestore to define when paths are restored after path failure.

- If `reactive_automorestore=on` and `periodic_automorestore=on`, then paths will be automatically restored at the next opportunity in either context. This is the default for these settings.
- If `reactive_automorestore=on` and `periodic_automorestore=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_automorestore=off` and `periodic_automorestore=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. Do not set this 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.

Options

`reactive_autorestore=on`

Enables reactive autorestore, which reactively tests dead paths and, if they pass the test, restores them to service. For non-disruptive upgrades (NDU), reactive autorestore should be set to `on`. The default is `on`.

`reactive_autorestore=off`

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

`class={all|symm|vnx|clarion|vplex|invista|netapp|hitachi|hpxp|ess|xtremio|unity|sc|powermax|generic|powerstore}`

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

`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 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 streamio_threshold

Sets the I/O threshold values for when to switch paths while the Stream I/O (si) load-balancing and failover policy is in effect.

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>]
```

Description

The `streamio` value is accepted only if the load-balancing policy in effect on the device is `si`. If the PowerPath license is invalid, setting the `streamio_threshold` will generate an error message. You can check whether Stream I/O threshold is set by viewing the `rpowermt display dev` output.

Options

`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 logical device. `all` specifies all logical devices. The default is `all`.

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

Limits the command to the specified type of storage system. `all` specifies all storage-system types. 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.

rpowermt set device_perf_to_array_report

Enables device performance statistics reporting to the array.

Syntax

```
rpowermt set device_perf_to_array_report={on|off} [class=<class>|all]
```

Description

User has to specify class option to enable or disable device performance reporting for an array.

This feature is available for PowerMaxOS 5978 microcode and above. PowerMax collects statistics for number of reads/writes, IO bandwidth of reads/writes, and latency of reads/writes and present in Unisphere. For more information see Unisphere documentation.

PowerPath sends performance statistics to PowerMax once in every minute for a maximum of 400 devices. This feature is persistent across host restart.

Options

on

Enables device performance statistics reporting to the array.

off

Disables device performance statistics reporting to the array.

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

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

rpowermt set dev_inuse_to_array_report

Enables device report to the array.

Syntax

```
rpowermt set dev_inuse_to_array_report={on|off} [class=<class>|all]
```

Description

User has to specify class option to enable or disable this feature for an array.

This feature is available for PowerMaxOS 5978 microcode and above. PowerPath sends the VM name, which is using the device. PowerPath also sends if the device is mounted on the host, the current usage status of the device, and the last time the device was in use. This information is sent to array once in 24 hours.

Options

on

Enables device reports to the array

off

Disables device reports to the array

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

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

rpowermt set port_mode

Sets the mode of all paths of a port on any storage system.

Syntax

```
rpowermt set port_mode={active|standby} dev=<path> [force] host=<hostname> [username=<username> [password=<password>|no_password]] [cim_sessionid=<CIM sessionid>]
```

Description

The `rpowermt display port_mode` command shows the mode.

Options

port_mode=active

User I/O is delivered to a path. If you designate a path as active, it is continuously scheduled for I/O according to the load-balancing and failover policy. An active port operates as it normally would within a PowerPath environment.

port_mode=standby

The path is held in reserve. Being set to standby does not mean a path will not be used. It only means the weight of the path is heavily adjusted to preclude its use in normal operations. A standby path still can be selected if it is the best path for a request. For example, if the active paths are so heavily loaded that the weighting factor in favor of the active path is overcome, the I/O will be sent to the standby path.

dev=path

Sets the mode for the specified path.

force

If a previous `rpowermt display port_mode` 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.

Set port mode

```
# rpowermt set port_mode=standby dev=vmhba4:C0:T6:L33 host=<host>
```

rpowermt setup add_host

Adds a vSphere host and its corresponding username/password pair to the specified lockbox (if the `host_file` option 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>]
```

Options

host_file=filename

Lockbox file on the rpowermt server. If the specified lockbox does not exist, it is created. If the `host_file` option 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

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.

Add host

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

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>]
```

Options

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.

List hosts

If the specified lockbox exists and it has stored entries, a list of host IPs is displayed.

```
# rpowermt setup list_hosts
  IP address                Hostname
  xx.xx.xx.x1              <host1.domainname.com>
  xx.xx.xx.x2              <host2.domainname.com>
  xx.xx.xx.x3              <host3.domainname.com>
  xx.xx.xx.x4              <host4.domainname.com>
```

rpowermt setup remove_host

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>
```

Description

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

Options

host_file=filename

Removes 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.

rpowermt setup update_host

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>]
```

Options

host_file=filename

Lockbox file on the rpowermt server. If the specified lockbox does not exist, it is created. If this 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

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.

Update hosts

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

Lists the hosts/IPs from the `host_file` (that is, the rpowermt lockbox) that are reachable.

Syntax

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

Description

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` option is specified).

Options

host_file=filename

Lockbox file on the rpowermt server. If the specified lockbox does not exist, it is created. If this option 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.

Verify hosts

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

```
# rpowermt setup verify_hosts host=10.249.216.18
      Host Verified          Status
10.249.216.18              OK
```

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

```
# rpowermt setup verify_hosts
      Host(s) Verified      Status
10.63.12.29                OK
10.249.216.21              OK
10.249.216.22              OK
10.249.216.19              OK
10.249.216.15              OK
10.249.216.17              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 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=10.249.216.18
Host Verified      Status
10.249.216.18     OK
```

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

```
# rpowermt setup verify_hosts host=10.249.216.56
ERROR: Host not found in Lockbox file.
```

rpowermt unregister

Removes a PowerPath license for a vSphere host.

Syntax

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

 **NOTE: This command is not applicable if host is registered using FX.**

Options

- 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 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.

Unregister

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 FX License:

```
# rpowermt unregister host=x.x.x.x
ERROR:PowerPath is licensed using FX license,can not unlicense the PowerPath driver.
```

rpowermt update host_registration

Updates host registration.

Syntax

```
rpowermt update host_registration host=<hostname> [username=<username> [password=<password>|no_password]] [cim_sessionid=<cim sessionid>]
```

Description

rpowermt update host_registration command triggers a host registration to all connected arrays that support host registration. This command can be used immediately to update an array whenever there is any change in host details.

This command always returns a success even if there are no arrays connected, if auto_host_registration is disabled on the host, or if auto_host_registration is not supported/enabled on the connected arrays.

This command does not return any status when run. Ensure that auto_host_registration feature is enabled before running this command.

rpowermt update lun_names

Retrieves the latest EMC Unity, VNX, CLARiiON, XtremIO, and IBM XIV user-assignable LUN names.

Syntax

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

Description

A user-assignable LUN name is a character string that a user or system manager associates with a LUN and assigns through Unisphere or Navisphere. CLARiiON LUN nice names are updated automatically every hour. Use this command to manually update CLARiiON nice names.

Options

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.

Update LUN names

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

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

rpowermt version

Displays the 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>]
```

Options

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.

Version

```
# rpowermt version host=x.x.x.x  
Dell EMC rpowermt for PowerPath (c) client Version 7.0 (build 194)  
Dell EMC PowerPath (c) host=x.x.x.x Version 7.0 (build 194)  
License search path: /etc/emc:/etc/emc/licenses:/opt/EMCpower:/opt/EMCpower/licenses  
Host file: /etc/emc/lockbox/root.clb
```

PowerPath family environment variables

This chapter contains the following topics:

Topics:

- [Disable autosave of configuration](#)
- [Display time stamp](#)
- [Display width](#)
- [Set rpowermt server to location](#)
- [Migration Enabler Remote SE port](#)
- [Migration Enabler Remote SE server](#)
- [Disable autoregistration](#)
- [Disable the lockbox](#)
- [Change lockbox name and location](#)
- [Set a non-default license path directory](#)

Disable autosave of configuration

Disables autosave of the PowerPath configuration.

Syntax

```
set PP_DISABLE_AUTOSAVE=<0|1>
```

Description

On PowerPath for Windows 5.7 SP2 and later, the `PP_DISABLE_AUTOSAVE` environment variable is set to 1 by default.

Options

```
PP_DISABLE_AUTOSAVE=<0|1>
```

0 (False) enables autosave. 1 (True) disables autosave.

Display time stamp

Displays a time stamp to the `powermt display` command output.

Syntax

```
set PP_DISPLAY_TIME_STAMP=TIME_VERBOSE|TIME_SECONDS
```

Description

You can add a time stamp to the `powermt display every=<#seconds>` command output by defining the `PP_DISPLAY_TIME_STAMP` environment variable and setting the time display format to `TIME_VERBOSE` or `TIME_SECONDS`. If the environment variable is not defined or if it is defined but not assigned a proper value the time stamp is not displayed.

Options

TIME_VERBOSE

Displays complete details such as day, month, and year, time in hours, minutes, and seconds that increments until you exit the `powermt display` command.

TIME_SECONDS

Starts a counter from zero that increments until you exit the `powermt display` command.

Time stamp

The following example shows `TIME_VERBOSE` value output:

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

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

The following example shows `TIME_SECONDS` value output:

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

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

Display width

Changes the width setting from 80 (the default) to the value you specify.

Syntax

```
set PP_DISPLAY_WIDTH=<#columns>
```

Description

You can set a persistent value for the `width=<#columns>` option by using the environment variable `PP_DISPLAY_WIDTH`. If the environment variable value differs from the value you set with the command line, the command line value takes precedence.

Options

PP_DISPLAY_WIDTH=*#columns*

The acceptable range is 80 to 160 columns. The output widens or narrows within the range depending on the width of the window.

Set rpowermt server to location

If the `rpowermt` CST libraries have been moved, set the environment variable to point the `rpowermt` server to the correct location.

Syntax

```
set PP_LB_LIB=<location of CST libraries>
```

Description

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.

Options

PP_LB_LIB=location of CST libraries

Sets the location of the CST library.

Set the location

To point the rpowermt server to `/usr/lib` on a Linux system, type:

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

To point the rpowermt server to `C:\program files\EMC\PowerPath\rpowermt` on a Windows system, type:

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

Migration Enabler Remote SE port

In rare cases if the remote port changes after a migration has started, this environment variable allows you to update values that were set during `powermig setup`.

Syntax

```
set PPME_REMOTE_SE_PORT=<port>
```

Description

This variable can only be used with post-setup `powermig` commands (that is, this cannot be used with `powermig setup`).

An error is returned if this variable is present during `powermig setup`.

Once used, values are saved and used for the duration of that migration (variables can and should be removed).

Options

PPME_REMOTE_SE_PORT=port

Identifies the port used for Remote SE.

Migration Enabler Remote SE server

In rare cases if the remote server changes after a migration has started, this environment variable allows you to update values that were set during `powermig setup`.

Syntax

```
set PPME_REMOTE_SE_SERVER=<server>
```

Description

This variable can only be used with post-setup `powermig` commands (that is, this cannot be used with `powermig setup`).

An error is returned if this variable is present during `powermig setup`.

Once used, values are saved and used for the duration of that migration (variables can and should be removed).

Options

`PPME_REMOTE_SE_SERVER=server`

Identifies the name of the server.

Disable autoregistration

Disables automatic license registration.

Syntax

```
set PPMT_DISABLE_LICAUTO=1
```

Description

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

Options

`PPMT_DISABLE_LICAUTO=1`

1 (True) disables autoregistration.

Disable the lockbox

Disables writing to and reading from the lockbox.

Syntax

```
set PPMT_LB_DISABLE=1
```

Description

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

Options

`PPMT_LB_DISABLE=1`

1 (True) disables the lockbox.

Disable the lockbox

To disable the lockbox on a Linux system, type:

```
export PPMT_LB_DISABLE=1
```

To disable the lockbox on a Windows system, type:

```
set PPMT_LB_DISABLE=1
```

Change lockbox name and location

Changes the lockbox name and location.

Syntax

```
set PPMT_LB_FILE=<file>
```


Options

PPMT_LB_FILE=*file*

Sets the lockbox name and location.

Change the lockbox name and location

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

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

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

```
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
```

Set a non-default license path directory

Sets a non-default license path if you want to point to a single license file instead of a general directory.

Syntax

```
set PPMT_LIC_PATH=<license_file_directory>
```

Options

PPMT_LIC_PATH=*license_file_directory*

Directory of the license file.

Set a non-default directory

To change the directory on a Linux system, type:

```
setenv PPMT_LIC_PATH=license_file_directory
```

To change the directory on a Windows system, type:

```
set PPMT_LIC_PATH=license_file_directory
```

PowerPath error codes

PowerPath kernel error messages are prefixed with error codes in PowerPath for Windows 7.0.

Topics:

- [PowerPath error messages](#)

PowerPath error messages

This section contains the list of error codes and messages that appear in PowerPath.

Table 11. PowerPath error messages

Error Code	Error Message	Action
EKBC001	EKBC001: pDeviceIdState refCnt overflow. opens and closes will be ignored.	-
EKBC002	EKBC002: pDevState refCnt overflow. opens and closes will be ignored.	-
EKBC003	EKBC003: SwapTop stuck userDev1 %lx userDev2 %lx.	Check if the devices are not in use before unmap and swap is performed.
EKBC004	EKBC004: PxCdbRwX: Unable to get sector size information for disk %p	Confirm that the device is not faulty.
EKBC005	EKBC005: PxCdbRwX: Unable to get sector size information for disk %p	Confirm that the device is not faulty.
EKBW001	EKBW001: DriverEntry: Cannot create driver object extension.	Check if sufficient memory is available on the system.
EKBW002	EKBW002: DriverEntry: Cannot allocate paged memory.	Check if sufficient memory is available on the system.
EKBW003	EKBW003: DriverEntry: Cannot create directory object for power devices.	Check if sufficient memory is available on the system.
EKBW004	EKBW004: Cannot create daemon thread	-
EKBW005	EKBW005: Cannot create daemon thread	-
EKBW006	EKBW006: PowerWinIsIrpSync:Unknown Major function %d	-
EKBW007	EKBW007: WinGetDeviceNumber returning status 0x%x	-
EKBW008	EKBW008: PowerWinCreateDirObject: Could not create directory %s	-
EKBW009	EKBW009: PowerWinCreateIoctlPDO(): Cannot create a Ioctl communication device	-
EKBW010	EKBW010: PowerWinCreateIoctlPDO: Cannot create symbolic link	-

Table 11. PowerPath error messages (continued)

Error Code	Error Message	Action
EKBW011	EKBW011: Cannot create bus fdo. Cannot load driver	-
EKBW012	EKBW012: AddDevice: Cannot add native device (0x%x)	-
EKBW013	EKBW013: AddDevice: Unable to attach the device to target	-
EKBW014	EKBW014: PowerWinCreateScsiPortFiDO: IoCreateDevice failed with status (0x%x)	-
EKBW015	EKBW015: PowerWinCreateScsiPortFiDO: Unable to attach the device to target	-
EKBW016	EKBW016: PseudoDevice: Cannot create psuedo device	-
EKBW017	EKBW017: ClassClaimDevice: Can't allocate Irp	-
EKBW018	EKBW018: PowerWinAddDevice: Cannot attach to PDO (0x%x)	-
EKBW019	EKBW019: PxWinIocmdSwapTop: Iocmd(%x) fails with %x.	-
EKBW020	EKBW020: PowerStartFilterDevice(): DevObj(0x %x) Lower DO failed start.	-
EKBW021	EKBW021: PowerStartFilterDevice(): DevObj(0x %x) GetScsiAddress fails	-
EKBW022	EKBW022: PowerWinStartFilterDevice(): DevObj(0x%x) Read Cap fails	-
EKBW023	EKBW023: PowerWinStartScsiLunFiDO: DevObj(0x %x) Lower DO failed start	-
EKBW024	EKBW024: PowerStartScsiLunFiDO: DevObj(0x %x) GetScsiAddress fails	-
EKBW025	EKBW025: PowerStartScsiLunFiDO: DevObj(0x %x) IsMpioPseudoDevice fails	-
EKBW026	EKBW026: PowerStartScsiLunFiDO: DevObj(0x %x) GetBusType fails	-
EKBW027	EKBW027: PowerWinStartScsiLunFiDO(): DevObj(0x%x) Read Cap fails	-
EKBW028	EKBW028: PowerWinScsi: %s: 0x%x; Receives an invalid request(0x%x)	-
EKBW029	EKBW029: Cannot attach to pseudo bus	-
EKDC001	EKDC001: DmRecognize- unable to set up device.	-
EKDC002	EKDC002: FUNNEL_IO status 0x%x kmd key 0x %lx dmd key 0x%lx.	-
EKDC003	EKDC003: LIMIT_ACCESS status 0x%x kmd key 0x %lx dmd key 0x%lx.	-

Table 11. PowerPath error messages (continued)

Error Code	Error Message	Action
EKDC004	EKDC004: DmRecognize: unable to manage device.	-
EKDC005	EKDC005: DmRecognize- unable to set up device (deferred).	-
EKDC006	EKDC006: REDIRECT_IO status 0x%x kmd key 0x %lx dmd keys 0x%lx 0x%lx.	-
EKDC007	EKDC007: CLONE_WRITES status 0x%x kmd key 0x %lx dmd keys 0x%lx 0x%lx.	-
EKDC008	EKDC008: DmRecognize- unable to manage device	-
EKDW001	EKDW001: DM installation is failed.	-
EKDW002	EKDW002: _fini DmUninstall failed.	-
EKGW001	EKGW001: _fini GpxUninstall failed.	-
EKMC001	EKMC001: SymmetrixKULam_getPNTData cmdStatus %x ScsiStatus %x SKey %x ASC %x ASCQ %x	-
EKMC002	EKMC002: Failed to allocate 3b/3c chain	Check if sufficient memory is available on the server.
EKMC003	EKMC003: Device Mount Stats Rpt Failed with S:%x ASC:%x ASCQ:%x	-
EKMC004	EKMC004: Bus instance overflow; increase MP_MAX_BUSPORTS_ORDER.	-
EKMC005	EKMC005: Path latency monitor daemon creation failed	Check if the server is out of memory.
EKMC006	EKMC006: Bus %d to %s port %s is dead.	Check the connectivity between the server and the array.
EKMC007	EKMC007: Invalid initiator info: tag=0x%x	-
EKMC008	EKMC008: PowerPath Auto Host Registration to %s-%s failed on Bus %d Port %s	-
EKMC009	EKMC009: PowerPath Auto Host Registration to %s-%s failed on Bus %d Port %s as the array's host record table is full.	-
EKMC010	EKMC010: PowerPath Auto Host Registration to %s-%s failed on Bus %d Port %s	-
EKMC011	EKMC011: MpxIocmdGetElmsLicInfo : size mismatch.	-
EKMC012	EKMC012: MpxIocmdSetElmsLicInfo : size mismatch.	-
EKMC013	EKMC013: Performance monitoring global enable failed	Check if sufficient memory is available on the server.
EKMC022	EKMC022: HostIoLmt LookupFailed pArraySpecificData NOT initialized	-

Table 11. PowerPath error messages (continued)

Error Code	Error Message	Action
EKMC023	EKMC023: Could not set PMI timer	-
EKMC024	EKMC024: Performance Stats Rpt Failed with S:%x ASC:%x ASCQ:%x	-
EKMC025	EKMC025: Could not set Storage perf stats reporting timer	-
EKMC026	EKMC026: Updating size on path %s to %s	-
EKMC027	EKMC027: SP%c Not Responding on %s: Trespassing volumes.	During trespass, verify that the SPs of the arrays are up and responsive.
EKMC028	EKMC028: SP%c Not Responding on %s: Trespass %s.	During trespass, verify that the SPs of the arrays are up and responsive.
EKMC029	EKMC029: Volume %s is unbound.	Check if the LU is bound.
EKMC030	EKMC030: Path %s to %s is marked as dead;volume ID has changed.	Check if the existing data LUN is not replaced with another valid data LUN.
EKMC031	EKMC031: %s is dead.	Check the status of the LUN on the array.
EKMC032	EKMC032: Path %s to %s is dead.	Check the connectivity between the server and the array.
EKMC033	EKMC033: MpxRegister: PRO re-Register failed, zero key for %s	-
EKMC034	EKMC034: MpxRegister: PRO service action %x failed, 0x5/0x%x/0x%x for %s	-
EKMC035	EKMC035: Unable to allocate memory, skipping mount status	Check if sufficient memory is available on the server.
EKMC036	EKMC036: SCSI-2 reservation failed with SCSI status 0x%x	-
EKMC037	EKMC037: SCSI-2 reservation failed with SCSI status 0x%x	-
EKMC038	EKMC038: SCSI-2 release failed with SCSI status 0x%x	-
EKMC039	EKMC039: ATS Error Response ScsiStatus %x SKey %x ASC %x ASCQ %x	-
EKMC040	EKMC040: MpxSendScsiCmd: CDB %x %x %x %x %x %x %x %x %x %x %x %x %x cmdStatus %x ScsiStatus %x SKey %x ASC %x ASCQ %x	-
EKMC041	EKMC041: MpRecognize: MpxGetVolIdFromInq failed	-
EKMC042	EKMC042: Unable to start device %x	-
EKMC043	EKMC043: SetPortMode: wrong size!	-
EKMC044	EKMC044: SetPortModeStandby: wrong size!	-
EKMC045	EKMC045: MpxIocmdSetResetDevLossTmo: pData is NULL or size does not match	-

Table 11. PowerPath error messages (continued)

Error Code	Error Message	Action
EKMC046	EKMC046: Wrong size: got %d must be > %d	-
EKMC047	EKMC047: Wrong sbuffer ize: got %d must be %d	-
EKMW001	EKMW001: DriverEntry(Mpx): failed to create pp generic daemon	Check if sufficient memory is available on the system.
EKMW002	EKMW002: PxDsmAddDevice(): Cannot create device object	-
EKMW003	EKMW003: PxDsmAddDevice: Cannot create symbolic link	-
EKMW004	EKMW004: PxDsmInterpretError: fails Irp 0x01 due to no suitable path	This error message is due to a volume going into "All paths dead." Check for the dead paths to the volume and restore it.
EKMW005	EKMW005: PxDsmInterpretError: retry Irp 0x01 on low memory	-
EKMW006	EKMW006: MpxDeviceUsageProcessNameRequest: Unable to allocate memory. Skipping. vol %s	Check if sufficient memory is available on the system.
EKMW007	EKMW007: PxDsmLBGetPath: Srb 0x08 CPath(%x) SPath(0x%x) with status 0x00000001L	-
EKMW008	EKMW008: PxDsmRwIrpComplete: pPirp(0x%x)-Irp(0x%x) is done with status(0x%x)	-
EKMW009	EKMW009: PxDsmSrbDeviceControl: Irp(0x%x) failed with status(0x%x)	-
EKMW010	EKMW010: PxDsmCategorizeRequest: Irp(0x%x) PrevPathId(0x%x) SelectedPathId(0x%x) Status(0x%x)	-
EKMW011	EKMW011: PxDsmGetPathFromMpioPassThruIoctl: SPath(0x%x) with status(0x%x)	-
EKMW012	EKMW012: PxDsmRemovePath: DSM-%s, Path(0x%x) is removed	-
EKMW013	EKMW013: PxDsmRemoveDevice: Path(0x%x) of DSM-%s on PathId(0x%x) is removed	-
EKMW014	EKMW014: PxDsmRemovePending: Path(0x%x) of DSM-%s is in remove pending	-
EKMW016	EKMW016: PxDsmPathVerify: Path 0x%x on Busport 0x%x is tested dead	-
EKMW019	EKMW019: PxDsmTopIodone: pPirp(0x%x)-Irp(0x%x) is done with status(0x%x)	-
EKMW020	EKMW020: PxDsmDereferenceMPIOIoctlPDO: ref count is 0!	-
EKMW021	EKMW021: PxDsmGetMPIOIoctlPDO: failed to get ioctl PDO with status(0x%lx)	-
EKMW022	EKMW022: MpxIodoneUp: %s %s on Vol %s.	-
EKMW023	EKMW023: MpxIodoneUp: Reset failed on Vol %s	-

Table 11. PowerPath error messages (continued)

Error Code	Error Message	Action
EKMW024	EKMW024: MpxIodoneUp: PR CDB(%02x%02x%02x%02x%02x%02x%02x%02x) %s on Vol %s.	-
EKMW025	EKMW025: Path %s to %s is dead	-
EKMW026	EKMW026: All paths to volume %s are dead.	Application that is configured on this volume will have I/O errors as all the paths are dead. Check the array, fabric, and host configuration to know the reason for dead paths.
EKMW027	EKMW027: The current owner of volume %s changed to SP %s.	-
EKMW028	EKMW028: Path %s to %s is dead.	-
EKMW029	EKMW029: MpxAssignAllVolumes: SP%c Hang on %s: Trespassing volumes.	-
EKMW030	EKMW030: MpxAssignAllVolumes: SP Hang on %s, SP%c: Trespass %s.	-
EKMW031	EKMW031: Volume %s is unbound.	-
EKMW032	EKMW032: MpxTestPath: %s-%s %s is dead	-
EKMW033	EKMW033: MpxMarkPathsOnPortOffline Unable to allocate memory, aborting port offline	Check if sufficient memory is available on the system during port offline.
EKMW034	EKMW034: Unable to allocate memory, skipping mount status	Check if sufficient memory is available on the system mount status request.
EKMW035	EKMW035: Unable to follow volume %s to %s	-
EKMW036	EKMW036: Unable to restore volume %s to default Controller	-
EKMW037	EKMW037: Unable to follow volume %s to %s	-
EKMW038	EKMW038: MpxDekDecodePirp: DEK: Checksum mismatch opcode %X; device 0x%p.	-
EKMW039	EKMW039: MpxDekDecodePirp: DEK: Invalid Handshake signature %I64X; device 0x%p.	-
EKMW040	EKMW040: MpxDekDecodePirp: DEK Disassoc: Invalid MPx handle %I64X; device 0x%p.	-
EKMW041	EKMW041: MpxDekDecodePirp: DEK Disassoc: Pending IO detected. Failing request; device 0x%p.	-
EKMW042	EKMW042: MpxDekDecodePirp: DEK Assoc: Pending IO detected. Failing request; device 0x%p.	-
EKMW043	EKMW043: MpxDekDecodePirp: DEK Assoc: %i is not supported number of encryption ranges. Failing request; device 0x%p.	-
EKMW044	EKMW044: MpxDekDecodePirp: DEK Assoc: Multiple association for single device is not supported. Failing request; device 0x%p.	-

Table 11. PowerPath error messages (continued)

Error Code	Error Message	Action
EKMW045	EKMW045: MpxDekDecodePirp: DEK Assoc: Multiple association for single device is not supported. Failing request; device 0x%p.	-
EKMW046	EKMW046: MpxDekPostprocessIodone: DEK failed with opcode %X; status %X; %s; device 0x%p.	-
EKMW047	EKMW047: MpxDekPostprocessIodone: DEK failed to verify handshake payload device 0x%p.	-
EKMW048	EKMW048: MpxDekPostprocessIodone: DEK Disassoc: Reset eHBA handle %I64X; device 0x%p	-
EKMW049	EKMW049: MpxDekIodoneUp: DEK: Invalid Handshake signature %I64X; device 0x%p.	-
EKMW050	EKMW050: PxDsmPnPNotificationPPDisk: DEK Disassoc encountered unexpected error %x for device 0x%p on handle %I64X.	-
EKMW051	EKMW051: PowerAllocatePirp: Allocates a new Pirp 0x%x	-
EKMW052	EKMW052: PowerBottomDispatch: Dispatch Irp 0x%x on DevObj 0x%x failed with status(%x)	-
EKMW053	EKMW053: PxDsmWmiQueryLBPolicy: No Path DsmId is given.	-
EKMW054	EKMW054: PxDsmWmiQueryLBPolicy: Output buffer too small.	-
EKMW055	EKMW055: PxDsmWmiQuerySupportedLBPolicies: Output buffer too small.	-
EKMW056	EKMW056: PxDsmWmiValidateSetLBPolicyInput: Insufficient buffer in SetLB.	-
EKMW057	EKMW057: PxDsmWmiValidateSetLBPolicyInput: Insufficient buffer in SetLB.	-
EKMW058	EKMW058: PxDsmWmiValidateSetLBPolicyInput: WMI Version mismatch.	-
EKMW060	EKMW060: PxDsmWmiValidateSetLBPolicyInput: Invalid LB Policy %d.	-
EKMW061	EKMW061: PxDsmWmiValidateSetLBPolicyInput: Path Count %d not equal to DSM IDs count %d.	-
EKMW062	EKMW062: PxDsmWmiValidateSetLBPolicyInput: Duplicate path id %I64x at %d and %d.	-
EKMW063	EKMW063: PxDsmWmiSetLBPolicy: No DSM Ids given.	-
EKMW066	EKMW066: PxDsmWinSetLBPolicy: Failed to validate input. Status %x.	-
EKMW067	EKMW067: PxDsmQueryData: GuidIdx-%d, Buffer too small.	-

Table 11. PowerPath error messages (continued)

Error Code	Error Message	Action
EKMW068	EKMW068: PxDsmQueryData: GuidIdx-%d, Buffer too small.	-

Multipathing messages

This chapter contains the following topics:

Topics:

- <device_id> is alive
- <device_id> is dead
- /etc/powermt: Permission denied
- /opt/emcpower/powercf does not exist
- /opt/emcpower/powercf: you have to be root to execute this program
- Can't get Pirp_t!
- Cannot install the new configuration files
- Cannot mount root on <>:<>
- Cannot obtain lock on <filename>
- Cannot obtain driver state
- Cannot open /dev/emcp, errno <number>
- Cannot open /etc/path_to_inst. Cannot continue
- Cannot open <filename>, errno <number>
- Cannot open configuration file <filename>
- Cannot open dev dir for reading, errno <number>
- Cannot open string Initialization error
- Cannot produce new files
- Cannot remove device <device_id>, name in use by application
- Cannot write to file
- Cannot write to <filename>
- CLARiiON device path c3t1d0s0 is currently misconfigured
- Configuration database corrupted
- Configuration files contains a dead path (device)
- Current/Active controller for volume has changed
- Device is unavailable
- Error number adding pnp path <hdisk_name> for <hdiskpower_name> to CuAt
- Error number getting pnp path attribute for <hdiskpower_name>
- Error number updating CuDv attribute for <device_id>
- Error retrieving minor number(s)
- Failed opening adapter <device_number>
- Failover mode PNR for CLARiiON LUN <WWN#> is not recommended. Please consider PAR or ALUA
- File deleted. Use -add to create a new file
- File not deleted. No licenses will be available until file is deleted and re-created
- Get adapter <devno> failed number
- Incomplete installation detected
- Invalid base minor number
- Invalid minor number
- Key <key_number> in license file is bad. Use -edit option to repair
- Key <key_number> is invalid
- Key <key_number> was not found
- Key <key_number> is already present, ignored
- License file is full -- cannot add number keys
- mknod of device failed
- MPxIO is enabled. Verify that no arrays are managed by both MPxIO and PowerPath
- Not a multiple of 8
- Number of minors out of range 1-1000

- One or more expired PowerPath license key(s) detected. These are typically used...use emcprep utility to review your license keys
- PRO service, action 46 failed, 0x5/0x80/0x20
- Pseudo Device Name Major# Minor# <ERROR: unexpected internal error>
- Read of file <file> failed
- There are no keys presently registered. Some functions and capabilities may be unavailable
- This version of PowerPath is NOT binary compatible with your kernel
- Forceload of drv/sd failed
- Get hdisk path data failed
- Internal DB read failed
- Kernel symbol versioning is not enabled in your linux kernel
- Killing bus <bus> to Symmetrix <id> port <port>
- Failed to load extension errno = <number>
- Recognize (loop) error number on device <device_id>
- Storage system <storage_system_name> PowerPath mode disabled
- Not all paths to <volume> are dead
- SPx Not Responding on array : Trespassing volumes
- ODM PdDv database contains duplicate uniquetypes. Refer to EMC Knowledge Base Article Number: 000084491 or contact EMC Customer Support personnel
- Path <path_name> to <device_id> is dead
- <storage_system> path <path_name> is dead. Do you want to remove it (y/n/a/q)?
- Path <bus> <tgt> <lun> to <device> is dead
- Open powerpath0 failed number
- PRO service action 6 failed, <Sense Key>/<ASC>/<ASCQ> for Bus <bus> Tgt <target> Lun <lun>
- Restored volume <clariion/vnx_volume> to default: SP<clariion/vnx_storage_processor>
- Reviving bus <bus> to Symmetrix <id> port <port>
- Sleeping for memory
- Syscall 8133 to <device_id> failed
- Trespassed volume <clariion_volume> to SP<clariion_storage_processor>
- Unable to restore volume <volume> to default controller
- Unable to restore volume <clariion_volume> to default SP
- Unable to follow volume <volume> to controller
- Encountered unexpected error 0x<code> on Dek operation command for <device id> on handle <association_handle>
- Unmanage of device <device> failed
- Getting unmanaged devices failed with error
- Unsupported configuration (R1/R2 setup detected)
- Unsupported HBA
- Volume <clariion_volume> followed to SP<clariion_storage_processor>
- Volume <clariion_volume> is unbound
- Write failed for lack of space, check for filesystem full
- Cannot open PowerPath. Initialization error
- Configuration files corrupted/out of date
- PowerPath configuration files present - exiting
- Major number is already in use
- PowerPath modules are not loaded. Ensure that the modules are loaded and try again
- An older EMCpower package is already installed. Please exit the installation and remove the old EMCpower package first
- PowerPath pseudo bus <bus_name> cannot be configured. Uninstall and reinstall PowerPath
- Unable to create kernel process
- Unable to open /etc/PowerPath Extensions
- Unable to start kernel process
- EMC registration data is missing or corrupt
- EMCP registration file does not exist. Use -add to create it
- EMCP registration file is corrupt
- EMCP registration file is corrupt. Delete it?
- License for <storage_system_name> storage system support is missing or expired
- There are no license keys now registered
- There are no license keys now registered
- License not installed

- License not installed
- Temporary upgrade key has been superseded
- Key found in old registration file is invalid
- Mismatch in configuration
- Missing configuration files
- All paths to vol <volume> are dead
- Attempted action for <storage_array> storage system support is not allowed with the current license level
- Attempted action is not allowed for volumes with BF or NR policy.
- "autostandby agingperiod" must be an integer greater than or equal to 0 and less than 24855 days (596520 hours)
- All paths to <device_id> are dead
- All paths to <volume> are dead
- Array <array> Port <port> is Disabled
- Array <array> Port <port> is Enabled
- Assigned volume <clarion/vnx_volume> to SP<clarion/vnx_storage_processor>
- <path> is in Autostandby, force option required.
- Bad dev value <device>, or not under Powerpath control.
- Bad HBA value <hba>
- Bus hba is Disabled.
- Bus <bus_id> to <storage_system_name> port <port_name> is alive.
- Bus <bus_id> to <storage_system_name> port <port_name> is dead.
- Cannot assemble drivers for root <:>
- Cannot manage. verify state of the class
- Device cannot be managed at this time. Please verify the state of the class <class>
- Cannot open filename
- Cannot remove alive device <device_id>
- Cannot unmanage. Verify state of the class <class>
- Cannot update persistent store
- <storage_system_class> device(s) not found
- Complete management of class <class> failed
- Configuration file created by newer version of powermt
- Device will be configured on closing all the applications to the device or upon reboot
- Get hdiskpower CuAt failed number
- Get hdiskpower CuDv failed number
- Device specified must be a native device.
- Device(s) not found.
- Device not found
- Device(s) not found
- disable/enable is not allowed on this hba type
- Error occurred loading saved driver state from file /etc/powermt_custom.xml Loading continues
- Error occurred loading saved driver state from file <filename>. Loading continues
- Error reading configuration file / custom file IO failure
- Failed to update CLARiON user assignable LUN names
- Length of the filename exceeds the maximum limit of: <number> characters
- Invalid option: dev=all is required for this version
- "autostandby iopflimit" must be a number greater than 1 and less than or equal to 100000000
- Path Bus <bus> Tgt <tgt> Lun <lun> to <storage_system> is set to asb:iopf
- powermt load failed
- Manage of device <device> failed
- Manage of device <device> failed
- Please run powermt save to persist PowerPath settings in custom file
- Performance measurement is not enabled & their powermt display perf command exits
- Performance monitoring already enabled.
- Policy is not a valid policy for <storage_system_class> storage systems
- bosboot: missing proto file: /usr/lib/boot/network/chrp.hdiskpower.proto
- The current boot logical volume, <device>, does not exist on <device>
- powermt save failed
- SCSI error on device <device>

- Latency Threshold exceeded by UserDev...with a new high latency of <latency_value> msec
- Unable to unconfigure path Bus X Tgt Y Lun Z as firstpath policy is set
- Unmanage of device <device> failed
- Renaming of temporary file to file:<filename> failed
- PowerPath cannot disable some of its devices because...two host reboots are required to uninstall PowerPath
- An unidentified program wants to access your computer. Don't run the program unless you know where it's from or you've used it before.

<device_id> is alive

The specified device was killed at some point, due to a critical I/O error.

Explanation

The device was just re-enabled, so PowerPath no longer inhibits I/O to it. The prior I/O error may have caused data on the device to become inconsistent.

Action

Run a data integrity checker to verify the integrity of the data. Fix any inconsistencies.

This must be done before re-enabling the application using the device, to avoid worsening any integrity problem and potentially making the data on the device unrecoverable.

<device_id> is dead

The logical device failed.

Action

Wait for automatic restoration of the logical device or, to avoid delay, run the `powermt restore` command to restore the device.

/etc/powermt: Permission denied

A required file could not be opened.

Explanation

A required file could not be opened or PowerPath attempts to open a device file without required permissions.

If string is a device file, PowerPath is not running or is running without permission to access the device.

Action

Perform one of the following steps:

- Log in as superuser and then run `powermt config`.
- Start PowerPath.
- Run the command as root.

/opt/emcpower/powercf does not exist

By default, the PowerPath installation process sets file permissions such that only the superuser can run `powermt`. (`powercf` is run automatically when the user runs `powermt config`.)

Action

1. Log in as superuser.
2. Run `powermt config`.

/opt/emcpower/powercf: you have to be root to execute this program

By default, the PowerPath installation process sets file permissions such that only the superuser can run `powermt`. (`powercf` is run automatically when the user runs `powermt config`.)

Action

1. Log in as superuser.
2. Run `powermt config`.

Can't get Pirp_t!

There was an error allocating a request buffer.

Action

Contact your Customer Support Representative.

Cannot install the new configuration files

The `emcpupgrade` command could not install new configuration files.

Action

Ensure that the `/etc` directory is not full.

Cannot mount root on <>:<>

Booting from an emcpower device has failed.

Action

Perform one of the following steps:

If booting from an R1 device or an array device:

1. Boot from the native device.
2. Remove the `/kernel/drv/emcp.conf` file.
3. Run `powercf -q` to rebuild the `emcp.conf` file.
4. Run `powermt config`
5. Ensure that the `/etc/vfstab` and `/etc/system` files are pointing to the emcpower device used for the boot.

6. Reboot from the emcpower device.

If booting from an R2 device:

1. Do failback.
2. Boot from the R1 external device.
3. Run **powercfs -z**
4. Run **powermt save**

Cannot obtain lock on <filename>

Multiple `powermt` sessions are occurring; unable to get lock.

Action

Avoid multiple `powermt` sessions.

Cannot obtain driver state

Explanation

Installation: You are making kernel modifications or installing kernel patches post-install.

`powermt`: No storage-system logical devices are configured, one or more PowerPath drivers is not running, or the PowerPath control device file may not exist.

Action

Ensure the driver's device file exists. If it does, make sure the user has permission to access the driver through that file.

Cannot open /dev/emcp, errno <number>

Explanation

- errno 13: By default, the PowerPath installation process sets file permissions such that only the superuser can run `powermt`.
- errno other than 13: A required file cannot be opened for some reason other than access permission.

Action

- errno 13: Log in as superuser and then run **powermt config**.
- errno other than 13: Report the error number to your Customer Support Representative.

Cannot open /etc/path_to_inst. Cannot continue

The `emcpupgrade` command could not open the system configuration file `path_to_inst`.

Action

Repair the system and then run **emcpupgrade** again.

Cannot open <filename>, errno <number>

Explanation

- errno 13: By default, the PowerPath installation process sets file permissions such that only the superuser can run `powermt`.
- errno other than 13: A required file cannot be opened for some reason other than access permission.

Action

- errno 13: Log in as superuser and then run `powermt config`.
- errno other than 13: Report the error number to your Customer Support Representative.

Cannot open configuration file <filename>

The `emcpupgrade` command could not open the specified configuration file.

Action

Perform the following procedure:

1. Change (su) to root.
2. Run `emcpupgrade`.
3. Check that the specified configuration file exists. If the file does not exist, contact your Customer Support Representative.

Cannot open dev dir for reading, errno <number>

Explanation

- errno 13: By default, the PowerPath installation process sets file permissions such that only the superuser can run `powermt`.
- errno other than 13: A required file cannot be opened for some reason other than access permission.

Action

- errno 13: Log in as superuser and then run `powermt config`.
- errno other than 13: Report the error number to your Customer Support Representative.

Cannot open string Initialization error

If string is a device file, PowerPath is not running or is running without permission to access the device.

Action

Perform one of the following steps:

- Start PowerPath
- Run the command as root

Cannot produce new files

The `emcpupgrade` command could not produce new configuration files.

Action

Ensure that the `/tmp` directory is not full.

Cannot remove device <device_id>, name in use by application

An application or mounted filesystem is using the device.

Action

Stop the application or unmount the filesystem that is using the device, and then run `powermt check` or `powermt remove`.

Cannot write to file

The user may not have permission to write to the file.

Action

1. Check the file's write permissions.
2. Log in as superuser if needed.
3. Run `powermt save`.

Cannot write to <filename>

The file is missing, the user does not have the correct access to read the file, or there was a filesystem error.

Action

Perform the following procedure:

1. Check the state of the configuration file and the filenames spelling.
2. Check the files read and write permissions, and log in as superuser if needed.
3. Run `powermt save`.

CLARiiON device path c3t1d0s0 is currently misconfigured

If you remove an iSCSI LUN from a VNX and CLARiiON storage group (but not from the fabric), `powermt check reconfig` sees the paths to the removed LUN as misconfigured and prompts you to reconfigure (that is, remove) them.

Action

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

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

Configuration database corrupted

The `emcpupgrade` command could not read the configuration database.

Action

Contact your Customer Support Representative.

Configuration files contains a dead path (device)

The specified device is present in the configuration files, but cannot be validated using SCSI inquiry.

Action

Reconfigure any application that was configured using this device with the previous installation of PowerPath.

Current/Active controller for volume has changed

The controller that owns the volume has changed. PowerPath was unable to switch to the set of paths attached to the new owning controller.

Action

1. Repair physical paths to the new owning SP.
2. Run `powermt restore` to restore the paths.

Until paths are restored, performance may be degraded significantly.

Device is unavailable

Device is inaccessible (system log may indicate whether message refers to source or target).

Action

Bring problem device back online.

Error number adding pntpath <hdisk_name> for <hdiskpower_name> to CuAt

The system cannot add to the CuAt ODM table.

Action

Remove `hdiskpower` devices and then run the `powermt config` command. Check the ODM state. Verify the root filesystem has sufficient free space.

Error number getting pnp_path attribute for <hdiskpower_name>

The system cannot find the definition for a PowerPath ODM object.

Action

Reinstall PowerPath.

Error number updating CuDv attribute for <device_id>

The system cannot update the CuDv ODM table.

Action

Remove hdiskpower devices and then run `powermt config`. Check the ODM state. Verify the root filesystem has sufficient free space.

Error retrieving minor number(s)

Either the base number you specified is too large or an entry exists in `/devices/pseudo` that corresponds to a free minor number.

Action

Rerun the `emcpminor` command specifying a base number in the allowable range.

If an error still occurs, for each device in `/devices/pseudo`, run the `emcpminor check` command, supplying that device's corresponding minor number. If an error occurs, remove that device from `/devices/pseudo`.

If the error persists, contact your Customer Support Representative.

Failed opening adapter <device_number>

There was a failure opening a bus adapter device.

Action

Check the adapter and bus.

Failover mode PNR for CLARiiON LUN <WWN#> is not recommended. Please consider PAR or ALUA

The failover mode is set to PNR, which is not recommended.

Action

Set the failover mode to either PAR or ALUA.

File deleted. Use -add to create a new file

You chose to delete a corrupt registration file.

Explanation

This message appears after the `EMCP registration file is corrupt` message, if you type **y** to delete the file.

Action

To use PowerPath again, add a registration key.

File not deleted. No licenses will be available until file is deleted and re-created

You chose not to delete a corrupt registration file.

Explanation

This message appears after the `EMCP registration file is corrupt` message, if you type **n** to not delete the file.

Action

No action is required. If you want to delete the corrupt file, run `emcpreg -edit` again.

Get adapter <devno> failed number

The system cannot determine the device number for the hdisk parent.

Action

Remove hdiskpower devices and then run `powermt config`. Check the ODM state

Incomplete installation detected

The `emcpupgrade` command cannot find some files/parameters that should have been set when `pkgadd` was run.

Action

Ensure `pkgadd` completed successfully before you run `emcpupgrade`.

Invalid base minor number

The base minor number specified is not in a valid format (numeric) or is greater than the maximum possible minor number.

Action

Rerun `emcpminor`, specifying as the base minor number a numeric value in the allowable range.

Invalid minor number

The minor number that you specified is not in a valid format (numeric), is greater than the maximum possible minor number, or is thought to be free and has an entry in /devices/pseudo.

Action

Rerun `emcprminor`, specifying as the minor number a numeric value in the allowable range of values.

If an error still occurs, check /devices/pseudo for a device corresponding to that minor number, and remove that device.

If the error persists, contact your Customer Support Representative.

Key <key_number> in license file is bad. Use -edit option to repair

A key in the registration file is invalid. For example, in cases where a key with encryption capability already exists the addition of a new key will result in this message.

Action

Run `emcpreg -edit` and then remove the bad key. Alternately, remove the key by running `emcpreg -remove key_number`

Key <key_number> is invalid

The specified key is invalid or a license with encryption capability was added.

Action

Run the command again, specifying a valid key. The key will display as invalid if the key includes encryption capability.

Key <key_number> was not found

The specified key is not in the registration file.

Action

Enter a different key to be removed.

Key <key_number> is already present, ignored

The specified key is already registered.

Action

Check the key you intended to register, and reenter it if it differs from what you specified.

License file is full -- cannot add number keys

The registration file cannot hold any more keys.

Action

Identify and delete unneeded or unused keys and then add the new key.

mknod of device failed

The `emcpupgrade` command was unable to create the `emcpower` device identified by *device*.

Action

Contact your Customer Support Representative.

MPxIO is enabled. Verify that no arrays are managed by both MPxIO and PowerPath

MPxIO is enabled on the host where you are installing PowerPath. Co-management of devices by PowerPath and MPxIO is not supported.

Action

Disable MPxIO or run `powermt unmanage` to remove the array class being managed by MPxIO from PowerPath management.

Not a multiple of 8

The minor number you specified is not a multiple of 8.

Action

Rerun `emcpminor`, specifying a minor number that is a multiple of 8.

Number of minors out of range 1-1000

The number of minor numbers specified for retrieval is out of the specified range.

Action

Rerun `emcpminor`, specifying as the number of minor numbers to be retrieved a number in the range 1–1000.

One or more expired PowerPath license key(s) detected. These are typically used...use emcpreg utility to review your license keys

Your PowerPath evaluation license has expired.

Action

Use the `emcpreg` command to add your new license key.

PRO service, action 46 failed, 0x5/0x80/0x20

An initiator with no registered Symmetrix path group ID sends a SCSI-3 PERSISTENT RESERVE OUT registration to a LUN.

Explanation

A physical configuration change (such as reconnecting a dead path) causes a bus or device reset, which clears group registrations.

This applies only in Symmetrix or VNX OE block and CLARiiON environments.

Action

No action is required.

An automatic retry register, including a group register, follows this message. In most cases, this restores all paths without intervention.

If cluster problems persists, check the system and array logs for indications of persistent hardware problems.

Pseudo Device Name Major# Minor# <ERROR: unexpected internal error>

The `rendev` command failed to rename the minor number of the pseudo device.

Action

Use the `emcpadm renamepseudo` command instead of the `rendev` command.

Read of file <file> failed

File read failed.

Action

Ensure permissions are set correctly on the specified file.

There are no keys presently registered. Some functions and capabilities may be unavailable

No key is registered; that is, no key was added or upgraded, or the last key was removed.

Action

Add a valid key.

This version of PowerPath is NOT binary compatible with your kernel

You are installing PowerPath when running a kernel version other than the supported kernel. (Or Linux kernel symbol versioning has yielded false results.)

Action

Ensure that the supported kernel is running.

Forceload of drv/sd failed

If the sd driver does not exist on the system, you see this message during boot.

Action

You can safely ignore this message.

Get hdisk path data failed

The system cannot get hdisk path data from the ODM.

Action

Remove hdiskpower devices and then run `powermt config`. Check the ODM state

Internal DB read failed

The configuration database is corrupted.

Action

Contact your Customer Support Representative.

Kernel symbol versioning is not enabled in your linux kernel

You are installing on a kernel that does not have CONFIG_MODVERSIONS enabled. PowerPath will not check for binary compatibility with the kernel.

Action

If checking for binary compatibility of kernel modules is necessary, perform the following steps.

1. Reconfigure the kernel, enabling CONFIG_MODVERSIONS.
2. Rebuild the kernel.
3. Reboot.

Otherwise, continue with the install.

Killing bus <bus> to Symmetrix <id> port <port>

A bus to a storage port is dead.

Action

Wait for automatic restore. Otherwise, repair the failed paths and then run the `powermt restore` command.

Failed to load extension errno = <number>

The system cannot load the specified PowerPath extension.

Action

Check the operating system level and PowerPath version. Verify the extension file is in `/usr/lib/drivers`. Check memory usage.

Recognize (loop) error number on device <device_id>

PowerPath failed to recognize the specified device.

Action

Verify the system is a supported configuration and the EMC ODM package is loaded.

Storage system <storage_system_name> PowerPath mode disabled

PowerPath is not managing the specified VNX and CLARiiON system.

Explanation

PowerPath manages only those VNX OE and CLARiiON systems with PNR (Passive Not Ready) or, with some platforms, ALUA (Asymmetric Logical Unit Access) set.

Action

If you want to manage this storage system with PowerPath, set the PNR or ALUA and LUNZ modes for it, and reboot. Otherwise, no action is required.

Not all paths to <volume> are dead

Some paths to volume are dead.

Action

Wait for automatic restore. Otherwise, repair the failed paths and then run the `powermt restore` command.

SPx Not Responding on array : Trespassing volumes

The current SP is inaccessible. PowerPath issues a trespass to the peer SP and switches to the set of paths attached to the peer SP.

Action

No action is required.

ODM PdDv database contains duplicate uniquetypes. Refer to EMC Knowledge Base Article Number: 000084491 or contact EMC Customer Support personnel

Multiple entries are present for a uniquetype in PdDv.

Action

Contact your Customer Support Representative.

Path <path_name> to <device_id> is dead

A path's state transitioned from alive to dead.

Action

Repair the physical path.

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

The path failed the last path test. If a path is marked dead or the serial numbers encoded in the path configuration information do not match the serial numbers on the logical device, powermt check prompts you to remove the path.

Action

Valid responses are:

- y — removes the dead path and continues checking remaining paths.
- n — does not remove the dead path but continues checking the remaining paths.
- a — removes the dead path and any subsequent paths marked dead.
- q — does not remove the dead path and exits the command. Any paths that were already removed remain removed.

Path <bus> <tgt> <lun> to <device> is dead

A path to a device is dead.

Action

Wait for automatic restore. Otherwise, repair the failed paths and then run the `powermt restore` command.

Open powerpath0 failed number

The system cannot open the powerpath0 device.

Action

Verify that `/dev/powerpath0` exists and has correct permissions.

PRO service action 6 failed, <Sense Key>/<ASC>/<ASCQ> for Bus <bus> Tgt <target> Lun <lun>

The PRO command failed for the specific device shown in the message.

Action

If Sense Key != 6 (unit attention, informational), then report the values to support.

Restored volume <clariion/vnx_volume> to default: SP<clariion/vnx_storage_processor>

The specified volume was trespassed to back to the default storage processor.

Action

No action is required.

Reviving bus <bus> to Symmetrix <id> port <port>

A bus to a storage port has come alive.

Action

No action is required.

Sleeping for memory

There is insufficient system memory.

Action

Check your system configuration and memory use.

Syscall 8133 to <device_id> failed

An RDF control command to the Symmetrix system failed. This message is followed by a buffer dump of the failed call.

Action

Check your configuration and path states. If the message persists, contact your Customer Support Representative.

Trespassed volume <clariion_volume> to SP<clariion_storage_processor>

The current SP is inaccessible. PowerPath issues a trespass to the peer SP and switches to the set of paths attached to the peer SP.

Action

Repair physical paths or the SP. Wait for automatic restoration of the volume or, to avoid delay, run **powermt restore** to restore the volume.

Unable to restore volume <volume> to default controller

An attempt to trespass the volume back to the default SP or controller failed.

Explanation

PowerPath attempts this when **powermt restore** is run or a path to the default SP or controller transitions from dead to alive.

Action

Examine the log file for message 100 or 112. Take the appropriate action for whatever message you find.

Unable to restore volume <clariion_volume> to default SP

An attempt to trespass the volume back to the default storage processor (SP) failed.

Explanation

PowerPath attempts this when **powermt restore** is run or a path to the default SP transitions from dead to alive.

Action

Examine the log file for message 100 or 112. Take the appropriate action for whatever message you find.

Unable to follow volume <volume> to controller

The SP or controller that owns the volume has changed. PowerPath was unable to switch to the set of paths attached to the new owning SP or controller.

Action

Repair physical paths to the new owning SP or controller. Then run **powermt restore** to restore the paths as soon as possible. Until paths are restored, performance may be degraded significantly.

Encountered unexpected error 0x<code> on Dek operation command for <device id> on handle <association_handle>

There was a problem with the DEK Management protocol.

Action

Contact your Customer Support Representative..

Unmanage of device <device> failed

The unmanage of a device will not take effect until all applications using the device are closed or the host is rebooted.

Action

Perform one of the following steps:

- Close all applications using the specified device and re-issue the command.
- Reboot the host.

Getting unmanaged devices failed with error

The command failed to retrieve unmanaged devices.

Action

There may be problems with some of the unmanaged devices. Check the configuration and connectivity of devices.

Unsupported configuration (R1/R2 setup detected)

Upgrade cannot take place on R1/R2-enabled systems.

Action

Contact your Customer Support Representative.

Unsupported HBA

The HBA on the system is not supported by PowerPath.

Action

The E-Lab Interoperability Navigator provides the supported HBAs.

Volume <clariion_volume> followed to SP<clariion_storage_processor>

The SP that owns the volume has changed. PowerPath is switching to the set of paths attached to the new owning SP.

Action

No action is required.

Volume <clariion_volume> is unbound

All paths to the volume are marked dead, because the volume is unbound.

Action

No action is required.

Write failed for lack of space, check for filesystem full

The file system containing /etc does not contain enough free space.

Action

Increase the amount of free space in the file system containing /etc.

Cannot open PowerPath. Initialization error

The PowerPath modules are not loaded. This situation can occur when the PowerPath service (on Linux) is stopped.

Action

To view the modules that are installed, use lsmod (modinfo for Solaris).

If the PowerPath modules are already installed, start the PowerPath service and then retry the command.

Configuration files corrupted/out of date

The configuration files are corrupted. Upgrade is not possible.

Action

Contact your Customer Support Representative.

PowerPath configuration files present - exiting

Explanation

You either:

- Upgraded once and then tried to upgrade again.
- Rebooted the host after you installed PowerPath but before you ran `emcpupgrade`.

Either of these actions would result in a valid PowerPath configuration on the host, and you cannot use `emcpupgrade` to convert files while there is valid PowerPath configuration on the host.

Action

Perform the following procedure:

1. Uninstall PowerPath.
2. Copy the saved configuration files from the archive directory to the appropriate directory.
3. Reinstall PowerPath.

Major number is already in use

When installing PowerPath for a clustered environment, you chose a major number that is already in use. This aborts the installation.

Action

Reinstall PowerPath, specifying a unique major number.

PowerPath modules are not loaded. Ensure that the modules are loaded and try again

The PowerPath modules are not loaded. This situation can occur when the PowerPath service (on Linux) is stopped.

Action

To view the modules that are installed, use `lsmod` (`modinfo` for Solaris).

If the PowerPath modules are already installed, start the PowerPath service and then retry the command. Type `/etc/init.d/PowerPath start`

Refer to the appropriate PowerPath installation and administration guide for steps to install the PowerPath modules.

An older EMCpower package is already installed. Please exit the installation and remove the old EMCpower package first

An older version of PowerPath is already installed on the host. This aborts the installation.

Action

Remove the earlier version of PowerPath and begin the installation again.

PowerPath pseudo bus <bus_name> cannot be configured. Uninstall and reinstall PowerPath

An error is preventing creation of the software bus.

Action

Uninstall and then reinstall PowerPath.

Unable to create kernel process

The system cannot create the PowerPath daemon.

Action

Check file access settings for the daemon and system resource use.

Unable to open /etc/PowerPath Extensions

The system cannot open the configuration file.

Action

Reinstall PowerPath.

Unable to start kernel process

The system cannot create the PowerPath daemon.

Action

Check file access settings for the daemon and system resource use.

EMC registration data is missing or corrupt

The registration file is missing, corrupt, or open in an editor.

Action

Verify the existence of the registration file. If the file is open, close the editor and then run the command again.

EMCP registration file does not exist. Use -add to create it

There is no registration file.

Action

Use the `emcpregr -add` command to add a registration key, which will cause a registration file to be created.

EMCP registration file is corrupt

The license registration file was corrupted and cannot be used.

Action

Contact your Customer Support Representative.

EMCP registration file is corrupt. Delete it?

The license registration file was corrupted and cannot be used.

Action

If desired, type `y` to delete the file, and then add back your registration keys. (Adding the first key will re-create the file.) Otherwise, type `n` and contact your Customer Support Representative.

License for <storage_system_name> storage system support is missing or expired

The PowerPath license is invalid or has expired, and the `powermt` command that was run requires a valid license.

Action

Run `powermt check_registration` to determine the problem with the PowerPath license.

There are no license keys now registered

No key is registered; that is, no key was added or upgraded, or the last key was removed.

Action

Add a valid license key.

There are no license keys now registered

The registration file is missing, corrupt, or open in an editor.

Action

Verify the existence of the registration file. If the file is open, close the editor and then run the command again.

License not installed

The PowerPath license is invalid or has expired, and the `powermt` command that was run requires a valid license.

Action

Run `powermt check_registration` to determine the problem with the PowerPath license.

License not installed

The PowerPath license is invalid or has expired. `powermt` requires a valid license.

Action

Add a valid PowerPath license. Run the `powermt check_registration` command to determine the problem with the PowerPath license.

Temporary upgrade key has been superseded

The temporary upgrade key was removed, because a site-specific key was added.

Action

No action is required.

Key found in old registration file is invalid

The old key is corrupt or is an evaluation key. Evaluation keys are not accepted for automatic upgrade.

Action

Add a valid registration key before installing PowerPath. If needed, get a new license key from your Customer Support Representative.

Mismatch in configuration

A mismatch was detected. The configuration files are not in sync.

Action

Contact your Customer Support Representative.

Missing configuration files

Some of the configuration files are not present. Cannot complete the action.

Action

Contact your Customer Support Representative.

All paths to vol <volume> are dead

Explanation

In a Solaris Cluster with Hitachi Lightning, Hitachi TagmaStore, or HP XP arrays: Unregistered cluster disk paths are dead after the node has booted into the cluster.

In all other configurations: All paths to a volume have failed.

Action

In a Solaris Cluster with Hitachi Lightning, Hitachi TagmaStore, or HP XP arrays: Run the `powermt display` command to check path status. If paths are dead, check the physical connection and then reboot the cluster node.

In all other configurations: Check the physical path state and then run the `powermt restore` command. If `powermt restore` fails to restore the paths, perform a reconfiguration reboot.

Attempted action for <storage_array> storage system support is not allowed with the current license level

The PowerPath license is invalid or has expired. powermt requires a valid license.

Action

Add a valid PowerPath license. Run the `powermt check_registration` command to determine the problem with the PowerPath license.

Attempted action is not allowed for volumes with BF or NR policy.

This message appears when you attempt to disable an HBA or port and the load balancing and failover policy is set to bf or nr.

Explanation

The `powermt disable hba` and `powermt set_port_disable` commands fail when the bf or nr policy is set.

Action

Change to another load balancing and failover policy supported with the storage array, and then retry the command to disable an hba or port.

"autostandby agingperiod" must be an integer greater than or equal to 0 and less than 24855 days (596520 hours)

This message is shown when incorrect value is set for the aging period.

Action

Ensure that the value is within the acceptable range.

All paths to <device_id> are dead

All paths to a logical device are dead. The logical device failed.

Action

Wait for automatic restoration of the logical device or, to avoid delay, run `powermt restore` to restore the device. Otherwise, Repair the physical paths.

All paths to <volume> are dead

All paths to a volume are dead.

Action

Wait for automatic restore. Otherwise, repair the failed paths and then run the `powermt restore` command.

Array <array> Port <port> is Disabled

The specified port has been disabled with the `powermt set port_disable` command.

Action

To enable the port, use the `powermt set port_enable` command.

Array <array> Port <port> is Enabled

The specified port is enabled.

Action

To disable the port, use the `powermt set port_disable` command.

Assigned volume <clariion/vnx_volume> to SP<clariion/vnx_storage_processor>

The current storage processor (SP) is inaccessible. PowerPath issues a trespass to the peer SP and switches to the set of paths attached to the peer SP.

Action

Repair physical paths or the SP. Wait for automatic restoration of the volume or, to avoid delay, run `powermt restore` to restore the volume.

<path> is in Autostandby, force option required.

Attempt to change the mode of a path group, in which some paths are in autostandby mode, has failed. The absence of the `force` option has caused the failure.

Action

Override the autostandby mode with the `force` option.

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

The dev option is invalid.

Action

Run the command again, specifying a valid device name. Use `powermt display dev` to list valid device names.

Bad HBA value <hba>

The HBA option is invalid.

Action

Run the command again, specifying a valid HBA number. Use `powermt display` to list valid HBA numbers.

Bus hba is Disabled.

The HBA has been disabled using the `powermt disable hba` command.

Action

To enable the HBA, use the `powermt enable hba` command

Bus <bus_id> to <storage_system_name> port <port_name> is alive.

The HBA is alive after a path failure or the HBA was enabled with the `powermt enable hba` command.

Action

No action is required.

Bus <bus_id> to <storage_system_name> port <port_name> is dead.

PowerPath proactively killed a path due to a path test failure or the HBA was disabled with the `powermt disable hba` command (Linux).

Action

Repair the physical path and then run `powermt restore` to restore it.

Cannot assemble drivers for root <>:<>

Booting from an emcpower device has failed.

Action

Perform one of the following steps:

- If booting from an R1 device or an array device:
 1. Boot from the native device.
 2. Remove the `/kernel/drv/emcp.conf` file.
 3. Run `powercfg -q` to rebuild the `emcp.conf` file.
 4. Run `powermt config`
 5. Ensure that the `/etc/vfstab` and `/etc/system` files are pointing to the emcpower device used for the boot.
 6. Reboot from the emcpower device.
- If booting from an R2 device:
 1. Do fallback.
 2. Boot from the R1 external device.
 3. Run `powercfg -Z`
 4. Run `powermt save`

Cannot manage. verify state of the class

The status of the specified class is `unmanage_incomplete`.

Action

Before running the `powermt manage class` command, perform one of the following steps:

- Close all the applications using the devices belonging to the storage class and then run the `powermt unmanage class` command.
- Reboot the host.

Device cannot be managed at this time. Please verify the state of the class <class>

The class status is either `unmanaged` or `unmanage_incomplete`.

Action

Perform the following procedure:

1. Run `powermt manage class` to manage the class.
2. Run `powermt manage dev`

Cannot open filename

The file is missing, the user does not have the correct access to read the file, or there was a filesystem error.

Action

Perform the following procedure:

1. Check the state of the configuration file and the filename's spelling.

2. Check the files read and write permissions, and log in as superuser if needed.
3. Run **powermt save**.

Cannot remove alive device <device_id>

An application or mounted filesystem is using the device.

Action

Stop the application or unmount the filesystem that is using the device, and then run `powermt check` or `powermt remove`.

Cannot unmanage. Verify state of the class <class>

The class status is `manage_incomplete`.

Action

Perform the following procedure:

1. Close all applications using the devices of the class.
2. Run **powermt manage class=class**
3. Run **powermt unmanage class=class**

Cannot update persistent store

The `powermt manage class` or `powermt unmanage class` command could not update persistent store.

Action

Perform the following procedure:

1. Run **powermt save**
2. Run **powermt manage class=class** or **powermt unmanage class=class**

<storage_system_class> device(s) not found

The `powermt` command specified a specific class, and no devices were found for the specified type of storage system.

Action

No action is required. If desired, run the command again with a different class value.

Complete management of class <class> failed

The `powermt manage class` command failed.

Action

Perform one of the following steps:

- Stop all applications using the devices belonging to the class and re-issue the command.
- Reboot the host.

Configuration file created by newer version of powermt

There is a mismatch between the version of `powermt` that is running and the version that created the configuration file. Some information in the file may not be valid or recognized.

Action

No action is required.

You can run `powermt config` (to ensure the file is current with respect to the host system's configuration) and `powermt save` (to update the file to the new `powermt` format).

Device will be configured on closing all the applications to the device or upon reboot

The specified device could not be managed with the `powermt manage dev` command.

Action

Perform one of the following steps:

- Stop all applications using the devices belonging to the class and re-issue the command.
- Reboot the host.

Get hdiskpower CuAt failed number

The system cannot get CuAt for an hdiskpower device.

Action

Remove hdiskpower devices and then run `powermt config`. Check the ODM state

Get hdiskpower CuDv failed number

The system cannot get CuDv for an hdiskpower device.

Action

Remove hdiskpower devices and then run `powermt config`. Check the ODM state

Device specified must be a native device.

A device was specified that is not a native device.

Action

Specify a native device.

Device(s) not found.

There is no device specified in the `powermt display perf dev=<dev|all>` or `powermt display perf bus` command.

Action

Specify a device in the `powermt display perf dev=<dev|all>` or `powermt display perf bus` command.

Device not found

The `dev` option is invalid.

Action

Run the command again, specifying a valid device name. Use `powermt display dev=<device>` to list valid device names.

Device(s) not found

The `class=all` option was specified, and no storage-system devices were found.

Action

No action is required.

disable/enable is not allowed on this hba type

The message is returned by the `powermt enable hba` and `powermt disable hba` commands in an unsupported configuration.

Action

No action required.

Error occurred loading saved driver state from file /etc/powermt_custom.xml Loading continues

If the custom file has information of the volumes which are not managed by PowerPath, that is using `powermt unmanage dev` or after removing luns from the host, the error message is displayed.

Action

Contact your Customer Support Representative.

Error occurred loading saved driver state from file <filename>. Loading continues

The configuration file is corrupted, or there was an I/O error adding a device in the file.

Explanation

There is a problem reading the configuration file, because of insufficient permissions, corrupted configuration file, or there was an I/O error while adding a device in the file.

Action

Check the file permissions. Check validity of the configuration file. Contact your Customer Support Representative.

Error reading configuration file / custom file IO failure

An error occurred while reading the configuration file.

Action

Run **powermt load** again. Also check the files read and write permissions, and log in as superuser if needed.

Failed to update CLARiiON user assignable LUN names

The command failed.

Action

Contact your Customer Support Representative.

Length of the filename exceeds the maximum limit of: <number> characters

Filename length has exceeded the set limit.

Action

Provide a filename within the expected length limitation.

Invalid option: dev=all is required for this version

A device ID value is specified for the device option in the `powermt display perf bus` command.

Action

The `device=<device ID>` option is not supported because the default setting is `device=all`. Use the `powermt display perf bus` command with the `device=all` option.

"autostandby iopflimit" must be a number greater than 1 and less than or equal to 100000000

An incorrect value is set for the average number of I/Os between failures.

Action

Ensure that the value is within the acceptable range.

Path Bus <bus> Tgt <tgt> Lun <lun> to <storage_system> is set to asb:iopf

The logging for IOsPerFailure-based autostandby (`asb:iopf`) identifies why the paths are moved to standby.

Action

None. This message is informational.

Set to asb:iopf

```
Mar 3 16:41:59 libj164 emcp: [ID 801593 kern.notice] Info: Path Bus 3074 Tgt 500014424000CF11
Lun 8 to VPLEX ID FCNCH092905441 is set to asb:iopf.
Mar 3 16:41:59 libj164 emcp: [ID 801593 kern.notice] Info: Path Bus 3074 Tgt 500014424000CF11
Lun 8 moved to iopf with following error sequence
Mar 3 16:41:59 libj164 emcp: [ID 801593 kern.notice] Info: Path Bus 3074 Tgt 500014424000CF11
Lun 8, ErrInst 1150a, TimeStamp 2015 Mar 03 11:11:19 GMT, ErrCode SCSI IO Error
Mar 3 16:41:59 libj164 emcp: [ID 801593 kern.notice] Info: Path Bus 3074 Tgt 500014424000CF11
Lun 8, ErrInst 1150b, TimeStamp 2015 Mar 03 11:11:19 GMT, ErrCode SCSI IO Error
...
```

```
PowerPath:Path vmhba2:C0:T2:L0 to VNX ID APM00141314672 is set to asb:iopf.
PowerPath:Path vmhba2:C0:T2:L0 moved to iopf with following error sequence
PowerPath:Path vmhba2:C0:T2:L0, ErrInst cbe8, TimeStamp 2015 Feb 05 06:34:59 GMT, Errcode
Host Retry(0xc)
PowerPath:Path vmhba2:C0:T2:L0, ErrInst cbe9, TimeStamp 2015 Feb 05 06:34:59 GMT, Errcode
Host Retry(0xc)
PowerPath:Path vmhba2:C0:T2:L0, ErrInst cbea, TimeStamp 2015 Feb 05 06:34:58 GMT, Errcode
Host Retry(0xc)
```

powermt load failed

This can occur after an upgrade. The loading of previously configured PowerPath settings failed.

Action

The `powermt-custom` file, from which the load is attempted, is saved in the format `/etc/powermt.custom.yymmdd:hhmmss.load_failed` (for example, `/etc/powermt.custom.030408:142021.load_failed`). Try to load this file manually by running `powermt load` from the command line. If this fails, contact your Customer Support Representative.

Manage of device <device> failed

The specified device could not be managed with the `powermt manage dev` command.

Action

Perform one of the following steps:

- Stop all applications using the devices belonging to the class and re-issue the command.
- Reboot the host.

Manage of device <device> failed

The `powermt manage` command failed to place the specified device under PowerPath control.

Action

Check the status of the class using the `powermt display options` command. If the status is `manage_incomplete` or `unmanage_incomplete`, the status will be updated after you close all applications using devices belonging to the class, or after you reboot the host. You can then re-issue the `powermt manage dev` command.

For other statuses, the device might be unreachable.

Please run powermt save to persist PowerPath settings in custom file

After running `powermt set persist_path_info save` the value to the custom file.

Action

Run `powermt save` to save the persistence value to the custom file.

Performance measurement is not enabled & their powermt display perf command exits

The `powermt display perf` command is run when performance monitoring is disabled by another user.

Action

Before enabling and using the performance monitoring, check the required settings of your environment and then contact the user who has disabled the performance monitoring regarding why monitoring is disabled. Enable performance monitoring.

Performance monitoring already enabled.

The `powermt set perfmon` command is run on the system where the performance monitoring is already enabled.

Action

Perform the following steps:

1. Run `powermt set perfmon=off` to disable performance monitoring.
2. Run `powermt set perfmon=on interval=<#seconds>` to change the performance monitoring interval.

Policy is not a valid policy for <storage_system_class> storage systems

The specified load-balancing and failover policy is invalid for the specified class of storage system.

Action

Run `powermt set policy` again, specifying a valid policy and storage-system class combination.

bosboot: missing proto file: /usr/lib/boot/network/chrp.hdiskpower.proto

After a system boots from a PowerPath device, the bosboot tool cannot function correctly.

Explanation

This is because of the state of the configuration after booting from a PowerPath device and the fact that bosboot expects the boot device to be an hdisk, not an hdiskpower device.

Action

Run `pprootdev fix` after the system boots using a PowerPath device.

The current boot logical volume, <device>, does not exist on <device>

The `pprootdev fix` and `pprootdev off` commands may display the warning message.

Action

You can ignore this message; it does not reflect an actual problem. You need to invoke each of these `pprootdev` commands only once.

powermt save failed

The attempt to save currently configured PowerPath settings failed. This can occur after an upgrade.

Action

Run `powermt save`. If that fails, ensure that the `powermt_custom.xml` file exists and check its permissions; if the file is bad, remove it and run `powermt save` again.

SCSI error on device <device>

Could not perform SCSI inquiry on the device. The device might be offline.

Action

Check that all devices are online. Bring any offline devices online. If some of the paths have been removed from the system but not from the previously installed version of PowerPath:

1. Uninstall the new package.
2. Reinstall the old package.
3. Remove the dead paths from PowerPath by running the `powermt remove dev` command.
4. Reinstall the new package.

Latency Threshold exceeded by UserDev...with a new high latency of <latency_value> msec

The I/O completion time has exceeded the threshold value.

Action

Estimate expected latencies for your environment and set the value suitable for your environment. To estimate, check the `powermt display latency` output.

Unable to unconfigure path Bus X Tgt Y Lun Z as firstpath policy is set

You tried to remove the first path an application opened to a multipathed volume.

Explanation

In certain situations (for example, when early versions of VxVM are running without DMP), removing this first path can cause data corruption and thus is prohibited.

Action

If you are sure you want to remove the path, rerun the `powermt remove` command using the `firstpath` option: `powermt remove firstpath dev=device`.

Unmanage of device <device> failed

A failure occurred while unmanaging the specified device.

Action

Check the status of the class using the `powermt display options` command. If the status is `manage_incomplete` or `unmanage_incomplete`, the status will be updated after you close all applications using devices belonging to the class, or after you reboot the host. You can then re-issue the command.

For other statuses, the devices might be unreachable.

Renaming of temporary file to file:<filename> failed

The file is missing, the user does not have the correct access to read the file, or there was a filesystem error.

Action

Perform the following procedure:

1. Check the state of the configuration file and the filenames spelling.
2. Check the files read and write permissions, and log in as superuser if needed.
3. Run `powermt save`.

PowerPath cannot disable some of its devices because...two host reboots are required to uninstall PowerPath

A page file located in a LUN is under PowerPath control or the host is configured to boot from SAN.

Action

Two host reboots are required. Click **OK** to proceed, or **Cancel** to exit.

An unidentified program wants to access your computer. Don't run the program unless you know where it's from or you've used it before.

This message displays (at the UAC prompt) when a user attempts to uninstall PowerPath (although the user installed a digitally signed copy of PowerPath setup program).

Action

No action is required.

Migration Enabler messages

This chapter contains the following topics:

Topics:

- Already recovered
- Already in the state requested
- Couldn't read disk label
- Couldn't write disk label, be sure target is at least as large as source
- Current node cannot be excluded
- Daemon communication error
- Daemon failed to create thread
- Daemon semaphore error
- Daemon socket connection failed
- Device is inaccessible
- Devices already involved in a migration
- An error occurred while using the base daemon, please see the log file
- An internal error occurred during a SYMAPI operation. Please report to EMC
- Failed to fix disk label of the target device
- Failed to get disk layout of device
- Failed to get target size
- Failed to load EFI support library
- Failed to set disk layout of device
- Failed to validate geometry of target device
- Failed to transfer the source reservation
- The feature being requested is not currently licensed
- A function call exceeded allocated time
- The gatekeeper device (while using the Base Daemon) has an error (Please see the log file)
- Handle not found
- Host can still access source device
- Migration stopped due to IO error - will attempt to restart
- Needs recovery
- No migration is in the proper state to perform the requested operation
- Node <nodename> not found
- PPME Resource <Cluster Disk Name> created with powermigcl
- PR Keys found on Target disk during Setup
- Not in proper state to perform this operation
- Possible alternate cylinder overlap detected using native source geometry
- PPME cannot be used with a temporary license
- Remote SE environment variable(s) not allowed in setup
- Source must be a cluster shared volume
- Source is a CSV (use '-cluster mscsv' to setup a cluster migration)
- Source disk's dependent PPME resource was changed
- Source device is in use (-force will override)
- Source faulted, must abort migration
- Source faulted, must commit migration
- Source and target must be from same array
- Source and target are the same device
- Target faulted, must abort migration
- Target device is in use
- Target must be at least as large as the source

- Target max_transfer attribute is too small
- Target is not accessible
- Target device is not attached to some nodes
- Target must be a pseudo
- UMD database not initialized
- Unable to scrub the abandoned device
- Unable to scrub the abandoned device. Device is offline
- Unknown device

Already recovered

The `powermig recover` command was run and the migration is not in needsRecovery state.

Action

Perform the following procedure:

1. Run the `powermig info` command to identify the state of the migration.
2. Resume the migration at the appropriate point in the migration process.

Already in the state requested

The migration is in the state to which you tried to transition.

Action

Proceed to the next step in the migration process.

Couldn't read disk label

This applies only to Solaris and occurs when the source or target EFI or VTOC label cannot be obtained. The system error log displays whether it was a source or target and EFI or VTOC problem.

Action

Use the `format` or `powerformat` utility to verify that the source or target logical unit involved in the migration has a valid disk label and are of the same label type (EFI or VTOC).

Couldn't write disk label, be sure target is at least as large as source

This applies only to Solaris and occurs when a copy of the label fails from source to target during `powermig setup`. The copy could fail because the target is not large enough to accept the label.

Action

Use the `format` or `powerformat` utility to verify the target logical unit is accessible, is as large as the source, and has a valid disk label of the same type as the source.

Current node cannot be excluded

Nodes cannot be excluded if they have migration devices connected during a migration.

Action

Select another node to exclude, if needed.

Daemon communication error

This applies only to Host Copy and occurs when read or write fails to migration daemon socket (UNIX) or migration Service named pipe (Windows).

Action

Perform one of the following steps:

- Run **powermig recover** to retry the failed command.
- Check that the emcpmigd daemon or the Service is running.
- If the problem persists, contact your Customer Support Representative.

Daemon failed to create thread

This applies only to Host Copy and occurs during the syncing state when thread creation fails in the migration daemon.

Action

Perform one of the following steps:

- Run **powermig recover** to retry the failed command.
- If the problem persists, contact your Customer Support Representative.

Daemon semaphore error

This applies only to Host Copy and occurs during the syncing state when Migration Enabler is unable to obtain semaphore for obtaining Host Copy copy slot semaphore.

Action

Perform one of the following steps:

- Run **powermig recover** to retry the failed command.
- If the problem persists, contact your Customer Support Representative.

Daemon socket connection failed

This occurs during the syncing state when failing to set up connection to the migration daemon socket (UNIX) or migration Service named pipe (Windows).

Action

If the problem persists, contact your Customer Support Representative.

Device is inaccessible

Migration Enabler is not able to access the source or target device.

Action

Review the system error log to determine which device is inaccessible and to verify that the logical unit is not opened exclusively by another application.

Devices already involved in a migration

The source or target (pseudo or native) is being used in another migration. A logical unit can only be involved in one migration at a time.

Action

Perform one of the following steps:

- Review the system error log to determine which device is part of another migration.
- Ensure you entered the correct device name.

An error occurred while using the base daemon, please see the log file

This applies to Open Replicator and TimeFinder/Clone.

Action

The Solutions Enabler documentation provides additional information.

An internal error occurred during a SYMAPI operation. Please report to EMC

During a SYMAPI operation with Solutions Enabler, the host has crashed causing the symapi database corruption.

Action

Add the following line in the symapi options file (/var/symapi/config/options) to flush all the symapi database writes to disk:

```
SYMAPI_DB_FSYNC_MODE = ENABLE
```

Failed to fix disk label of the target device

Migration Enabler did not update the ASCII text label or disk signature, which includes the disk vendor, product, and revision.

Explanation

This update is usually done as part of a migration when the migration is a heterogeneous migration in a Solaris host environment.

Action

Perform one of the following steps:

- Use the `powermig cleanup -force` command to ignore the failure.
- For Solaris, use the `powerformat` utility to update the ASCII disk label information.
- Review the system error log to determine which error is the cause.

Failed to get disk layout of device

Explanation

For Solaris hosts, after reading the label during the setup state, the disk geometry cannot be obtained.

For Linux and Windows hosts, this may occur during setup if the disk geometry or disk size cannot be retrieved, or during cleanup (without the `-format` option), if the disk size cannot be retrieved as part of erasing the disk layout.

Action

For Solaris hosts, use the `format` or `powerformat` utility to verify the disk label.

For Linux and Windows hosts, verify the disk is accessible.

Failed to get target size

Unable to obtain the size of the target.

Explanation

For Solaris hosts, there was a failure to obtain the size of the target from SCSI `read_capacity` IOCTL.

Action

Check the accessibility of the target.

Failed to load EFI support library

Explanation

For Solaris hosts, the EFI shared library failed to open.

For AIX hosts, this may occur after running `powermig cleanup` when opening attributes class in ODM.

Action

Contact your Customer Support Representative.

Failed to set disk layout of device

Explanation

For Solaris hosts, unable to set the target disk geometry during setup.

For Windows hosts, this may occur during setup if the disk geometry or disk size cannot be retrieved, or during cleanup (without the `-format` option), if the disk size cannot be retrieved as part of erasing the disk layout.

Action

For Solaris hosts, use the `format` or `powerformat` utility to verify the disk label.

For Windows hosts, verify the disk is accessible.

Failed to validate geometry of target device

During setup, after the source VTOC label is copied to the target, the target geometry validation fails. Such as when Migration Enabler is unable to determine that the source label matches the target.

Action

Perform one of the following steps:

- If you see this message when setting up a migration, re-run the `powermig setup` command (the problem may be intermittent).
- If the error still occurs, provision a larger target device.

Failed to transfer the source reservation

Explanation

For Windows hosts, if the target device has PR key already assigned to it, `powermig setup` will throw this error message.

Action

For Windows hosts, please clear the PR Key in the target device manually to proceed.

The feature being requested is not currently licensed

This is not a Migration Enabler licensing error. The technology that you are using with Migration Enabler is not licensed.

Action

Ensure that all technologies that you use with Migration Enabler are licensed. The documentation for the technology provides more information.

A function call exceeded allocated time

This applies only to Open Replicator.

Action

The Solutions Enabler documentation provides additional information.

The gatekeeper device (while using the Base Daemon) has an error (Please see the log file)

This is a Solutions Enabler error that applies to Open Replicator and TimeFinder/Clone.

Explanation

You may need to modify the semaphore parameters for SYMCLI. The semaphore parameters control access to the database file and gatekeepers.

Action

The Solutions Enabler documentation provides additional information.

Handle not found

The specified handle could not be found in any migration in the UMD.

Action

Perform one of the following steps:

- Verify you are entering the correct handle for the source and target device-name pair.
- Use the `powermig getHandle` command to determine the correct handle. However, once a migration has been cleaned up, the handle is no longer valid.

Host can still access source device

This applies only to Invista encapsulation and occurs when the source is still available to the host after successful encapsulation.

Action

Remove host access to the source logical unit before running the `powermig cleanup` command.

Migration stopped due to IO error - will attempt to restart

This applies only to HostCopy and will occur when there is no application I/O and the source or target is unavailable.

Action

Resolve the I/O problem or abort the migration.

Needs recovery

A process crash or error occurred during the execution of a `powermig` command, and the command needs to be recovered.

Action

Resolve the problem and then run the `powermig recover` command.

No migration is in the proper state to perform the requested operation

This may occur when using the `-a11` option if there is no migration in a valid state for the command.

Action

Choose a different command that is valid for the state of at least one or more of the migrations.

Node <nodename> not found

The excluded node is not present in the cluster or not reachable.

Action

None.

PPME Resource <Cluster Disk Name> created with powermigcl

The Migration Enabler Cluster resource for the cluster disk was created using the `powermigcl` utility rather than `powermig`.

Action

None. Informational.

PR Keys found on Target disk during Setup

A stale Persistent Reservation key was found on target disk at the setup of the migration and this needs to be cleared in order to proceed with migration.

The complete message is as follows.

```
PR Keys found on Target disk during Setup: Target Disk is <TargetDisk used for migration> Source Disk is <SourceDisk used for migration> <TargetDisk>s has PR key : <PersistantReservation Key of the Target> please clear the key manually to proceed.
```

Action

Clear the keys on the Target disk using the Microsoft ClearPr PowerShell command or by another means.

Not in proper state to perform this operation

The migration is not in the required state for the command being executed.

Action

Perform one of the following steps:

- To determine the state of the migration, run the `powermig info` or `powermig query` command.
- Run the proper command that is valid for the current state of the migration.

Possible alternate cylinder overlap detected using native source geometry

On Solaris, you may see the following warning message in a heterogeneous migration when the disks are close in size.

Action

None. The migration succeeds.

PPME cannot be used with a temporary license

Temporary licenses cannot be used with Migration Enabler.

Action

Register Migration Enabler using a permanent license.

Remote SE environment variable(s) not allowed in setup

This may occur during the setup state if either environment variable `PPME_REMOTE_SE_SERVER` or `PPME_REMOTE_SE_PORT` is set.

Action

Perform one of the following steps:

- Use the `-techHost` or `-techPort` command line option during setup.
- Delete the environment variables.

Source must be a cluster shared volume

The `-cluster mscsv` option could be used only for migrating cluster shared volumes.

Action

Use the `-cluster` or `-cluster mscd` option if you are not migrating cluster shared volumes.

Source is a CSV (use '-cluster mscsv' to setup a cluster migration)

The `-cluster` option is used to setup a migration for Cluster Disk.

Action

To setup a migration for CSV use `-cluster mscsv`.

Source disk's dependent PPME resource was changed

It was detected that the Migration Enabler cluster resource for the migration was changed or renamed.

Action

Restore the original Migration Enabler cluster resource.

Source device is in use (-force will override)

This may occur after running the `powermig undoRedirect` command if the source device is still in use.

Action

Perform the following procedure:

1. Bring down any application sending I/O to the target.
2. Run the `powermig undoRedirect` command.
3. Reconfigure the application to use the target device name.

Otherwise, use the `-force` option to stop the redirection of I/O from the source to the target even if the source device is in use.

Source faulted, must abort migration

This occurs during the syncing state when a source logical-unit fault occurred before the synchronization completes.

Action

Run the `powermig abort` command to abort the migration. After you resolve the fault situation, restart the synchronization.

Source faulted, must commit migration

This occurs during the `sourceSelected` or `targetSelected` state when a source logical-unit fault occurred after the synchronization completes.

Action

Perform one of the following steps:

- If you are in the `sourceSelected` state, run the `powermig selectTarget` command and then the `powermig commit` command.
- If you are in the `targetSelected` state, run the `powermig commit` command.

Source and target must be from same array

This applies only to TimeFinder/Clone and occurs when the devices specified are from different arrays. TimeFinder/Clone requires that devices are from the same Symmetrix.

Action

Ensure that the source and target names point to devices on the same array.

Source and target are the same device

The same device name was specified as the source and target when you set up the migration (`powermig setup`).

Action

Ensure that the source and target names point to different devices, and then retry the migration setup.

Target faulted, must abort migration

This occurs during the syncing, `sourceSelected`, or `targetSelected` state when a target logical-unit fault occurred.

Action

Run the `powermig abort` command to abort the migration. After you resolve the fault situation, restart the migration session.

Target device is in use

This occurs during setup when the target is accessed by an application.

Action

Perform one of the following steps:

- Ensure that no application is accessing the target. Such as in a VxVM environment.
- On UNIX hosts, run `fuser` to discover which process is holding open the target.
- Use `-force` if you are sure that the device can be used for a migration.

Target must be at least as large as the source

During setup, the target is smaller than the source.

Action

Specify a target that is the same size or larger than the source.

Target max_transfer attribute is too small

This occurs during setup when the `max_transfer` attribute of the target is smaller than that of the source.

Action

Modify the target's `max_transfer` size to be at least equal to that of the source.

Target is not accessible

This occurs before transitioning to the commit state if Migration Enabler cannot read one block from the target. The migration remains in the `targetSelected` state.

Action

Perform one of the following steps:

- Check the availability of the target.
- If the problem cannot be resolved, run `powermig abort` to abort the migration.

Target device is not attached to some nodes

It was detected that target device is not attached to some nodes.

Action

Attach target device to all nodes and retry.

Target must be a pseudo

A native device name was entered as the target, and a pseudo device name was entered as the source in a `powermig setup` command. Pseudo-named to native-named device migrations are not supported.

Action

Specify a pseudo target device.

UMD database not initialized

This may occur on any `powermig` command if the userspace metadata (UMD) is not opened properly.

Action

Contact your Customer Support Representative.

Unable to scrub the abandoned device

This occurs during cleanup if there is a problem scrubbing the abandoned device.

Action

Review the system log for more detail.

Unable to scrub the abandoned device. Device is offline

This occurs during cleanup if there is a problem scrubbing the abandoned device. Device is offline or cannot be opened.

Action

Check the state of the abandoned device.

Unknown device

This may occur during setup when the source or target device name entered cannot be found.

Action

Enter a valid device name, and then retry the command.

PowerPath/VE messages

This chapter contains the following topics:

Topics:

- Allocation of dummy <ASyncloDaemon/ TestDaemon/ SyncloDaemon> entry failed. Not aborting it
- Bad dev value <invalid_value>, or not under PowerPath control
- Bad dev value emcpower <invalid value>, or not under PowerPath control
- Bad hba value <invalid_value>, or not under PowerPath control
- Cannot communicate with host, version mismatch too great
- Cannot communicate with host, version mismatch too great
- Could not allocate SCSI device
- Could not claim the path. Status: <status>
- Could not create vol Daemon
- Could not initialize log name: <status>
- Device <device_bus_path> is currently dead
- Device(s) not found
- <array_type> device(s) not found
- <class> device(s) not found
- Device(s) not found
- Enqueued from <ASyncloDaemon/ TestDaemon/ SyncloDaemon> to Abort queue Cmd <command>
- Failed to Create Pirp
- Failed to create Pirp
- Failed to create Platform scsicmd
- Failed to create scsi Pirp
- Failed to find IP address for <unresolved hostname>
- Failed to open host file
- Failed to remove host information from host file
- Failed to reset the license server count
- Failed to retrieve host information from host file
- Failed to retrieve host list from host file
- Failed to store host information into host file
- Failed to store host list into host file
- Failed to alloc Pirp
- Failed to create Char device name. Status: <status>
- Global PMI disabled
- Host not found
- Inquiry failed. Status: <status>
- Invalid option '-d' when upgrading from a previous version of installer
- Invalid passphrase during host file creation
- License autoregistration is disabled
- License checkout failed
- License data does not match host data
- License feature or attribute not found
- License file does not exist or not found
- License file format error or misspelling
- License file is invalid (inconsistent)
- License key or feature not found
- License not installed
- License not installed
- License type mismatch. Hyper-V enabled: PowerPath/VE for Windows license key not found

- Mpx: pData <data> size <size>
- MpxRecognize failed. No resources- check system maximum path limits
- MpxRecognize failed
- No valid licenses found
- No active path to the device found
- No PMI data available for this VM
- No standard UID. Using the WWN as volume UID
- Not a valid sense buffer
- The path has already been claimed
- Path is in Autostandby, force option required
- Path <path> could not be disabled. It is the last working path to volume
- The Path has not been claimed
- Performance Monitor is not enabled
- Persistence of SCSI Res fail Value is 0, generating new one
- PowerPath driver is already registered
- PowerPath driver is already unlicensed
- PowerPath Not Found
- Probe of <path> path failed. Status <status>
- ProbePath Failed path
- Path <path>: could not issue sync/async IO command: <status>
- SCSI-2 reservation/release failed with SCSI status <status>
- SCSI-2 Translation command failed after the maximum number of retries
- Sense Info S: <sc> ASC: <asc> ASCQ: <ascq>
- Status <status>. Failed to register device. User should unclaim paths.
- Unable to send cmd (cleanup PRI) to device. SCSI status <status>
- Unable to send cmd (cleanup PRO) to device. SCSI status <status>
- Could not unregister the device
- User Authentication Invalid
- VM either does not exist or is not registered on this host
- vmk_LogRegister failed: <status>
- Volume is being destroyed. Cannot open
- Volume not found
- Volume to be opened not found
- The volume is open. Cannot unclaim last path <path>
- The volume is being registered or unregistered. Cannot unclaim last path <path>
- Wait for n losPending count (count) to become 0
- Wait for volume to be unquiesced

Allocation of dummy <ASyncloDaemon/ TestDaemon/ SyncloDaemon> entry failed. Not aborting it

Task management operation has failed.

Action

None. PowerPath/VE will automatically retry the command.

Bad dev value <invalid_value>, or not under PowerPath control

A non-alpha character was specified.

Action

Specify a valid value.

Bad dev value emcpower <invalid value>, or not under PowerPath control

An invalid device number was specified.

Action

Specify a valid value.

Bad hba value <invalid_value>, or not under PowerPath control

An invalid HBA number was specified.

Explanation

For values greater than 2^{31} , the overflowed value is indicated in the error message.

Action

Specify a valid value.

Cannot communicate with host, version mismatch too great

The RTOOLS version is not compatible with the PowerPath/VE version on the host.

Action

Check RTOOLS version compatibility and then upgrade the RTOOLS package, if necessary.

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.

Action

None.

Could not allocate SCSI device

Maximum number of SCSI devices, 256, has already been claimed.

Action

PowerPath/VE load failed. Check if enough memory is available.

Could not claim the path. Status: <status>

PowerPath/VE failed to claim path.

Action

None. PowerPath/VE will automatically retry the command.

Could not create vol Daemon

PowerPath/VE failed to create a daemon or to allocate memory.

Action

Check if server is out of memory.

Could not initialize log name: <status>

Initializing the log handle name failed.

Action

PowerPath/VE load failed. Check if enough memory is available.

Device <device_bus_path> is currently dead

New dead paths were found.

Action

Refer to `rpowermt restore` for more information.

Device(s) not found

A management command (for example, a set command) was issued, but no devices are being managed by PowerPath.

Action

None.

<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.

Action

None.

<class> device(s) not found

Both the `class` and `dev` options were used, but the specified values do not match a currently configured PowerPath volume and device path.

Explanation

A management command (for example, a set command) was issued for a vSphere host, but no devices were configured for the specified class.

Action

None.

Device(s) not found

There is no data for the device or there is no device specified in the `rpowermt display perf dev=device|all` or `rpowermt display perf bus` commands.

Action

Specify a device in the `rpowermt display perf dev=device|all` or `rpowermt display perf bus` commands.

Enqueued from <ASyncloDaemon/ TestDaemon/ SyncloDaemon> to Abort queue Cmd <command>

PowerPath/VE has aborted a command.

Action

None. PowerPath/VE will automatically retry the command.

Failed to Create Pirp

PowerPath/VE failed to create I/O request.

Action

Check if enough memory is present on the server.

Failed to create Pirp

Insufficient memory on the system.

Action

Check if enough memory is present on the server.

Failed to create Platform scsicmd

Insufficient memory on the system.

Action

Check if enough memory is present on the server.

Failed to create scsi Pirp

Insufficient memory on the system.

Action

Check if enough memory is present on the server.

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.

Action

None.

Failed to open host file

The `rpowermt setup verify_hosts` command verified that a lockbox exists but specified host is not present or the `rpowermt setup list_hosts` command could not obtain a lock on the specified lockbox.

Explanation

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

Action

None.

Failed to remove host information from host file

An encryption-specific error occurred. Such errors, while extremely rare, may prevent rpowermt from properly managing the lockbox.

Action

None.

Failed to reset the license server count

License server is not responding or unreachable.

Action

Repeat the command with the `force` option after communication to the license server is restored.

Failed to retrieve host information from host file

An encryption-specific error occurred. Such errors, while extremely rare, may prevent rpowermt from properly managing the lockbox.

Action

None.

Failed to retrieve host list from host file

An encryption-specific error occurred. Such errors, while extremely rare, may prevent rpowermt from properly managing the lockbox.

Action

None.

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.

Action

None.

Failed to store host list into host file

An encryption-specific error occurred. Such errors, while extremely rare, may prevent rpowermt from properly managing the lockbox.

Action

None.

Failed to alloc Pirp

Insufficient memory on the system.

Action

Check if enough memory is present on the server.

Failed to create Char device name. Status: <status>

PowerPath/VE character device could not be created.

Action

PowerPath/VE load failed. Check if enough memory is available.

Global PMI disabled

You attempt to enable VMPMI when PMI is not enabled.

Action

1. Enable performance monitoring.

```
rpowermt set perfmon
```

2. Enable VM performance monitoring.

```
rpowermt set vm_perfmon
```

Host not found

An 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.

Action

None.

Inquiry failed. Status: <status>

The inquiry failed.

Action

Check if the device is accessible.

Invalid option '-d' when upgrading from a previous version of installer

The `elms_install` option that you used to upgrade ELMS was incorrect.

Explanation

The command options change based on the ELMS version that you have installed. The ELMS versions can be rpm based (11.10 or later) or non-rpm based (earlier than 11.10) and therefore might require a different option.

Action

Follow the procedure in the *PowerPath/VE for VMware vSphere Installation and Administration Guide* when upgrading the ELMS.

Invalid passphrase during host file creation

The passphrase entered for a lockbox did not meet the passphrase requirements.

Explanation

The passphrase requirements are:

- Be eight or more characters in length.
- Contain at least one numeric, one uppercase, one lowercase character, and one non-alphanumeric character (such as # or !).

Action

Type a valid passphrase.

License autoregistration is disabled

The `PPMT_DISABLE_LICAUTO` environment variable is set to 1 thereby disabling autoregistration on an unlicensed host.

Action

Remove the environment variable to enable automatic license registration.

License checkout failed

The `rpowermt register` command is located the license file, but the attribute value within the license file is corrupted.

Action

None.

License data does not match host data

License information in persistent string does not match host hardware data.

Action

Persisted license data is not valid on this host. Add a new PowerPath/VE License.

License feature or attribute not found

There is a clock skew of more than 48 hours between the `rpowermt` server and PowerPath/VE Electronic License Manager System (ELMS) or the ESX host.

Action

Synchronize the `rpowermt` server and the PowerPath/VE ELMS (for served licenses) or the ESX host (for unserved licenses) within 48 hours of each other.

License file does not exist or not found

The `rpowermt register` command could not locate the license file.

Action

None.

License file format error or misspelling

The `rpowermt register` command is located the license file, but the license file is corrupted.

Action

None.

License file is invalid (inconsistent)

The `rpowermt register` command is located the license file, but the file did not contain a key for the specified vSphere host.

Action

None.

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.

Action

None.

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.

Action

None.

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.

Action

None.

License type mismatch. Hyper-V enabled: PowerPath/VE for Windows license key not found

The specified key that was installed on a Hyper-V server is a license for physical platforms.

Action

Please contact EMC, obtain a PowerPath/VE for Windows license key, and then install it. Remove the physical platform license from the server.

Mpx: pData <data> size <size>

Invalid ioctl buffer size.

Action

Check if kernel module version matches mpapi version.

MpxRecognize failed. No resources- check system maximum path limits

Maximum allowed path limit has already been reached.

Action

Verify that the number of devices configured on this setup is fewer than the maximum supported device/path limit.

MpxRecognize failed

PowerPath/VE failed to claim path.

Action

None. PowerPath/VE will automatically retry the command.

No valid licenses found

There is a clock skew of more than 48 hours between the rpowermt server and PowerPath/VE Electronic License Manager System (ELMS) or the ESX host.

Action

Synchronize the rpowermt server and the PowerPath/VE ELMS (for served licenses) or the ESX host (for unserved licenses) within 48 hours of each other.

No active path to the device found

No active path to the dump device was found.

Action

Check the paths to storage array that has the dump device.

No PMI data available for this VM

You need to wait for two sample intervals after VMPMI (performance monitoring) is enabled before requesting VMPMI data.

Action

Rerun the `rpowermt display perf vmname` command after two sample intervals.

No standard UID. Using the WWN as volume UID

Standard UUID for device could not be found.

Action

None.

Not a valid sense buffer

Not a valid sense buffer for I/O status.

Action

None. PowerPath/VE will automatically retry the command.

The path has already been claimed

The path has already been claimed by another multipathing plugin.

Action

Unclaim this path from the other multipathing plugin before trying to claim the path using PowerPath/VE. The *PowerPath/VE for VMware vSphere Installation and Administration Guide* provides more information.

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.

Action

Use the `force` option.

Path <path> could not be disabled. It is the last working path to volume

PowerPath/VE failed to change path state.

Action

Ensure that the path that you are trying to disable is not the last active path to the volume. There must be at least one active path to the volume.

Retry the operation after any existing I/O has been completed.

The Path has not been claimed

Path being unclaimed has not been claimed by PowerPath/VE.

Action

None. The path is already unclaimed.

Performance Monitor is not enabled

Running any perf command when the performance monitor is off.

Action

Ensure that performance monitoring is enabled.

Persistence of SCSI Res fail Value is 0, generating new one

Generating new reservation key.

Action

None.

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.

Action

None.

PowerPath driver is already unlicensed

The `rpowermt unregister` command was issued for a vSphere host that does not have a valid license.

Action

None.

PowerPath Not Found

PowerPath is not installed on the specified host.

Action

If you get the error message after running an `rpowermt` command, ensure that the `/root` partition is not full. If the `/root` partition is full PowerPath/VE might be only partially installed.

Probe of <path> path failed. Status <status>

The path probe has failed.

Action

None. PowerPath/VE will automatically retry the command.

ProbePath Failed path

PowerPath/VE path probe has failed.

Action

Check for path availability.

Path <path>: could not issue sync/async IO command: <status>

PowerPath/VE failed to issue asynchronous I/O.

Action

None. PowerPath/VE will automatically retry the command.

SCSI-2 reservation/release failed with SCSI status <status>

SCSI 2 reservation/release has failed.

Action

None. PowerPath/VE will automatically retry the command.

SCSI-2 Translation command failed after the maximum number of retries

The SCSI 2 translation command has failed after maximum number of retries.

Action

None. PowerPath/VE will automatically retry the command.

Sense Info S: <sc> ASC: <asc> ASCQ: <ascq>

Non-retryable error. Some I/O failed.

Action

None. PowerPath/VE will handle this error.

Status <status>. Failed to register device. User should unclaim paths.

Paths to the device have been claimed but the device could not be registered. The user should unclaim paths to this device.

Action

Check if maximum number of devices has already been claimed or if some paths to this device are claimed by other multipathing plugins like NMP. If so, unclaim paths to this device.

Unable to send cmd (cleanup PRI) to device. SCSI status <status>

PowerPath/VE is unable to determine device state.

Action

None. PowerPath/VE will automatically retry the command.

Unable to send cmd (cleanup PRO) to device. SCSI status <status>

PowerPath/VE is unable to unregister the device.

Action

Check for device connectivity and check the system log for SCSI status.

Could not unregister the device

PowerPath/VE cannot unclaim path because this is the last available path and is being used.

Action

Stop all I/O on the device before trying to unclaim all paths to the device.

User Authentication Invalid

An `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.

Action

None.

VM either does not exist or is not registered on this host

The VM is not registered on this ESXi host

Action

Register the VM or ensure that the VM exists.

vmk_LogRegister failed: <status>

PowerPath/VE failed to register log handle.

Action

PowerPath/VE load failed. Check if enough memory is available.

Volume is being destroyed. Cannot open

Device could not be opened.

Action

None. PowerPath/VE will automatically retry the command.

Volume not found

SCSI device could not be found.

Action

None. PowerPath/VE will automatically retry the command.

Volume to be opened not found

Device could not be opened.

Action

None. PowerPath/VE will automatically retry the command.

The volume is open. Cannot unclaim last path <path>

Cannot unclaim path because it is the last available path and is being used.

Action

Stop all I/O on the device before trying to unclaim all paths to the device.

The volume is being registered or unregistered. Cannot unclaim last path <path>

PowerPath/VE cannot unclaim path because this is the last available path and is being used.

Action

PowerPath/VE load failed. Check if enough memory is available.

Wait for n losPending count (count) to become 0

Until I/O count becomes zero, PowerPath/VE will not add or remove path. There is a delay in adding or removing path.

Action

None. PowerPath/VE will automatically retry the command.

Wait for volume to be unquiesced

Device is currently quiesced; waiting for it to be unquiesced.

Action

None. PowerPath/VE will automatically retry the command.

PowerPath Management Appliance messages

This chapter contains the following topics:

Topics:

- Although served license added, there was an error contacting License Server. Check License Server status
- License already present
- License key not valid for the domain / IP address of this Virtual Appliance
- Unable to add new license. Reason: License is invalid or expired

Although served license added, there was an error contacting License Server. Check License Server status

Electronic License Manager System (ELMS) failed to start or re-read licenses.

Action

Check if the license server is running.

```
ps aux|grep elms
```

The reason for the license server's failure will be available either in `/opt/emc/elms/elms_debug.log` or `/opt/emc/elms/elms_service.log`

License already present

Same license is already present in PowerPath Management Appliance.

Action

No action is required. The license is already available in PowerPath Management Appliance.

License key not valid for the domain / IP address of this Virtual Appliance

License contains invalid IP or Domain

Action

Get a new license with the correct domain name/IP.

Unable to add new license. Reason: License is invalid or expired

License signature is invalid or the overdraft count is 0, or the license being uploaded is expired.

Action

Perform one of the following steps:

- Get license with valid signature and non-zero overdraft count.
- Get a new valid license.

PowerPath family new and changed commands

This appendix contains the following topics:

Topics:

- [powermig new and changed commands](#)
- [powermt new and changed commands](#)
- [rpowermt new and changed commands](#)

powermig new and changed commands

The following table lists the new and changed PowerPath Migration Enabler commands added in versions 6.0 and later, with the earliest supported version listed by platform. Support in later versions (including service packs) can be assumed if no exception in the table notes otherwise. Version 5.0 is the first Migration Enabler release.

Table 12. New and changed commands in 6.0 and later

Command	Description	AIX	Linux	Solaris	Windows
powermig setup	Updated to accept <code>-cluster -cluster mscd</code> (for cluster disk migration) and <code>-cluster mscsv</code> (for CSV migration)	N/A	N/A	N/A	6.0
powermig setup	Updated to exclude nodes in a SQL AlwaysOn environment	N/A	N/A	N/A	6.0 SP1

powermt new and changed commands

The following table lists the new `powermt` commands added or changed in versions 6.0 and later, with the earliest supported version listed by platform. Support in later versions (including service packs) can be assumed if no exception in the table notes otherwise.

Table 13. New and changed commands in 6.0 and later

Command	Description	AIX	Linux	Solaris	Windows
powermt (multiple commands)	Addition of Device WWN for Symmetrix devices	6.0	6.0	6.0	6.0
powermt (multiple commands)	Addition of Unity ID	6.1	6.1	N/A	6.1
powermt (multiple commands)	Updated to notify when PowerPath is unlicensed	N/A	6.2	N/A	N/A
powermt display	Updated to display the IP address of the NIC	N/A	6.2	N/A	N/A
powermt display [dev]	Addition of the verbose option for XtremIO long nice names	6.0 SP1	6.0 SP1	N/A	6.0 SP1
powermt display [dev]	Addition of VPLEX nice names	6.0 SP1	6.0 SP1	N/A	6.0 SP1
powermt display [dev]	Updated to display the SRDF/Metro logical device ID for the distributed volume	6.1	6.1	N/A	6.0 SP1

Table 13. New and changed commands in 6.0 and later (continued)

Command	Description	AIX	Linux	Solaris	Windows
powermt display [dev]	Updated to display the non-spoofed ID for Symmetrix devices which are part of SRDF/Metro	N/A	6.1	N/A	N/A
powermt display options	Information on auto host registration	6.0	6.0	6.0	6.0
powermt display port_mode	Modified to support setting array backend port mode	6.0	6.0	6.0	6.0
powermt set all_paths_dead_retry_timeout	New command	N/A	6.0	N/A	N/A
powermt set auto_host_registration	Modified to support class option	6.0	6.0	6.0	6.3
powermt set autostandby	Addition of setting asb:prox threshold values	N/A	N/A	N/A	6.1
powermt set autostandby agingperiod	Addition of setting aging period in hours	6.0	6.0	6.0	6.0
powermt set autostandby agingperiod	The default aging period changed from 7 days to 1 day	6.0	6.0	6.0	6.0
powermt set device_perf_to_array_report	New command	N/A	N/A	N/A	6.3
powermt set dev_inuse_to_array_report	New command	N/A	N/A	N/A	6.3
powermt set path_retry_timeout	New command	N/A	6.0	N/A	N/A
powermt set port_mode	New command	6.0	6.0	6.0	6.0
powermt update host_registration	New command	N/A	N/A	N/A	6.3

rpowermt new and changed commands

The following table lists the new `rpowermt` commands added or changed in versions 6.0 and later, with the earliest supported version listed by platform. Support in later versions (including service packs) can be assumed if no exception in the table notes otherwise.

Table 14. New and changed commands in 6.0 and later

Command	Description	PowerPath/VE
rpowermt (multiple commands)	Addition of Device WWN for Symmetrix devices	6.0
rpowermt (multiple commands)	Addition of Unity ID	6.1
rpowermt check_registration	Removal of License Count and Overdraft Count	6.0
rpowermt check_served	New command	6.0
rpowermt display [dev]	Addition of the <code>type</code> field	6.0
rpowermt display [dev]	Addition of the verbose option for XtremIO long nice names; addition of VPLEX nice names; updated to display the SRDF/Metro logical device ID for the distributed volume; updated to display the non-spoofed ID for Symmetrix devices which are part of SRDF/Metro	6.1
rpowermt display options	Information on auto host registration	6.0
rpowermt display perf bus vmname	New command	6.0 SP1

Table 14. New and changed commands in 6.0 and later (continued)

Command	Description	PowerPath/VE
rpowermt display perf vmname	New command	6.0 SP1
rpowermt display port mode	New command	6.0
rpowermt set auto_host_registration	Modified to support class option	6.3
rpowermt set autostandby	Addition of setting asb:prox threshold values	6.1
rpowermt set autostandby agingperiod	Addition of setting aging period in hours	6.0
rpowermt set autostandby agingperiod	The default aging period changed from 7 days to 1 day	6.0
rpowermt set device_perf_to_array_report	New command	6.3
rpowermt set dev_inuse_to_array_report	New command	6.3
rpowermt set migration_enabler	New command	6.1
rpowermt set port_mode	New command	6.0
rpowermt update host_registration	New command	6.3
rpowermt set vm_perfmon	New command	6.0 SP1

List of the PowerPath load balancing and failover policies


This appendix contains the following topics:

Topics:

- [List of the PowerPath load balancing and failover policies](#)

List of the PowerPath load balancing and failover policies

Displays the list of the PowerPath load balancing and failover policies.

- ad (adaptive)** I/O requests are assigned to paths based on an algorithm that takes into account path load and logical device priority.
- This policy is available for the following storage classes: ess, hitachi, hpxp, invista, netapp, sc, vplex, xtremio
 - This policy is the default policy with a PowerPath license for Hitachi, HP, Invista, Netapp, VPLEX, Dell SC, and XtremIO storage systems.
- bf (basic failover)** The basic failover policy is supported only in configurations with one HBA connected to a storage system through a switch and one HBA port zoned to each SP on a VNX and CLARiiON system, or to a port on two separate FAs on a Symmetrix system. Using this policy in other configurations is not supported and may result in loss of access to data in the event of path failures. Load balancing is not in effect. I/O routing on failure is limited to one HBA and one port on each storage system interface. When a host boots, it designates one path (through one interface) for all I/O. If an I/O is issued to a logical device that cannot be reached via that path (that is, the I/O cannot reach that logical device through the device's assigned interface), the logical device is assigned to the other interface.
- This policy protects against VNX and CLARiiON SP failures, Symmetrix FA port failures, and back-end failures, and it allows non-disruptive upgrades to work when running PowerPath without a license key. It does not protect against HBA failures. HP-UX hosts designate the first two paths to a volume for use as basic failover paths; the remaining paths are unlicensed.
- This policy is available for the following storage classes: clariion, invista, symm, vnx, vplex
 - This policy is the default policy without a PowerPath license for EMC storage arrays.
- 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 available for the following storage classes: clariion, unity, vnx
 - This policy is the default policy with a PowerPath license for Unity, VNX, and CLARiiON storage systems.
- 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 fewest queued blocks, regardless of the number of requests involved.
- This policy is available for the all storage classes.
 - This policy is not a default policy on any storage system.
- 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 fewest queued requests, regardless of total block volume.
- This policy is available for the all storage classes.
 - This policy is not a default policy on any storage system.
- nr (no redirect)**  **NOTE: Do not use this policy in production environments. Use only for diagnostic purposes.**

Neither load balancing nor path failover is in effect. If nr is set on a pseudo device and the I/O path fails, data errors can occur. If nr is set on a native device and the I/O path fails, data errors will occur. If the policy is nr and there is only one path, I/O to that path is not guaranteed.

- This policy is available for the following storage classes: ess, hitachi, hpxp
- This policy is the default policy without a PowerPath license for third-party storage arrays.

re (request)

For native devices, this policy uses the path that would have been used if PowerPath were not installed. For pseudo devices, it uses one arbitrary path for all I/O. For all devices, path failover is in effect, but load balancing is not.

- This policy is available for the following storage classes: All (PowerPath license), CLARiiON only (PowerPath Base license)
- This policy is not a default policy on any storage system.

rr (round robin)

I/O requests are assigned to each available path in rotation.

- This policy is available for the all storage classes.
- This policy is not a default policy on any storage system.

si (streamIO)

For each I/O to a volume, streamlo policy uses the same path that was used for the previous I/O to the volume. However, if the volume's I/O count exceeds the volume's stream I/O threshold since the last path selection, the policy selects a path using the adaptive policy algorithm. The volume I/O count is rezeroed on each path change.

- This policy is available for the all storage classes.
- This policy is not a default policy on any storage system.

so (Symmetrix optimization)

I/O requests are routed to paths based on an algorithm that takes into account path load. Load is a function of the number, size, and type of I/O queued on each path.

- This policy is available for the following storage classes: symm
- This policy is the default policy with a PowerPath license for Symmetrix storage systems.

List of fields used in the powermt and rpowermt display command output

This appendix contains the following topics:

Topics:

- [List of all fields used in the display output](#)

List of all fields used in the display output

Displays the full list of all available fields found in the various output for the `powermt` and `rpowermt display` commands.

The fields vary based on the command that is run.

<code>storage_system_type ID</code>	<p>Identification number for the storage system on which the logical device is located. Each <code>powermt/rpowermt display</code> table contains information about one type of storage system, such as Symmetrix.</p> <p>For VNX and CLARiiON systems only, the user-configurable storage group name, if available, is by default displayed in brackets after this identification number.</p>
<code>storage_system_type logical device count</code>	<p>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 <code>powermt/rpowermt display</code> table contains information about one type of storage system, such as Symmetrix.</p>
<code>Array failover mode</code>	<p>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.</p> <p>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.</p> <p>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.</p> <p>Modes supported with PowerPath include:</p> <ul style="list-style-type: none"> • 1—passive not ready (PNR) mode—Supported with all PowerPath platforms. • 3—passive always ready (PAR) mode—Supported only with PowerPath for AIX in specific configurations. • 4—asymmetric logical unit access (ALUA) mode—Supported with PowerPath on some platforms, as described in the E-Lab Interoperability Navigator. <p>Refer to the VNX and CLARiiON documentation for more information about the array failover mode.</p>
<code>Device class name count</code>	<p>Storage device class name and the total number of devices accessed through the buses.</p>
<code>Device WWN</code>	<p>This applies to EMC Symmetrix arrays only. The Device WWN field displays the logical unit NAA id from VPD page 0x83.</p> <p>In cases where the WWN is not available, the Device WWN will be displayed as N/A. For example on VMAX arrays the device WWN cannot be retrieved if the SPC-2 and SC3 flags are not set.</p>
<code>Error delta</code>	<p>Number of alive to dead transitions during the interval.</p>
<code>Host Bus Adapters ###</code>	<p>PowerPath number for the HBA. This number is preserved across boots but is not preserved after configuration changes.</p> <p>When the <code>dev</code> option is used, the output of <code>powermt/rpowermt display</code> identifies the HBA by this same HBA number.</p>
<code>Host ###</code>	<p>PowerPath number for the HBA. This number is preserved across boots but is not preserved after configuration changes. When the <code>dev</code> option is used, the output of <code>powermt display</code> identifies the HBA by this same HBA number.</p>

Host Bus Adapters Displays the path to the host bus adapter.

HW Path

- AIX: The name of the device that is the parent of the hdisk; for example, scsi1.
- HP-UX: Hardware path as defined by the `ioscan` command; for example, 10/4.2.0.
- Linux: HBA name registered by the HBA driver, for example, lpfc. Registration of the HBA name is optional. For those drivers that do not register a name, Linux displays the string information returned from the driver. This string is truncated to 24 characters without parsing. To view the entire name of the HBA, under the `/proc/scsi` directory, change to the subdirectory containing copies of the HBA drivers for the HBA SCSI bus instance you are investigating. Use the `cat` command to view the full name of the HBA.
- Solaris: Modification of the `/dev/rdisk` symbolic link target; for example, sbus@1f,0/fcaw@3,0.
- UNIX: Operating-system-specific hardware descriptive name for the bus.
- VMware: vSphere-assigned name for the HBA.
- Windows: port#\path#, where # is the integer displayed in Properties in Disk Administrator or Disk Management; for example, port1\path0.

Host HW Path Platform-specific hardware descriptive name for the path.

- AIX: The name of the device that is the parent of the hdisk; for example, fscsi1.
- HP-UX: Hardware path as defined by the `ioscan` command; for example, 10/4.2.0.
- Linux: HBA name registered by the HBA driver. Registration of the HBA name is optional. For those drivers that do not register a name, Linux displays the string information returned from the driver, truncated to 24 characters without parsing. To view the entire HBA name, under the `/proc/scsi` directory, change to the subdirectory containing copies of the HBA drivers for the HBA SCSI bus instance you are investigating. Use the `cat` command to view the HBA's full name.
- Solaris: Modification of the `/dev/rdisk` symbolic link larger; for example, /pci@8,700000/QLGC,qla@3/sd@0,2.
- VMware: vSphere-assigned name for the HBA.
- Windows: port#\path#\tgt#\lun#, where # is the integer displayed in Properties in Disk Administrator or Disk Management. For example, port2\path0\tgt6\lun7. (A LUN, or Logical Unit Number, is a logical device.)

On Linux SLES 11 and later, UNKNOWN may appear in this column following a path failure. This occurs because the corresponding native device has been removed from the operating system kernel.

Host I/O Path The platform-specific device name for the path.

- On HP-UX and Solaris, if new devices are added, `powermt display every=#seconds` may display unknown in place of actual device names, due to the creation of device-name mappings at startup. If this occurs, restart `powermt display every=#seconds`
- On Linux SLES 11 and later, UNKNOWN may appear in this column following a path failure. This occurs because the corresponding native device has been removed from the operating system kernel.

HW Path Platform-specific hardware descriptive name of every path for the bus.

ID Platform-specific ID for the bus.

I/O Paths Dead Total number of paths originating from this HBA that are dead (not usable).

I/O Paths Mode (mode) Current path mode:


- active specifies a path that will be actively used for I/O for multipathing.
- standby specifies a path that has been manually placed in standby mode by the user. This path will not be used unless all other active paths are dead. However, note that with ALUA-enabled and active-passive devices, standby paths will be used in favor of active/non-optimized and passive paths, respectively.
- 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.
- unlic indicates that unlicensed PowerPath is running for a Symmetrix, VNX, and CLARiiON storage system (no license key has been installed). In this scenario, all paths are marked unlicensed except one path to each SP. Unlicensed paths cannot become candidates for path failover.

I/O Paths State 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 will not direct I/O to this path. After marking the path dead and incrementing the Errors count, PowerPath tests the path to see whether it is usable. If the test succeeds, PowerPath marks the path alive; the path is then available for I/O. If the test fails, the path remains dead, and

PowerPath ignores it for subsequent I/O operations. If all the paths to a logical device are dead, PowerPath retests each path.

I/O Paths Summary	Status of the paths originating from this HBA. <ul style="list-style-type: none">· 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.
I/O Paths Total	Total number of paths that originate from this HBA. The maximum number of logical devices supported by one HBA is platform specific. The total number of paths may exceed the number of logical devices in a complex SAN topology, due to zoning.
Invista ID	Identification number for the Invista or VPLEX system on which the logical device is located.
Latency	When the verbose option is specified, the Read and Write I/Os per second count for each device path is shown under each bins. The latency bins shown in the verbose output are a good place to verify that dead, standby (both manual and autostandby), and non-optimal paths are not handling application reads and writes under normal circumstances. Keep in mind, however, that even standby and non-optimal paths can have application reads and writes routed to them under extremely heavy loads, and such "overflow" I/Os would show up in the latency bins.
Latency Current	The amount of time, in microseconds, it took the last I/O on the path to complete.
Latency Max	The longest time, in microseconds, it took any I/O on the path to complete.
Latency Reads	The high and low water marks for Read counts per each device path or bus in the interval. In milliseconds.
Latency Writes	The high and low water marks for Write counts per each device path or bus in the interval. In milliseconds.
Logical device ID	Identification number for the logical device. Each logical device on each storage system has a unique ID. Each storage system, however, uses the same storage-system ID. Together, storage-system ID and logical-device ID create a unique ID for every logical device in the world. On a VNX and CLARiiON system, the logical-device ID is a 32-digit number. This is 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 brackets after this identification number.
Mode	The mode of the HBA or port. <ul style="list-style-type: none">· Enabled or Active means that the HBA or port is alive (usable) and operating normally.· Disabled or Standby means the HBA or port was disabled previously with the <code>powermt disable</code> or <code>powermt set port_disable</code> command. While an HBA is disabled, I/O is not sent to it and devices configured through it are closed.
NonVirtual Device ID	For VNX and CLARiiON systems, the identification number associated with the physical device and the VNX and CLARiiON storage system. For Symmetrix systems, the identification number associated with the physical device.
NonVirtual pseudo name	For the physical device associated with the Invista or VPLEX Virtual Volume, the platform-specific value assigned by PowerPath to the device. When using Migration Enabler (Virtual encapsulation) to set up a migration, this is the source device name corresponding to the target device name given in the Pseudo name field (in this command's output).
NonVirtual Symmetrix ID	For Symmetrix systems, the identification number for the Symmetrix system on which the physical device is located.
Owner	Default and current owners of the logical device (VNX and CLARiiON systems only).
Path status	Combination of the I/O Paths Mode and I/O Paths State columns of the normal devices display. Path Status is a two-part field, with the following valid values: a (active), s (standby), u (unlicensed), v (alive), d (dead). Standby in this situation includes asb:prox and asb:iopf standby created automatically by PowerPath algorithms. The u (unlicensed) state is not valid for PowerPath/VE.
policy	Current load-balancing and failover policy for the device.
Pseudo name	Platform-specific value assigned by PowerPath to the PowerPath device. For perf dev display, the storage device pseudo name for each storage device from which the metrics are gathered. The metrics for each storage devices are listed separately.
queued-I/Os	Number of I/O requests queued to this PowerPath device.

Retry delta	Number of retry attempts to restore a failed I/O through another path or bus to the storage device during the interval.
Sample Interval	Interval period when the metrics are gathered. In seconds.
state	State of the PowerPath device: <ul style="list-style-type: none"> • PowerPath marks a PowerPath device dead if it returns certain types of I/O errors and is judged unusable. Once a PowerPath device is marked dead (and until it is restored), PowerPath returns subsequent I/O requests with a failure status, without forwarding them to the associated logical device. This prevents further, unrecoverable corruption and allows the user to perform data recovery if needed. • If a PowerPath device is alive, either PowerPath never marked it dead, or it was marked dead but restored with <code>powermt restore</code> or an automatic path test. • For <code>perf dev</code> display, the state of the storage device from which the metrics are gathered. The state can be either dead or alive.
Stats Errors	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 <code>powermt restore</code> executes, when <code>powermt set port_disable=on</code> executes, or when <code>powermt unmanage</code> executes. PowerPath periodic autorestore, which is performed automatically whenever Plug and Play takes a device offline or brings a device back online, does not clear the error count.
Stats IO/sec	This field is blank for <code>powermt/rpowermt</code> display, unless it is used with the <code>every</code> parameter. Subsequent <code>powermt/rpowermt</code> display iterations display the average number of I/Os sent across this bus each second.
Stats Q-I/Os	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-I/Os 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 <code>ioctls</code> - these are generally unconditionally routed to the "request" path, and not redirected to optimal or active paths.
Storage System ID	Identification number for the storage system on which the logical device is located. This is the array serial number/ID typically set by EMC manufacturing. Note that the formatting or naming scheme of the array serial number/ID will look different in production environments.
Storage System Interface (Storage Interface)	Storage-system interface. <ul style="list-style-type: none"> • For Symmetrix, the Storage Interface field displays as <code><interface_type> <director_number>:<director_port></code> where <code>interface_type</code> is FA (Fibre Channel), SA (SCSI) or SE (iSCSI). For example, the first port in director 6f with FA interface displays as FA 6f:00 • For Unity, VNX, and CLARiiON, this has one part: Interface port is SP [A-B][0-3] • For VPLEX, the storage interface is displayed as <code><cluster id>-<target on the vplex array></code>. For example, CL1-09 • For XtremIO, the storage interface is displayed as <code><brick number>-<controller number>-<port></code>. For example, X1-C1-f1 • For Dell SC, the storage interface is displayed as <code><active controller>-<last 2 digits of target physical port WWPN on active controller></code>. For example, 12345-1f • For HP StorageWorks EVA this has one part: Interface port: [A-B][1-2] where A,B denotes the controller and 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.)
Storage System Wt_Q	Maximum number of write I/O requests that will be serviced before the I/O queue checks for any outstanding read I/O.
Timestamp	Period since when the last metrics were computed. Hours:Minutes:Seconds
type	The type of LUN: Conventional, which are the traditional LUNs used to store data or PE (protocol endpoint), which are virtual LUNs which indirectly access the actual data stored on VVOLs. PE is supported only on PowerPath/VE.
write_throttle	Write-throttling setting for this logical device.

Microsoft Windows IDs

This appendix contains the following topics:

Topics:

- [Microsoft Windows application event IDs](#)
- [Microsoft Windows system event IDs](#)

Microsoft Windows application event IDs

The following table lists the Windows application event messages that appear on Windows platform, along with their IDs.

Table 15. Application event IDs

Event ID	Message
0x0	Insufficient user privilege
0x1	PowerPath Success: %1
0x2	PowerPath Information: %1
0x3	PowerPath Warning: %1
0x4	PowerPath Error: %1
0x5	PowerPath Panic:%1
0x6	PowerPath Debug:%1
0x7	PowerPath Audit:%1: handle=%2, event=%3%4%5%6%7%8%9%10%11%12%13%14
0x8	PowerPath Info:%1: Console CTRL_SHUTDOWN Event Started at %2
0x9	PowerPath Info:%1: Console CTRL_SHUTDOWN Event Ended at %2
0xa	PowerPath Info: %1: Service SHUTDOWN Event Started at %2
0xb	PowerPath Info: %1: Service SHUTDOWN Event Ended at %2
0xc	PowerPath Warn: %1: Unable to flush volume for disk %2 at shutdown
0xd	PowerPath Warn: %1: Unable to get crash status at startup
0xe	PowerPath Error: %1: target disk contains logical volumes that are in use, error=%2
0xf	PowerPath Error: %1: scrub sector(0x%2) for ScrubDevice %3 could not be written at cleanup, error=%4
0x10	PowerPath Error: %1: scrub sector(0x%2) for ScrubDevice %3 could not be read at cleanup, error=%4
0x11	PowerPath Error: %1: scrub sector(0x%2) for NonScrub Device %3 could not be read at cleanup, error=%4
0x12	PowerPath Warn: %1: scrub failure for Device % was overridden by force flag
0x13	PowerPath Warn: %1: source disk contains a partition type(%2) that is not supported
0x14	PowerPath Info: %1: source disk contains a partition type(%2) that is supported

Table 15. Application event IDs (continued)

Event ID	Message
0x15	PowerPath Error: %1: Failed to get Source device name
0x16	PowerPath Error: %1: Failed to get Target device name
0x17	PowerPath Info: %1: Failed to get source name, device must be removed
0x18	PowerPath Error: %1: Verification failed: %2 is still accessible
0x19	PowerPath Info: %1: PowerPath Migration Service successfully started with ShutdownParam = 0x%2
0x1a	PowerPath Error: %1: The target device must be the same size as, or larger than, the source
0x1b	PowerPath Error: %1: For this type of device, the target must be the same size as the source
0x1d	PowerPath Error: %1: db_recover failed at startup with erro code %2
0x1e	PowerPath Error: %1: PPME NamedPipe Server could not impersonate client
0x1f	PowerPath Error: %1: PPME NamedPipe client needs local computer admin privileges
0x20	PowerPath Error: %1: PPME NamedPipe Server could not revert to self
0x21	PowerPath Error: %1: PPME Unable to Enable all Devices at Stratup, error code: %2, result: %3.
0x22	PowerPath Error: %1: PPME unexpected exception cought in main HostCopy thread.
0x23	PowerPath Error: %1: PPME unexpected exception cought in %2.
0x24	PowerPath Info: %1: Scrub Device: %2, signature: %3, Partition[%4] = %5.
0x25	PowerPath Error: %1: Maximum transfer length check failed, SrcMaxLen=%2, TgtMaxLen=%3.
0x26	PowerPath Error: %1: PPME Unable to Configure Devices at Stratup, error code: %2, result: %3.
0x201	PowerPath Error: %1: umd_OpenDbase failed rc = %2 errCode = %3
0x202	PowerPath Error: %1: umd_OpenDbase failed rc = %2 errCode = %3
0x203	PowerPath Error: %1: umd_CloseDbase failed rc = %2 errCode = %3
0x204	PowerPath Error: %1: umd_swapinstance failed rc = %2 errCode = %3
0x205	PowerPath Error: %1: DB Error: %2:%3
0x206	PowerPath Error: %1: write zero failed, block_no=%2, %3
0x207	PowerPath Error: %1: USCSICMD(SCMD_INQUIRY) ioctl failed: %2
0x208	PowerPath Error: %1: cleanup failed. handle=%2 rc=%3
0x209	PowerPath Error: %1: patf_preShowScrub:: %2 Failed, error=%3
0x20a	PowerPath Error: %1:scrub:: Invalid %2 device:%3, error=%4
0x20b	PowerPath Error: %1:scrub:: Failed to applay disk lable
0x20c	PowerPath Error: %1:Source (or an alias) is already in a migration
0x20d	PowerPath Error: %1:Target (or an alias) is already in a migration
0x20e	PowerPath Error: %1:Setup failed. rc=%2
0x20f	PowerPath Error: %1:Sync failed. handle=%2 rc=%3

Table 15. Application event IDs (continued)

Event ID	Message
0x210	PowerPath Error: %1:selectSource failed. handle=%2 rc=%3
0x211	PowerPath Error: %1:selectTarget failed. handle=%2 rc=%3
0x212	PowerPath Error: %1:commit failed. handle=%2 rc=%3
0x213	PowerPath Error: %1:undoRedirect failed. handle=%2 rc=%3
0x214	PowerPath Error: %1: Error
0x215	PowerPath Error: %1:abort failed. handle=%2 rc=%3
0x216	PowerPath Error: %1:throttle failed. handle=%2 rc=%3
0x217	PowerPath Error: %1:getHandle failed. rc=%2
0x218	PowerPath Error: %1:getVersion failed. rc=%2
0x219	PowerPath Error: %1:info failed. rc=%2
0x21a	PowerPath Error: %1:query failed. handle=%2 rc=%3
0x21b	PowerPath Info: %1:Beginning AUTOMATIC RECOVERY
0x21c	PowerPath Info: %1:Successfully recovered handle=%2
0x21d	PowerPath Info: %1:Finished AUTOMATIC RECOVERY
0x21e	PowerPath Info: %1:Beginning AUTOMATIC DEVICE FAULT TRANSITIONS after system crash
0x21f	PowerPath Info: %1:Beginning AUTOMATIC DEVICE FAULT TRANSITIONS at system shutdown
0x220	PowerPath Info: %1:Successfully transitioned handle=%2
0x221	PowerPath Error: %1: Error
0x222	PowerPath Info: %1:Finished AUTOMATIC DEVICE FAULT TRANSITIONS
0x223	PowerPath Error: %1:Internal Error: Invalid TechType (%2)
0x224	PowerPath Info: %1:Failed to get source userdev from devId
0x225	PowerPath Info: %1:Failed to get target userdev from devId
0x226	PowerPath Info: %1:Failed to get source name from userdev(0x%2)
0x227	PowerPath Info: %1:Failed to get target name from userdev(0x%2)
0x228	PowerPath Info: %1:Get IO Cloning Control from kmd: reads=%2, writes=%3
0x229	PowerPath Info: %1:Set IO Cloning Control in %2: reads=%3, writes=%4
0x22a	PowerPath Info: %1:Turned %2 WriteCloning
0x22b	PowerPath Info: %1:Successful call to turn %2 redirection
0x22c	PowerPath Info: %1:Set access control on %2
0x22d	PowerPath Info: %1:Successful call to turn %2 funneling
0x22e	PowerPath Info: %1:Set IO Cloning Control in %2: reads=%3, writes=%4
0x22f	PowerPath Info: %1:Failed to get PPME license; error=0x%2
0x230	PowerPath Error: %1:Failed to get license for %2 shim
0x231	PowerPath Error: %1:SymInit failed with code %2
0x232	PowerPath Error: %1:SymExit failed with code %2
0x233	PowerPath Error: %1:SymRemoteCopyControl failed src %2 tgt %3 err %4 (%5)

Table 15. Application event IDs (continued)

Event ID	Message
0x234	PowerPath Error: %1:SymPdevSync failed with code %2 (%3)
0x235	PowerPath Error: %1:SymPdevShow failed with code %2 (%3)
0x236	PowerPath Error: %1:Unable to find %2 in %3 (rc: %4, status=%5)
0x237	PowerPath Error: %1:Unable to allocate session list: %2
0x238	PowerPath Error: %1:Unable to allocate session info (n=0): %2
0x239	PowerPath Error: %1:Unable to allocate session info (n=%2): %3
0x23a	PowerPath Error: %1:in function %2: %3=%4, errCode=0x%5
0x23b	PowerPath Error: %1:NdmePP_GpxIocmd, mgr=%2, cmd=%3, rc=%4, err=0x%5
0x23c	PowerPath Error: %1:NdmePP_Iocmd retval %2 err 0x%3
0x23d	PowerPath Error: %1:Failed to get device name from userdev(0x%2)
0x23e	PowerPath Error: %1:Failed to get source name from userdev(0x%2)
0x23f	PowerPath Error: %1:Failed to get target name from userdev(0x%2)
0x240	PowerPath Error: %1:LAM Initialization Error, rc=%2, error code=0x%3
0x241	PowerPath Error: %1:Failed to get inquiry data for %2.rc=%3, error code=0x%4",
0x242	PowerPath Error: %1:Unable to determine array type for %2.rc=%3, error code=0x%4
0x243	PowerPath Error: %1:Failed to get WWN for %2. rc=%3, error code=0x%4
0x244	PowerPath Error: %1:Inquiry data is not recognized for %2. rc=%3, error code=0x%4
0x245	PowerPath Error: %1:Unable to get source device info (rc=%2, *pErr=%3)
0x246	PowerPath Error: %1:Unable to get target device info (rc=%2, *pErr=%3)
0x247	PowerPath Error: %1:Unable to open INV Target Device
0x248	PowerPath Error: %1:Unable to query Invista state
0x249	PowerPath Error: %1:Target should be in "Not Ready" state instead of %2
0x24a	PowerPath Error: %1:Target encapsulates %2, not source (%3)
0x24b	PowerPath Error: %1:Target encapsulates %1, not source (%2)
0x24c	PowerPath Error: %1:Unable to open INV target device
0x24d	PowerPath Error: %1:Unable to query Invista state
0x24e	PowerPath Error: %1:Inv test target failed
0x24f	PowerPath Error: %1:Unable to set Invista state to TEST
0x250	PowerPath Error: %1:Target should be in TEST or COMMIT state instead of %2
0x251	PowerPath Error: %1:Invista commit failed to open target
0x252	PowerPath Error: %1:Unable to set Invista state to COMMIT
0x253	PowerPath Error: %1:Target should be in COMMIT state instead of %2
0x254	PowerPath Error: %1:Invista unable to open Target in abort
0x255	PowerPath Error: %1:Invista unable to open Target in abort
0x256	PowerPath Error: %1:Unable to query Invista state

Table 15. Application event IDs (continued)

Event ID	Message
0x257	PowerPath Error: %1:Unable to set Invista state to NOT_READY
0x258	PowerPath Error: %1:Unable to get source device info (rc=%2, *pErr=%3)
0x259	PowerPath Error: %1:Unable to get source device ID from name
0x25a	PowerPath Error: %1:Unable to get target device ID from name
0x25b	PowerPath Error: %1:Unable to determine redirection state
0x25c	PowerPath Error: %1:Unable to get Invista name from device ID
0x25d	PowerPath Info: %1: Beginning AUTOMATIC DEVICE FAULT TRANSITIONS at system startup
2	emcpAudit: Info: cmd=%s: %s Time Elapsed: %d seconds (user ID real=%s effective=%s)

Microsoft Windows system event IDs

The following table lists all messages that appear on Windows platforms, along with their IDs.

If a PowerPath message exceeds the maximum message size allowed by Windows, the Event Viewer truncates the message and places the data fields in the dump portion of the event (the hex information at the bottom of the viewer). For example, the message Path c3t1d10 to 60001440B06028A300A4A48002000010 is alive. exceeds the allowed maximum size. The message would be displayed as Path %1 to %2 is alive. with the data for %1 and %2 included at the bottom.

Table 16. System event IDs

Event ID	Message
0x2	PowerPath %2 driver loaded.
0x5	PowerPath psuedo bus cannot be configured. Uninstall then reinstall PowerPath.
0x6	PowerPath cannot be loaded in safeboot.
0x9	The flag FilterClusterIoctl is turned off, cluster service is not fully supported by PowerPath
0x10	PowerPath debug to memory is on, buffer size is %2 bytes
11	PowerPath debug to memory has been turned off
12	PowerPath debug to memory failed start because of memory allocation error
13	PowerPath parameter %2 was changed from %3 to %4
14	PowerPath Error removing records for the KMD %2 property
100	Path %2 to %3 is dead. Disk %2 has reached a logical block provisioning soft threshold.
101	Path %2 to %3 is alive. The logical block provisioning type for Disk %2 has changed.
102	%2 is dead.
103	%2 is alive.
104	paths to %2 are dead.
105	Not all paths to %2 are dead.
106	Bus %2 to %3 port %4 is dead.
107	Path %2 to %3 is %4 deferred.

Table 16. System event IDs (continued)

Event ID	Message
108	Volume %2 is unbound.
109	Restored volume %2 to default: %3.
110	Unable to restore volume %2 to default controller.
111	Volume %2 %3 owner followed to %4.
112	Unable to follow volume %2 to %3.
113	Trespassed volume %2 to %3.
114	scsi-2 reserve not allowed with policy other than NoRedirect. cmd = 0x%2.
115	Volume [%2] logging changed to: Error=%3 Warn=%4 Info=%5.
116	%2 is in an unsupported mixed ALUA configuration
117	Latency Threshold exceeded by Path %2 Vol %3 with a new high latency of %4 msec
120	Bus %2 to %3 port %4 is alive.
121	Paths to non-preferred %2 ID %3 set to asb:prox.
122	Preferred %2 ID is %3.
123	Min latency from %2 ID %3 is %4 micro seconds.
124	Path %2 is set to Auto standby due to Iopf.
125	Path %2 is moved to Active from Auto standby due to path aging.
126	Path %2 to %3 ID %4 is set to %5%6%7%8%9.
127	Path %2 to %3 ID %4 is set to %5%6%7%8%9.
128	Path %2 to %3 ID %4 is set to %5%6%7%8%9.
129	Path %2 to %3 ID %4 is set to %5%6%7%8%9.
200	New SCSI controller %2 with non compatible parameters added. To guarantee correct PowerPath functioning reboot is required

AIX error logging facility messages

This appendix contains the following topics:

Topics:

- [AIX error logging facility messages](#)

AIX error logging facility messages

The following table lists the messages that are generated for use with the AIX Error Logging Facility. These messages can be found in the `errpt -a report`.

Table 17. Logging facility messages

LABEL	IDENTIFIER	Description
EMCP_PATH_ALIVE	C6E26F3B	This event indicates a single Alive Path. A path is defined as the route through which the I/Os reach the LUN. A path is a unique connection between a host bus adapter and a LUN through a particular storage array port. Alive state indicates that a path is usable, I/Os can be directed on an alive path.
EMCP_PATH_DEAD	516A2BC4	An event indicating a single Dead Path. A dead path is not usable. I/Os cannot be directed on a dead path, as path test failed on it.
EMCP_VOL_ALIVE	DA1CBA8D	This event indicates a single Alive Volume. This event indicates a single Alive Volume. Volume in PowerPath terminology refers to a LUN.
EMCP_VOL_DEAD	2C121728	An event indicating a single Dead Volume. A volume is marked dead when a I/O failed on a path and thereafter all retry attempts on all possible paths too failed.
EMCP_ALL_PATHS_DEAD	1E67811B	An event indicating All Paths for Device Dead.
EMCP_NA_PATHS_DEAD	1BBD20F4	An event indicating Not All Paths for Device Dead.
EMCP_DUP_CONFIG	459BC757	This event indicates that the device is already Configured. A configured device will be in Available state, and is capable of handling I/Os.
EMCP_PIN	C72A0C06	This event indicates that pinning of the pages associated with PowerPath driver has failed. Pinned pages cannot be paged out of real memory.
EMCP_DEVSW	F18230A1	An event indicating that adding device entry of PowerPath driver to the kernel device switch table failed. The device switch table in kernel is a table of device switch structures indexed by the device driver's major device number.
EMCP_BUS_DEAD	3767AAFF	An event indicating a dead bus. A bus is a set of physical connections which can be shared by multiple hardware components to communicate with each other. Each path from the HBA port to the target array port is referred to as a bus.
EMCP_BUS_ALIVE	3B34B002	An event indicating an Alive bus. A bus is a set of physical connections which can be shared by multiple hardware components to communicate with each other. Each path from the HBA port to the target array port is referred to as a bus.
EMCP_VOL_TRESPASS	EA8F3B3F	An event indicating that the volume has been trespassed. Trespass is a mechanism to reassign a LUN from one SP to its peer SP. This event

Table 17. Logging facility messages (continued)

LABEL	IDENTIFIER	Description
		causes all I/Os to be quiesced on one SP, and the cache to be cleaned for this LU in preparation for switching I/O over to its peer.
EMCP_VOL_RESTORED	FE746689	An event indicating that the volume has been restored. A volume restore is a mechanism to assign a LUN back to its originally owning SP.
EMCP_MINOR_NUMBER	764D20E6	This event indicates that configuration method for the hdiskpower in question, failed to find available minor number for the device.
EMCP_LATENCY_MAX_EXCEEDED	95D0858D	An event indicating that maximum Path Latency Exceeded. Path latency is the time taken by an I/O to complete on a particular path. Maximum Path Latency acts as a high water mark for the I/O completion time.
EMCP_PMIG_BEGIN	C549554E	An event indicating that Migration Enabler has Begun. Migration Enabler is a migration tool that enables no disruptive or minimally disruptive data migration between storage systems or between logical units within a single storage system.
EMCP_PMIG_END	E9595B51	An event indicating that Migration Enabler has Ended.
EMCP_PMIG_ABORTED	F2A95EA4	An event indicating that Migration Enabler has been Aborted.

Index

A

array serial
 rpowermt display array serial [110](#)
auto_host_registration
 rpowermt [132](#)

C

CLI install directories [17](#)

P

powermet set
 all_paths_dead_retry_timeout [89](#)
powermt config [59](#)
powermt display bus-ports san-latency
 bus-ports san-latency [68](#)
powermt display perf dev [77](#)
powermt unmanage [102](#)

R

rpowermt check_registration
 check_registration [106](#)
rpowermt display
 rpowermt display for NVMe device [108](#)
rpowermt display alua
 alua [109](#)
rpowermt display bus [111](#)
rpowermt display options [113](#)
rpowermt display paths [116](#)
rpowermt display ports [125](#)
rpowermt unregister [146](#)
rpowermt update host_registration
 host_registration [147](#)

U

Unique Error Codes [154](#)