



EMC[®] PowerPath[™] for AIX

Version 6.0

Installation and Administration Guide

302-001-305

REV 01

EMC²

Copyright © 1997 - 2015 EMC Corporation. All rights reserved. Published in the USA.

Published December 2014

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

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

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

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

CONTENTS

PREFACE

Related documentation	7
Revision History	8
Conventions used in this guide	8
Typographical conventions.....	9
Where to get help	9
License key ordering and activation	10
Your comments	10

Chapter 1

Installing PowerPath on an AIX host

Before you install	12
Reviewing the Storage System Environment	12
Reviewing the Host Environment	12
Obtain up-to-date information.....	12
Locate your license key	13
Prepare the host and storage system.....	13
Installing PowerPath as part of a mksysb master image	16
Removing the hdiskpower devices	16
Removing all the hdiskpower devices and the powerpath0 device.....	17
Installation methods	17
Mounting the installation CD	18
Installing by using the command line	18
Installing by using SMIT	19
Installing by using the tar file	19
Installing a PowerPath service pack or patch	20
After you install.....	21
Register PowerPath on the host.....	21
Choose a major number	22
Initialize PowerPath hdiskpower devices	22
Remove the CD.....	23
Configure the management daemon.....	23
Vary on volume groups.....	23
Reconfigure applications.....	23
Reserve_policy setting in Dual VIO Server Environments.....	24
Upgrading PowerPath.....	25
Before upgrading PowerPath	25
During an upgrade	25
After an upgrade	25
Upgrade procedure	25
Upgrading PowerPath in a VIO server environment	26
Upgrading PowerPath in a PowerHA/HACMP concurrent VG environment....	29
Troubleshooting the upgrade	29
If a PowerPath device is the boot device.....	30
Upgrading the AIX operating system	30
Upgrading AIX operating systems for major releases	30
Upgrading AIX operating systems for Technology Level upgrades or Service Pack upgrades	30
Upgrading AIX operating system using nimadm command	31

	Coexistence with third-party path management software.....	32
Chapter 2	Installing PowerPath on Veritas Volume Manager	
	Installing with Veritas Volume Manager on new hosts.....	36
	Converting to PowerPath from DMP	36
	Adding a new path	37
Chapter 3	Configuring a PowerPath Boot Device on AIX	
	Setting up a PowerPath boot device	38
	Considerations for Symmetrix storage	38
	Considerations for VNX OE and CLARiiON storage.....	38
	PowerPath and the bosboot tool	39
	VNX OE and CLARiiON trespasses.....	39
	Configuring a new PowerPath installation.....	39
	Configuring an existing PowerPath installation	41
	Updating the system microcode	41
	Installing PowerPath	41
	Cloning the boot device.....	42
	Troubleshooting.....	43
	Create a rootvg mirror on AIX.....	44
	Install PowerPath and configure a boot device	44
	Create a Rootvg mirror.....	44
	PVID mismatch and Usage of 'pprootdev fixback'	44
	Adding paths	45
	Disabling PowerPath on a storage system boot device.....	45
	Booting an AIX host from a SAN device.....	45
Chapter 4	Configuring a PowerPath Logical Device	
	Introduction.....	48
	Adding new paths to a logical device.....	48
	Adding new logical devices	50
	Removing paths or logical devices.....	51
	Cleaning up incorrect path or logical volume removal.....	52
Chapter 5	PowerPath in a Cluster	
	PowerPath in a PowerHA/HACMP Cluster.....	54
	Installing PowerPath and PowerHA/HACMP on new hosts.....	54
	Integrating PowerHA/HACMP in a PowerPath environment.....	55
	cfgscsi_id binary in a CLARiiON environment.....	56
	Synchronizing volume group information between cluster nodes	56
	hdiskpower devices in a PowerHA/HACMP environment.....	57
	PowerPath in a Veritas Cluster Server cluster	57
	Installing PowerPath in a new VCS cluster	57
	Integrating PowerPath into a VCS cluster.....	58
	Changing the naming scheme for Veritas volumes.....	59
Chapter 6	Removing PowerPath from an AIX Host	
	Before removing PowerPath	62
	Removing PowerPath	62
	Using command line entries.....	62
	Using SMIT.....	63

After removing PowerPath 64
 When a storage system device is the boot device 64

Chapter 7 PowerPath Administration on AIX

Management daemon using SNMP..... 66
 Monitored events 66
 Configuring the management daemon..... 67
 Editing the event 68
 emc_cfgmgr script..... 68
 Changing the major number..... 69
 PowerPath hdiskpower devices 69
 PowerPath and the lsvg command..... 71
 PowerPath and the iostat command 72
 BCVs and cfgmgr operations 72
 Bringing hdiskpower-based BCV Symmetrix logical devices online 72
 Importing an LVM volume group from a remote host..... 73
 Changing the target/LUN address of a storage system logical device 73
 Adding new devices to an existing configuration 74
 Managing and unmanaging EMC Celerra iSCSI devices 74
 Manage EMC Celerra iSCSI devices under PowerPath 74
 Unmanage EMC Celerra iSCSI devices under PowerPath 74
 Manage IBM XIV devices under PowerPath 74
 Replacing an HBA that PowerPath is using online 75
 Moving LUNs in and out of a storage group..... 75
 Eliminating ghost devices 75
 Reconfiguring PowerPath devices online 76
 Failover in switched environments 76
 SMIT screens..... 77
 Audit and error messages 77
 Error log messages..... 77
 Troubleshooting..... 78

Appendix A Files Changed by PowerPath Installation

Index

PREFACE

As part of an effort to improve and enhance the performance and capabilities of its product line, EMC from periodically releases revisions of its hardware and software. Therefore, some functions described in this manual may not be supported by all revisions of the software or hardware currently in use. For the most up-to-date information on product features, refer to your product release notes.

If a product does not function properly or does not function as described in this manual, please contact your EMC representative.

About this guide

This guide is part of the PowerPath documentation set, and is intended for use by a system or storage administrator during installation, configuration, and administration of the product.

Readers of this guide are expected to be familiar with:

- ◆ Storage systems in their environment
- ◆ AIX operating system
- ◆ Applications (for example, clustering software) used with PowerPath

Related documentation

The complete set of EMC enterprise storage documentation for PowerPath, available from EMC Corporation, includes:

- ◆ *EMC PowerPath Family for AIX Release Notes*
- ◆ *EMC PowerPath Migration Enabler Version User Guide*
- ◆ *EMC PowerPath Family Product Guide*
- ◆ *EMC PowerPath Family CLI and System Messages Reference Guide*

These documents are updated periodically. Electronic versions of the updated manuals are available on the EMC Online Support site.

The following table provides a list of reference documents for specific storage environments:

If your storage environment includes	Refer to
Symmetrix storage system	<ul style="list-style-type: none">• <i>EMC Host Connectivity Guide for IBM AIX</i>
VNX OE block and CLARiiON storage system	<ul style="list-style-type: none">• <i>EMC Host Connectivity Guide for IBM AIX</i>• CLARiiON Support home page• VNX OE block Support home page

If your storage environment includes	Refer to
Invista system	<ul style="list-style-type: none"> • <i>EMC Invista Installation and Setup Guide</i> • <i>EMC Element Manager Administration Guide</i>
XtremIO storage system	XtremIO Storage Array Users Guide
Other vendors' storage system	The appropriate documentation from your vendor.

All documents are available on the EMC Online Support site.

Another related EMC enterprise storage publication is the *E-Lab Interoperability Navigator*, available on the EMC Online Support site.

Revision History

The following table presents the revision history of this document.

Table 1 Revision History

Revision	Date	Description
01	December 2014	Release of PowerPath 6.0 for AIX.

Conventions used in this guide

EMC uses the following conventions for notes, cautions, warnings, and danger notices.



DANGER indicates a hazardous situation which, if not avoided, will result in death or serious injury.



WARNING indicates a hazardous situation which, if not avoided, could result in death or serious injury.



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



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

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

Typographical conventions

EMC uses the following type style conventions in this guide:

Normal font	<p>In running text:</p> <ul style="list-style-type: none"> • Interface elements (for example, button names, dialog box names) outside of procedures • Items that user selects outside of procedures • Java classes and interface names • Names of resources, attributes, pools, Boolean expressions, buttons, DQL statements, keywords, clauses, environment variables, filenames, functions, menu names, utilities • Pathnames, URLs, filenames, directory names, computer names, links, groups, service keys, file systems, environment variables (for example, command line and text), notifications
Bold	<ul style="list-style-type: none"> • User actions (what the user clicks, presses, or selects) • Interface elements (button names, dialog box names) • Names of keys, commands, programs, scripts, applications, utilities, processes, notifications, system calls, services, applications, and utilities in text
<i>Italic</i>	<ul style="list-style-type: none"> • Book titles • New terms in text • Emphasis in text
Courier	<ul style="list-style-type: none"> • Prompts • System output • Filenames • Pathnames • URLs • Syntax when shown in command line or other examples
Courier bold	<ul style="list-style-type: none"> • User entry • Options in command-line syntax
<i>Courier italic</i>	<ul style="list-style-type: none"> • Arguments in examples of command-line syntax • Variables in examples of screen or file output • Variables in pathnames
⟨⟩	Angle brackets for parameter values (variables) supplied by user.
[]	Square brackets for optional values.
	Vertical bar symbol for alternate selections. The bar means or.
...	Ellipsis for nonessential information omitted from the example.

Where to get help

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

Product information

For documentation, release notes, software updates, or information about EMC products, licensing, and service, go to the EMC Online Support site.

Technical support

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

License key ordering and activation

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

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

Your comments

Your suggestions will help us continue to improve the accuracy, organization, and overall quality of the user publications. Please send your opinion of this guide to:

techpubcomments@EMC.com

If you have issues, comments, or questions about specific information or procedures, please include the title and, if available, the part number, the revision (for example, A01), the page numbers, and any other details that will help us locate the subject you are addressing.

CHAPTER 1

Installing PowerPath on an AIX host

This chapter includes the following sections:

◆ Before you install	12
◆ Installing PowerPath as part of a mksysb master image	16
◆ Installation methods	17
◆ After you install	21
◆ Upgrading PowerPath	25
◆ Upgrading the AIX operating system	30
◆ Coexistence with third-party path management software	32

Before you install

Before installing EMC® PowerPath® review the storage system and the host environments, and take necessary preliminary precautions.

Reviewing the Storage System Environment

The sequence in which you configure the storage system and install PowerPath on a host depends on which storage system you use:

- ◆ **EMC Symmetrix®, EMC XtremIO, Hitachi TagmaStore, HP StorageWorks XP, HP StorageWorks EVA, and IBM ESS systems**

Install PowerPath *after* you have set up the storage system and verified that it is working properly.

- ◆ **EMC VNX OE block and CLARiiON® storage systems**

PowerPath installation is an integral part of an EMC CLARiiON® setup and configuration procedure. The CLARiiON Support home page provides directions on installing PowerPath on a host with a CLARiiON array.

The VNX OE block Support home provides directions on installing PowerPath on a host with a VNX array.

Reviewing the Host Environment

Installation of PowerPath on a host demands removal of any other existing failover applications from EMC. Incorrect migration from other failover applications of EMC to PowerPath results in data loss. Therefore, EMC recommends a migration with the support of EMC Professional Services.

The other failover applications of EMC that may pre-exist could be: EMC Navisphere® Application Transparent Failover (ATF) or Common Desktop Environment (CDE).

Note: For VNX OE block systems and CLARiiON systems, the Navisphere® and the Unisphere applications support limited PowerPath functions.

If you choose to perform the migration on your own, see the *Removing ATF or CDE Software Before Installing Other Failover Software* document. This document is available on the EMC Online Support site.

Obtain up-to-date information

Check the EMC Online Support site for current information:

- ❑ PowerPath documentation — The documentation are updated with new features and fixed bugs.
- ❑ Service packs and upgrades — EMC releases a Service Pack for a set of new features. You can download PowerPath service pack software from EMC Online Support. Determine which service packs (if any) to install after PowerPath, and if that service pack has any additional installation prerequisites.

Locate your license key

The PowerPath license registration key is on the License Key Card that you received from EMC.

Note: Later versions of PowerPath can continue to use the license key of existing PowerPath. While upgrading, PowerPath will convert the existing 12-character license key to a 24-character key and run the application. Therefore, you can avoid repeated registrations for every upgrade.

Prepare the host and storage system

- ❑ Preview the environment. Verify if your environment meets the requirements described in the *EMC PowerPath Family for AIX Release Notes* and in E-Lab™ Interoperability Navigator. To verify the requirements, run PowerPath Configuration Checker (PPCC). The *EMC PowerPath Configuration Checker User Guide* provides detailed information about PPCC installation.
- ❑ Select the correct HBA initiator type. **Disconnect all non-root AIX Logical Volume Managers (LVM).** Before taking the LVMs offline, unmount any of the file system or application that uses the volume groups. Use the **AIX varyoffvg command** to take offline all AIX Logical Volume Manager (LVM) volume groups that use storage-system hdisk devices.

Note: Do not unmount root volume storage group (rootvg).

- ❑ Ensure that the required AIX Licensed Program Products (LPPs) are installed.

[Table 1 on page 13](#) lists and describes the required AIX LPPs.

Table 1 Required AIX LPPs (page 1 of 2)

LPP	Description	Required for
EMC.AIX.TPA.ODM	EMC TPA ODM kit	IBM XIV, HSV 210, HSV 110, HSV 450, HSV 111
EMC.Symmetrix.iscsi.rte	EMC driver kit	Symmetrix iSCSI devices
EMC.CLARiiON.iscsi.rte	EMC driver kit	CLARiiON iSCSI devices
EMC.Symmetrix.aix.rte EMC.CLARiiON.aix.rte EMC.INVISTA.aix.rte EMC.XtremIO.aix.rte	EMC driver kit	Fibre Channel devices
EMC.Symmetrix.fcp.rte EMC.CLARiiON.fcp.rte EMC.INVISTA.fcp.rte EMC.XtremIO.fcp.rte	EMC driver kit	Fibre Channel devices
EMC.Symmetrix.ha.rte EMC.CLARiiON.ha.rte	EMC driver kit Symmetrix or CLARiiON support for PowerHA/HACMP	PowerHA/HACMP clusters
Hitachi.aix.support.rte	Hitachi AIX support for xp disks	All Hitachi Tagmastore configurations

Table 1 Required AIX LPPs (page 2 of 2)

LPP	Description	Required for
HP.aix.support.rte	HP AIX support for disks	All HP StorageWorks xp configurations
ibm2105.rte	IBM 2105 disk device	All IBM ESS configurations
devices.fcp.disk.ibm.rte	IBM 2107 disk device	All IBM ESS configurations

The Symmetrix, VNX OE block and CLARiiON, and Invista LPPs are included in the EMC Object Data Manager (ODM) fileset. The EMC *E-Lab Interoperability Navigator* provides details on determining the ODM fileset version required for your system. You can download the EMC ODM fileset from:

ftp://ftp.EMC.com/pub/elab/aix/ODM_DEFINITIONS/.

Contact your vendor for the Hitachi, HP, and IBM ODM filesets.

- If logging is not enabled on the AIX host, enable it. PowerPath reports errors, diagnostic messages, and failover recovery messages through the syslog file specified by the administrator. EMC recommends using `/tmp/emcpsyslog.log` to ensure the support utility, EMC grab, collects the log files. “[Enabling logging on an AIX host](#)” on [page 77](#) provides more information.
- Ensure that the AIX hdisk devices are configured correctly. Each logical path that PowerPath uses to access a storage system device must have an hdisk configured for it. If the number of storage system hdisk devices is incorrect, complete the following procedure before installing PowerPath.

To configure hdisks:

1. Ensure all physical device connections are connected.
2. Remove the AIX hdisks corresponding to storage system devices.

To remove hdisks corresponding to Symmetrix devices, type:

```
lsdev -CtSYMM* -rname | xargs -n1 rmdev -dl
```

To remove hdisks corresponding to VNX OE block and CLARiiON devices, type:

```
lsdev -CtCLAR* -rname | xargs -n1 rmdev -dl
```

To remove hdisks corresponding to XtremIO devices, type:

```
lsdev -CtXtremIO -rname | xargs -n1 rmdev -dl
```

To remove hdisks corresponding to Hitachi Lightning device, type:

```
lsdev -CtHitachi* -rname | xargs -n1 rmdev -dl
```

To remove HP StorageWorks EVA devices, type:

```
lsdev -CtHSV* -rname | xargs -n1 rmdev -dl
```

To remove hdisks corresponding to IBM ESS devices, type:

```
lsdev -Ct2105* -rname | xargs -n1 rmdev -dl
```

To remove hdisks corresponding to Invista and VPLEX devices, type:

```
lsdev -CtINV* -rname | xargs -n1 rmdev -dl
```

To remove hdisks corresponding to HP StorageWorks XP devices, type:

```
lsdev -CtHP* -rname | xargs -n1 rmdev -dl
```

These commands cannot delete hdisks in use. These hdisks do not need to be removed, and you can ignore any error messages.

Ensure that every path has its hdisks. After all storage system hdisks are removed, run the `usr/lpp/Symmetrix/bin/emc_cfgmgr` script to ensure that hdisks are configured for each path. The **emc_cfgmgr** script invokes the AIX **cfgmgr** tool to probe each adapter bus separately. Running the script configures a storage system hdisk that was missing for each device on each path. You can run **lsdev -Cc disk** to confirm the correct configuration.

The **emc_cfgmgr** script is available at the following EMC FTP site:

ftp://ftp.EMC.com/pub/elab/aix/ODM_DEFINITIONS

- ❑ Configure the HBA drivers.



Follow HBA driver configuration guidelines outlined in the E-Lab Interoperability Navigator and product documentation. Improper settings can cause erratic failover behavior, such as greatly increased I/O delays.

- ❑ EMC recommends setting certain Object Data Manager (ODM) attributes to enable enhanced error recovery during certain failover scenarios. The *EMC Host Connectivity Guide for IBM AIX* on EMC Online Support provides more information on ODM attributes the Fast Fail feature.

Note: This does not apply to Cambex cards.

To configure the ODM attributes:

1. For Fibre Channel switched environments only, enable the adapter driver **fast_fail** feature as follows:

```
chdev -l fscsix -a fc_err_recov=fast_fail -P
```

where *x* is the specific HBA (fscsi0, fscsi1, and so forth).

2. Run **shutdown -Fr** to restart the host. In addition to changing the **fc_err_recov** ODM attribute, the **num_cmd_elems** should also be added:

```
chdev -l fcsx -a num_cmd_elems=2048 -P
```

For NDU it is recommended to enable following two attributes for all native devices along with **fc_err_recov=fast_fail** and **num_cmd_elems=2048**

```
lsattr -El hdisk3 | egrep "clr_q|q_err"
clr_q      no          Device CLEARS its Queue on error True
q_err     yes          Use QERR bit                True
```

3. To set dynamic tracking of devices:

```
chdev -l fscsi0 -a dyntrk=yes -P
```

4. To set the **max_xfer_size** for FC HBAs to 0x1000000:

```
# chdev -l fcs1 -a max_xfer_size=0x1000000 -P
```

5. Set the number of 4k pages in Block IO buffer cache. Where the maximum value is 1000 and the default is 20.

Note: For higher throughput, the recommended *maxbuf* value should be set to between 500-800.

```
# chdev -l sys0 -a maxbuf=500
sys0 changed
```

6. Set the maximum number of PROCESSES allowed per user.

Note: For faster device discovery on scale setups with a higher number of devices or paths, the recommended *maxuproc* value should be set to 1024 with the *maxbuf* value set to 100.

```
bash-4.0# chdev -l sys0 -a maxuproc=1024
sys0 changed
```

7. Reboot the host.

- ❑ Ensure that the ownership and permission attributes of all *hdisk* devices are correct. PowerPath configuration sets the ownership and permission values of each *hdiskpower* device to match the values of one of the path devices.

Installing PowerPath as part of a mksysb master image

You can create a master site backup image from which you can install the complete host images. These host images could include PowerPath software on target systems.

The *hdiskpower* devices and *powerpath0* device from the source system affects the creation of master site backup image. Therefore, remove them before creating the master site backup image.

CAUTION

Failure to remove these devices prior to creating the master site backup image may result in unpredictable behavior on the target host. It could also cause problems that are not immediately obvious, but could affect adding or removing devices in the future.

Removing the *hdiskpower* devices

If the source host is to create the master site backup image from the SAN, under PowerPath control, you must perform the following steps prior to removing the *hdiskpower* devices:

1. Run **pprootdev off**.
2. Restart the host.

If Naviagent is installed, stop the agent prior to removing the *powerpath0* driver device.
Type:


```
/etc/rc.agent stop
```

Removing all the hdiskpower devices and the powerpath0 device

To remove all hdiskpower devices and the powerpath0 device:

1. Deactivate all applications, file systems, volume groups, and paging space located on PowerPath devices.

2. Remove the hdiskpower devices:

```
lsdev -Cdisk -tpower -Fname | xargs -n1 rmdev -dl
```

3. Remove the powerpath0 device driver:

```
rmdev -dl powerpath0
```

4. Create the **mksysb** backup image from this source system.

5. To create the powerpath0 driver and hdiskpower devices, run the **powermt config** command on the target host after it is started.

Note: This procedure removes PowerPath device persistency from the source machine's configuration. This removal could result in hdiskpower devices changing the order when you run the **powermt config** command on the source machine.

Installation methods

You can install PowerPath on a host by using the installation CD, the command line, System Maintenance Interface Tool (SMIT), or the tar file.

Note: Ensure that ODM is installed before installing PowerPath for AIX.

Service packs and patches released on AIX are usually full installable packages. The installation and upgrade procedure of a service pack or patch is the same as that for the Major release on which the service pack or patch is based. For how to install such, see [“Installing a PowerPath service pack or patch” on page 20](#).

If you are upgrading from an earlier release of PowerPath, follow the instructions in [“Upgrading PowerPath” on page 25](#).

This section describes:

- ◆ [“Mounting the installation CD” on page 18](#).
- ◆ [“Installing by using the command line” on page 18](#).
- ◆ [“Installing by using SMIT” on page 19](#).
- ◆ [“Installing by using the tar file” on page 19](#).
- ◆ [“Installing a PowerPath service pack or patch” on page 20](#)

CAUTION

Do not attempt to run `powermt config` while using `installp` or SMIT to add, remove, or update PowerPath software. Doing so can result in a system crash.

2. Download the tar file to **/usr**.
3. Type the following command to extract the installation package:

```
gunzip EMCPower.AIX.<version>.<build>.tar.gz
```

Where <version> is the product version, for example PowerPath 6.0 and <build> is the software build number.

4. Type the following command to extract the *EMCpower_install* installation file:

```
tar -xvf EMCPower.AIX.<version>.<build>.tar
```

5. Before installing, check if all the required **ODM drivers** are available:

```
lslpp -l | grep -i emc
```

This lists all the ODM file sets.

6. Perform one of the following:
 - For existing PowerPath installations, run **rm .toc** from the install directory to delete the .toc file.
 - For new PowerPath installations, skip to [step 7](#).
7. For installing PowerPath, type:

```
installp -agXd . EMCpower
```

8. To view the version of the PowerPath installed, run:

```
powermt version
```

PowerPath is installed on the host, but before PowerPath can run, you must perform some administrative tasks (including registering PowerPath for first time users). [“After you install” on page 21](#) provides post-installation information.

Installing a PowerPath service pack or patch

To install a PowerPath service pack or patch:

1. On EMC Online Support, select **Support > Downloads**.

2. Download the tar file to **/usr**.

3. Extract the installation package. Type:

```
gunzip EMCPower.AIX.<version>.<build>.tar.gz
```

4. Extract **EMCpower_install**. Type:

```
tar -xvf EMCPower.AIX.<version>.<build>.tar
```

5. Perform one of the following:

- For existing PowerPath installations, run **rm .toc** to delete the .toc file.
- For new PowerPath installations, skip to [step 6](#).

6. Install the software:

```
installp -agXd . EMCpower
```

7. Run **powermt config**.

After you install

After installing the PowerPath software:

- ◆ “Register PowerPath on the host” on page 21 (first-time installation only).
- ◆ “Choose a major number” on page 22.
- ◆ “Initialize PowerPath hdiskpower devices” on page 22.
- ◆ “Remove the CD” on page 23.
- ◆ “Configure the management daemon” on page 23.
- ◆ “Vary on volume groups” on page 23.
- ◆ “Reconfigure applications” on page 23.
- ◆ “Reserve_policy setting in Dual VIO Server Environments” on page 24.

Register PowerPath on the host

If you are upgrading from an earlier version of PowerPath, you do not need to reregister. PowerPath will use your existing license key if necessary, converting the old 12-character license key to a 24-character key.

1. To register the PowerPath software, type:

```
emcprep -install
```

Output similar to the following is shown:

```
===== EMC PowerPath Registration =====
Do you have a new registration key or keys to enter? [n]
```

2. Type **y** and press **Enter**.

Output similar to the following is shown:

```
Enter the registration key(s) for your product(s),
one per line, pressing Enter after each key.
After typing all keys, press Enter again.
Key (Enter if done):
```

3. Type the PowerPath registration key and press **Enter**.

If you typed a valid registration key, output similar to the following is shown:

```
1 key(s) successfully added.
Key successfully installed.
```

Key (Enter if done):

If you type an invalid registration key, the screen shows an error message and prompts you to type a valid key. The *EMC PowerPath Family CLI and System Messages Reference Guide* provides a list of error messages returned by the emcprep license registration utility.

4. Press **Enter**.

Output similar to the following is shown:

```
1 key(s) successfully registered.
```

Choose a major number

By default, PowerPath uses a system-assigned major number for its **pseudo devices**. To support NFS applications in a clustered environment, the major number of PowerPath **pseudo devices (for example, /dev/*hdiskpower*) must be the same across all cluster nodes**. This must be a unique major number that is unused by any host in the cluster.

To choose a major number:

1. Determine a free major number common to all nodes. On each node in the cluster, type:

```
lvlstmajor
```

Compare the free ranges on each node to find a major number that is unused on any node.

2. Change the major number of the `powerpath0` device. On each node in the cluster, type:

```
chdev -l powerpath0 -a pp_major=x
```

where *x* is the major number you found in step 1.

3. Verify the change. On each node in the cluster, type:

```
lsattr -El powerpath0
```

The major number will take effect after you run the `powermt config` command on each node in the cluster. [“Initialize PowerPath hdiskpower devices” on page 22](#), which follows, provides additional information.

[“Changing the major number” on page 69](#) provides instructions on changing the major number.

Initialize PowerPath hdiskpower devices

Initialize the **PowerPath hdiskpower devices** and make them available to the host. You can initialize devices by using either command line entries or SMIT.

Using the command line

To initialize devices at the command line, type:

```
powermt config
```

Using SMIT

To initialize devices by using SMIT:

1. Run **smit** to open SMIT.
2. On the **System Management** window, select **Devices**.
3. Select **PowerPath Disk**.
4. Select **Configure All PowerPath Devices**.
5. Exit SMIT.

Remove the CD

The CD should be unmounted before removal.

1. To unmount the CD, type:

```
cd /umount /cdrom
```

2. Remove the CD from the CD-ROM drive.

Configure the management daemon

After installing that version of PowerPath supporting the management daemon, configure the management daemon. Configuring the management daemon allows you to receive SNMP traps when a path is dead for one minute, or when all paths to a device are dead for five minutes. [“Configuring the management daemon” on page 67](#) describes how to configure the management daemon.

Vary on volume groups

Vary on any existing LVM volume groups that you varied off before installing PowerPath:

1. To vary the LVM volume groups, type:

```
varyonvg vgroupname
```

2. Remount any file systems you unmounted.
3. Restart any applications you stopped.

You do not need to reconfigure these volume groups. The installation procedure migrates existing volume groups that use storage system devices from AIX hdisk devices to PowerPath hdiskpower devices.

Note, however, that if you failed to vary off a volume group before installing PowerPath, this migration would fail. A future vary-on attempts may result in errors.

To correct this state:

1. Vary off the volume group.
2. Run `powermt config`.

The `varyonvg` command should now succeed, and the volume group should be using hdiskpower devices.

When defining new volume groups, use PowerPath hdiskpower devices, not AIX hdisk devices.

Reconfigure applications

If an application accesses AIX hdisk devices directly, rather than through a volume group (a DBMS, for example), you must reconfigure that application to use PowerPath hdiskpower devices if you want PowerPath load balancing and path failover functionality. Run `powermt display dev=all` to determine the correspondence between PowerPath hdiskpower devices and AIX hdisk devices.

Note: You do not need to reconfigure applications that access hdisks through a volume group.

When adding new applications to your system that typically would access hdisks directly, configure them to use hdiskpower devices instead.

Reserve_policy setting in Dual VIO Server Environments

From PowerPath Versions 5.5 for AIX onwards, the **reserve_lock** attribute is replaced by the **reserve_policy** and **PR_Key Value** attributes.

On a fresh install of PowerPath 6.x

The **reserve_policy** value is set to **single_path** by default. Dual VIO Server setups with shared disk configurations require the **reserve_policy** attribute to be set to **no_reserve** before configuring them to the VIOS.

To change reserve_policy from 'single_path' to 'no_reserve' run

```
# chdev -l hdiskpower# -a reserve_policy=no_reserve
```

PowerPath supports the **reserve_enable** attribute for the PowerPath control device **powerpath0** which can be used to set the default value for **reserve_policy**.

To set reserve_policy setting to 'no_reserve' for all the devices by default run

1. Bring all the hdiskpower devices to Defined state:

```
# lsdev -Cc disk -t power -F name | xargs -n1 rmdev -l
```

2. Set **reserve_enable** attribute to **no**:

```
# chdev -l powerpath0 -a reserve_enable=no
```

3. Run **powermt config** to reconfigure all the devices:

```
# powermt config
```

For information on how to set 'reserve_policy' to 'no_reserve' before running **powermt config**, see [“Reserve_policy setting in Dual VIO Server Environments” on page 24](#).

Deleting redundant entries

During a fresh installation or upgrade, PowerPath deletes redundant entries from the ODM PdAt for array uniquetypes and retains appropriate entries.

When **powermt config** is run, redundant reserve attribute entries for each PowerPath managed uniquetypes are deleted, and **reserve_lock** attributes are converted to **reserve_policy**.

During installation or upgrade or when **powermt config** is run, the following warning message is sent to syslog (if configured) if PowerPath detects the presence of multiple entries for a uniquetype in PdDV:

```
WARNING!!! ODM PdDv database contains duplicate uniquetypes. Refer to
EMC Knowledge Base Article Number: 000084491 or contact EMC Customer
Support personnel.
```


Upgrading PowerPath

Based on your environment and requirements, you can upgrade PowerPath using:

- ◆ Full package installation: You can directly upgrade PowerPath from a previous version with the later version of PowerPath package.

Before upgrading PowerPath

- Download the latest version of EMC Grabs, available on EMC Online Support, and then run the PowerPath Configuration Checker (PPCC).
- Check the EMC Online Support site, for the most current release notes and service packs.
- You do not need to remove the previous version of PowerPath before upgrading to PowerPath 6.x.

During an upgrade

Do not add devices to your PowerPath configuration.

After an upgrade

You do not need to restart the host after the upgrade, if you have closed all applications that use PowerPath devices before you install PowerPath 6.x.

- You do not need to re-enter license information for PowerPath multipathing. PowerPath will use your existing license key, if necessary, converting the old 12-character license key to a 24-character key.
- If you are using PowerPath Migration Enabler and do not reboot after PowerPath upgrade, run **powermig transition -startup**.

Note: Disabling the cluster services before upgrading PowerPath is not a mandatory requirement. You can ensure continuous availability of the services despite the disk being closed for the PowerPath upgrade. For continuous service availability, fail over the resources of the disk to another node.

Upgrade procedure

Note: When upgrading in a dual VIO server environment, follow the procedure described in [“Upgrading PowerPath in a VIO server environment.”](#) When upgrading in PowerHA/HACMP configurations, follow the procedure described in [“Integrating PowerHA/HACMP in a PowerPath environment” on page 55.](#)

1. Close all applications that use PowerPath devices, and vary off all volume groups except the root volume group (rootvg).

Note: If any application is holding a PowerPath device open, the upgrade will fail. Note that certain applications such as EMC ControlCenter® agents or the Navisphere agent running on a host attached to a VNX OE block and CLARiiON system may have PowerPath devices open; and, this will cause the upgrade to fail.

- ❑ If EMC Solutions Enabler daemons, such as `storwatchd` and `storapid`, are running, run **stordaeomon shutdown** `<daemon>` to stop the daemons before you install PowerPath.
 - ❑ If EMC ControlCenter is running on the host, stop the ControlCenter agents.
 - ❑ In a VNX OE block and CLARiiON environment, if the Navisphere Host Agent is running, type **/etc/rc.agent stop**, to stop the agents before uninstalling PowerPath.
2. Install PowerPath as described in [“Installation methods” on page 17](#).
 3. When upgrading from PowerPath 5.3 SP1/5.5, the configuration data from the `powermt.custom` file needs to be loaded to the default `/etc/powermt_custom.xml` file. To do this, specify **file =/etc/powermt.custom** in the **powermt load** command. If this is not specified, the configuration data is loaded from the default `.xml` file.
 4. Run **powermt config**.
 5. Run **powermt save** to save configuration data to the `/etc/powermt_custom.xml` file.

When you upgrade from an unlicensed to a licensed version of PowerPath, the load balancing and failover device policy is set to `bf/nr` (BasicFailover/NoRedirect). You can change the policy by using the **powermt set policy** command.

Upgrading PowerPath in a VIO server environment

You can perform an upgrade with a full package installation or by uninstalling the existing version and then installing the later version. Each of these methods differently affects the mapping information in a Virtual I/O (VIO) server environment. Therefore, based on your choice, the upgrading procedure also differs in a VIOS environment.

When you upgrade by using the full package installation, the mapping information is maintained by default.

However, if you choose to uninstall PowerPath 5.3 SP1 and install PowerPath 6.x, you need to follow this procedure:

1. Before uninstalling PowerPath 5.3 SP1, back up the mapping information first.
2. Unmap the devices. Follow the procedure in [“Unmapping PowerPath” on page 27](#).
3. Uninstall PowerPath 5.3 SP1.
4. Install PowerPath 6.x.
5. Use the stored mapping information and map the devices.

Reserve policy setting

Upgrading PowerPath from 5.3 and its minor releases to PowerPath 6.x will change the **reserve_lock** attribute of the devices to **reserve_policy** and **PR_Key Value**. After the upgrade, the default value of the **reserve_policy** attribute will be set to **single_path**. Dual VIO Server setups with shared disk configurations require the **reserve_policy** attribute to be set to **no_reserve** before configuring them to the VIO Client.

For information on how to set 'reserve_policy' to 'no_reserve' before running **powermt config**, see [“Reserve_policy setting in Dual VIO Server Environments” on page 24](#).

Unmapping PowerPath

In the second method, you can back up the mapping information and unmap the devices either manually or by using the VIOS Map Manager tool.

Using VIOS Map Manager tool

The VIOS Map Manager tool is useful in an environment where a large number of servers and virtual disks are mapped. This tool helps you back up the huge amount of mappings before uninstalling PowerPath 5.3 SP1 or later, and to restore the same mappings after installing the later version.

The tool is available for download along with PowerPath packages on EMC Online Support. The package has a readme file that provides the usage details to back up and unmap the mapping information.

Note: This tool is not applicable on PowerPath 5.3 SP1 that uses LUNs from VMAX.

Manual unmapping and upgrading

To manually upgrade PowerPath in a VIO server environment:

1. On one of the VIO servers, run **lsmap -all**.

This command shows the mapping among physical, logical, and virtual devices.

Output similar to the following is displayed:

```
$ lsmap -all
SVSA
-----
vhost1          Physloc          Client Partition ID
                U8203.E4A.069BE95-V2-C12  0x00000007

VTD             vtscsi0
Status          Available
LUN             0x8100000000000000
Backing device  hdiskpower3
Physloc         U789C.001.DQD6R05-P1-C2-T2-L11
Mirrored        false

VTD             vtscsi1
Status          Available
LUN             0x8200000000000000
Backing device  hdiskpower5
Physloc         U789C.001.DQD6R05-P1-C2-T2-L17
Mirrored        false
```

2. Log in on the same VIO server as the **padmin** user.
3. Unconfigure all the VTDS using hdiskpower devices identified in step 1, by running:

```
# rmdev -dev <VTD> -ucfg
```

where **<VTD>** is the virtual target device.

For example: `rmdev -dev vtscsi1 -ucfg`

The VTD status changes to Defined.

Note: Run `rmdev -dev <VTD> -ucfg` for all VTDS identified in step 1.

4. Upgrade PowerPath as described in [“Upgrade procedure” on page 25](#).

5. Run **powermt config**.
6. Run **cfgmgr**
7. Run **lspath** on all clients to verify all paths are **Enabled**.
8. Perform steps 1 through 5 on the second VIO server.
9. At times the path status may be displayed as **Disabled** even if the 'lsmap-all' output in the VIOS is displayed as **Available**. If this happens, run the following:
 - a. `#chpath -l <hdisk#> -p <vscsi#> -s disable`
 - b. `#chpath -l <hdisk#> -p <vscsi#> -s enable`

Upgrading PowerPath on PowerHA7.1

You can upgrade PowerPath from version 5.5 or later to 6.x in an PowerHA7.1 cluster where two or more hosts run PowerPath.

To upgrade PowerPath on PowerHA7.1:

1. Select the node where you want to upgrade PowerPath.
 - a. Stop PowerHA Cluster Services through smitty on the passive node.
 - b. Enter `smitty hacmp`.
 - c. Select **System Management (CSPOC) > PowerHA System Mirror Services > Stop Cluster Services**
 - d. Press Enter.

Note: Ensure that cluster services are stopped and verify the status using **lssrc -ls clstrmgrES**

2. Set the passive node to maintenance mode, run:

```
clctrl -stop -n <clustername> -m <nodename>
```

Example: `clctrl -stop -n powerha_cluster -m aix154.emc.com`

Where:

- `clctrl` is the system administration function for managing a cluster
- `powerha_cluster` is the cluster name
- `aix154.emc.com` is the node name

This will deactivate the **Caavg_private** on the passive node.

3. Ensure that Naviagent and stordaemons are stopped.
4. Upgrade PowerPath on the passive node. Use either:


```
smit update_all or installp -agXd EMCpower_install all
```
5. Run **cfgmgr -v** and **powermt config**, after PowerPath is upgraded.
6. Bring cluster back to normal mode, on the passive node, run the command:

```
clctrl -start -n <clustername> -m <nodename>
```

Example: `clctrl -start -n powerha_cluster -m aix154.emc.com`

This will activate the **Caavg_private** on the passive node.

7. Start cluster services on the passive node.
 - a. Enter `smitty hacmp`.
 - b. Select **System Management (CSPOC) > PowerHA System Mirror Services > Start Cluster Services**
 - c. Press Enter.
8. Move the resource groups from an active node to passive node. Following are the steps:
 - a. Enter `smitty hacmp`.
 - b. Select **System Management (CSPOC) > Resource Groups and Applications > Move a Resource Group to Another Node / Site > Move Resource Groups to Another Node**
9. After resource groups get online on the passive node, repeat the same steps for all other nodes where PowerPath needs to be upgraded.

Upgrading PowerPath in a PowerHA/HACMP concurrent VG environment

When upgrading PowerPath for AIX on PowerHA/HACMP clusters, it is recommended to use concurrent volume groups (VGs). PowerPath for AIX does not support seamless upgrades on PowerHA/HACMP configuration with non-concurrent volume groups. Recommended setting is either `reserve_policy=no_reserve` or `reserve_lock=no`

1. Fail over all resource groups to one node (For example, node A).
2. Stop PowerHA/HACMP on node B.
3. Stop the following daemons and agents:
 - a. Vary off all volume groups on node B (except rootvg).
 - b. If EMC Solutions Enabler daemons are running, shutdown the daemons by running:


```
stordaeomon shutdown <daemon>
```
 - c. If NaviSphere Host Agent is running, stop the agents by running:


```
/etc/rc.agent stop
```
4. Install Powerpath following the steps described in [“Upgrading PowerPath” on page 25](#).
5. Bring up PowerHA/HACMP on this cluster.
6. Fail over all resource groups to node B and repeat steps 2 - 5 for node A.

Troubleshooting the upgrade

Upgrading from PowerPath 5.5 or later could fail for the following reasons:

- ◆ Not all applications that use PowerPath devices were closed before starting the upgrade procedure.

If you failed to close applications

If you failed to close all applications that use PowerPath devices before starting the upgrade, you would see a message indicating that the driver could not be uninstalled or a message stating that the Navisphere agent was still running.

In this case, stop all applications that use PowerPath devices and then restart the upgrade procedure.

If a PowerPath device is the boot device

If your previous PowerPath version uses a PowerPath device as the boot device, disable multipathing to the root device before upgrading.

1. Run `pprootdev off`.
2. Restart the host.

After upgrading, enable multipathing to the root device again:

1. Run `pprootdev on`.
2. Run `pprootdev fix` and `bosboot -ad /dev/ipldevice`.
3. Reboot the host.

Upgrading the AIX operating system

You can upgrade AIX operating systems on a host:

- ◆ AIX migration upgrade
- ◆ AIX Technology Level Upgrades or Service Pack Upgrades
- ◆ Upgrading AIX operating system using `nimadm` command

Upgrading AIX operating systems for major releases

When upgrading the AIX operating system on a host that has PowerPath installed:

1. Uninstall PowerPath, as described in [Chapter 6, “Removing PowerPath from an AIX Host.”](#)
2. Upgrade the AIX operating system. Refer to the relevant AIX documentation.
3. Reinstall PowerPath, as described in [“Installation methods” on page 17.](#)

Upgrading AIX operating systems for Technology Level upgrades or Service Pack upgrades

When upgrading the AIX operating system for a Technology Level upgrade or Service Pack on a host that has PowerPath installed, then:

1. Do not uninstall PowerPath.
2. Upgrade the AIX operating system. Refer to the relevant AIX documentation.

Note: In case a PowerPath hdiskpower device is a boot device, **pprootdev fix** needs to be run before the upgrade, for the bosboot command to function correctly.

Upgrading AIX operating system using nimadm command

1. Identify the large hdiskpowerX device for the AIX upgrade. The device is on the target host and not on the NIM master. This device will be unmanaged in the next step.

For example, **powermt display dev=hdisk80**

Note: If a PowerPath hdiskpowerX device is a boot device, **pprootdev fix** must be run before the nimadm migration upgrade in order for the bosboot command to function correctly.

2. Run **powermt unmanage** to unmanage the device from PowerPath control.

For example, **powermt unmanage dev=hdisk80**

Note: You can clear pvids for all the native disks of the chosen hdiskpowerX using **chdev -l <hdisk#> -a pv=clear**.

3. On the NIM server, create a Volume Group to be used for the upgrade.
4. On the NIM server, run the **nimadm** command, specifying the target host (for example, hdisk80), to upgrade to AIX 7.1. The **nimadm** command executes 12 phases. If an error is encountered during any phase, the issue needs to be resolved and the **nimadm** command should be repeated after a cleanup.
5. Run **shutdown -Fr** to reboot the device listed in step 1. The bootlist is now set to the hdisk. After reboot, the device is upgraded to AIX 7.1.
6. Run **powermt manage dev** to manage the hdisk.
For example, **powermt manage dev=hdisk80**
7. In a Boot from SAN environment, run **shutdown -Fr** to reboot the system from the PowerPath hdiskpower pseudo device. After reboot the host is booted from SAN and the AIX level is 7.1. In a non-Boot from SAN environment, reboot is not required.
8. Upgrade ODM if required, for example when upgrading from AIX 5.3 or 6.1 to AIX 7.1.

Note: If upgrade fails in a BFS environment, run the following and then upgrade ODM.

```
pprootdev on
pprootdev fix
bosboot -ad /dev/ipldevice
```

Coexistence with third-party path management software

PowerPath can coexist with the following third-party path management software:

- ◆ Hitachi Dynamic Link Manager (HDLM)
- ◆ IBM Subsystem Device Driver (SDD)
- ◆ HP StorageWorks Secure Path
- ◆ HP StorageWorks AutoPath XP Virtual Array

However, PowerPath cannot co-manage devices with third-party path management software. During installation of PowerPath, the installation script tests for the presence of third-party path management software on the system. If such software is installed on the system, the PowerPath installation script disables support for the corresponding array type in PowerPath.

[Table 2 on page 32](#) lists the third-party path management software and the corresponding array disabled during the installation.

Table 2 Third-party software and corresponding array support

If this software is present	Support for this class is disabled in PowerPath
Hitachi HDLM	hitachi
IBM SDD	ess
HP StorageWorks XP AutoPath	hpxp
Sun Solaris MPxIO	Class(es) that would ordinarily handle the MPxIO-configured array(s)

[Table 3 on page 32](#) lists the software and corresponding command that results in undefined multipathing behavior.

Table 3 Third-party software and corresponding powermt manage class command

If this software is present	Multipathing behavior is undefined if you execute this command
Hitachi HDLM	<code>powermt manage class=hitachi</code>
IBM SDD	<code>powermt manage class=ess</code>
HP StorageWorks XP AutoPath	<code>powermt manage class=hpxp</code>
Sun Solaris MPxIO	<code>powermt manage class=<class></code> Where <class> is any class that handles an MPxIO-configured array

Because PowerPath and the third-party software cannot co-manage devices, do not initiate co-management by executing the above commands on third-party array class machines. As long as these commands are not run, support for the relevant third-party arrays will remain disabled across restarts.

Similarly, before you install third-party path management software on a system on which PowerPath is installed, you should disable any support by PowerPath for the relevant third-party array devices by running `powermt unmanage class=<class>`.

[Table 4 on page 33](#) lists the commands to run before you install third-party path management software.

Table 4 Commands to run before installing third-party software

Before you install this software	Execute this command
Hitachi HDLM	<code>powermt unmanage class=hitachi</code>
IBM SDD	<code>powermt unmanage class=ess</code>
HP StorageWorks XP AutoPath	<code>powermt unmanage class=hpxp</code>

CHAPTER 2

Installing PowerPath on Veritas Volume Manager

This chapter includes the following sections:

- ◆ [Installing with Veritas Volume Manager on new hosts.....](#) 36
- ◆ [Converting to PowerPath from DMP](#) 36

Installing with Veritas Volume Manager on new hosts

When installing PowerPath and Veritas Volume Manager (VxVM) for the first time on a host attached to an EMC or third-party storage system, have the Veritas documentation available for reference.

To install PowerPath and Veritas Volume Manager:

1. Install and configure PowerPath following the instructions in [Chapter 1, “Installing PowerPath on an AIX host.”](#)
2. Install Veritas Volume Manager following the instructions in the Veritas documentation.

You must run the `vxdisk list <PowerPath pseudo device>` command to display the power devices. For MP1 and later, the power devices appear as components of a particular VxDMP node in the `vxdisk list <PowerPath pseudo device>` output as follows:

```
# vxdisk list hdiskpower10
Device:      hdiskpower10
devicetag:   hdiskpower10
type:        auto
info:        format=none
flags:       online ready private autoconfig invalid
pubpaths:    block=/dev/vx/dmp/hdiskpower10
              char=/dev/vx/rdmp/hdiskpower10
guid:        -
udid:        DGC%5FVRAID%5FAPM00114002720%5F6006016075802D00A4188133E1B3E111
site:        -
Multipathing information:
numpaths:    1
hdiskpower10 state=enabled
```

Converting to PowerPath from DMP

This section describes how to reconfigure an existing Veritas Volume Manager host configuration to use PowerPath hdiskpower devices instead of DMP devices, without loss of data.

For each Veritas Volume Manager disk group:

1. Unmount all volumes in the disk group.
2. Stop all volumes in the disk group. Type:

```
vxvol -g <disk_group> stopall
```

3. Restart the host. Type:

```
reboot
```

Veritas Volume Manager volumes are now present on PowerPath hdiskpower devices.

4. Start all the volumes in the disk group. Type:

```
vxvol -g <disk_group> startall
```

5. Remount the volumes.

Adding a new path

To add a new path to a storage system controlled by PowerPath and Veritas Volume Manager:

1. Run `cfgmgr`.
2. Run `powermt config`.
3. Run `vxdctl enable`.

CHAPTER 3

Configuring a PowerPath Boot Device on AIX

This chapter includes the following sections:

- ◆ [Setting up a PowerPath boot device](#) 38
- ◆ [PowerPath and the bosboot tool](#) 39
- ◆ [Configuring a new PowerPath installation](#) 39
- ◆ [Configuring an existing PowerPath installation](#) 41
- ◆ [Create a rootvg mirror on AIX](#) 44
- ◆ [Adding paths](#) 45
- ◆ [Disabling PowerPath on a storage system boot device.....](#) 45

Setting up a PowerPath boot device

On some storage systems, you can use a PowerPath hdiskpower device as a boot device (a device that contains the startup image). Using a PowerPath hdiskpower device as a boot device provides load balancing and path failover for the boot device.

Consult the *PowerPath for AIX Release Notes* to find out if your storage system supports PowerPath boot devices.

When you set up a PowerPath boot device, consider the following:

- ◆ All path devices that make up the hdiskpower device must be valid AIX boot devices.
- ◆ The boot device should not be visible to any other host attached to the same storage system. If using a storage system device as a boot device in a cluster environment (with or without PowerPath), other hosts should not be able to address the boot device.
- ◆ The host's boot list must contain all hdisks that compose the hdiskpower device being used as the boot device. Otherwise, the host may fail to boot if one or more paths are disabled while the machine tries to boot.

[“Considerations for VNX OE and CLARiiON storage” on page 38](#) contains additional information if you plan to configure a VNX OE and CLARiiON boot device.

- ◆ At startup, the system searches for an AIX boot image in the boot list, a list of hdisks stored in the hardware's NVRAM. If the system fails to boot, you can change the boot list.

Use one of the following methods:

- Boot the system from an installation device (CD or tape) into Maintenance Mode. Select the option to access the root volume group, and then run the AIX `bootlist` command from the shell.
- Enter the **System Management Services** menu when the system starts, and use the **Multiboot** menu options to change the boot list. This method is faster, but it is more difficult to determine the correspondence between devices listed in the menu and the storage system device you want to add to or remove from the boot list.

Considerations for Symmetrix storage

In a Symmetrix environment, booting from a PowerPath device is supported in Fibre Channel environments that include specific versions of EMC Enginuity™ software. The E-Lab Interoperability Navigator provides details. Contact your EMC Customer Support Representative for information about installing Enginuity software.

Considerations for VNX OE and CLARiiON storage

PowerPath can be used to enable multipathing and failover to an external boot device on a VNX series and CLARiiON array, but such a configuration has some functional limitations and extra configuration steps.

The primary limitation to VNX series and CLARiiON boot is that, for the bootlist search for a suitable boot device to succeed, the boot LUN must either be located on its owner SP or be completely trespassed to its secondary SP. If the boot path to the host is lost while in the process of booting, the host may fail to boot until access to the LUN through the default SP is restored.

In a VNX series and CLARiiON environment, the bootlist for the AIX host should include all hdisks that correspond to active and passive paths of the boot LUN on all associated SPs. If PowerPath is not configured, these devices will show a PVID in the output of the `lspv` command. Passive hdisks will show a PVID of `None`. If PowerPath is configured, use the `powermt display dev=device/all` command to examine the boot device and determine which hdisks are part of the boot LUN.

Once the AIX host is up and running, PowerPath will enable it to survive path failures and trespasses of the boot device.

PowerPath and the bosboot tool

After a system boots from a PowerPath device, the **bosboot** tool cannot function correctly. This is because of the state of the configuration after booting from a PowerPath device and the fact that **bosboot** expects the boot device to be an hdisk, not an hdiskpower device.

The **pprootdev fix** command corrects the configuration to allow **bosboot** to work. Run **pprootdev fix** before undertaking any administrative task that runs **bosboot**. This corrects the configuration for **bosboot** but does not change the PowerPath boot switch. The next system boot still uses PowerPath. Run *pprootdev fix* only once after the system that uses PowerPath restarts. **bosboot** then functions correctly until the system restarts again.

Note: Always use the command **bosboot -ad /dev/ipldevice** to create a boot image from the current boot kernel disk.

VNX OE and CLARiiON trespasses

When a trespass occurs in a VNX OE and CLARiiON environment, a passive interface becomes the active interface. In this situation, **bosboot** will fail unless you transfer the rootvg PVID to the newly active interface. To do so, run the command `emcpassive2active`. Once you have run the command, **bosboot** will succeed.

Run `emcpassive2active` whenever a trespass occurs.

Configuring a new PowerPath installation

If the system contains sufficient internal storage, install and configure the operating system on the internal devices. Use the procedure described in [“Configuring an existing PowerPath installation” on page 41](#) to clone the operating system image on the storage system.

If there is insufficient internal storage, use the following procedure to install AIX directly onto a storage system device and use PowerPath to manage multiple paths to the root volume group.

1. Start with a single connection to the storage subsystem. If you are using a switch, only one logical path should be configured.
2. Install AIX on a storage system device that is accessed through a fibre adapter.
3. Install the current storage system drivers.
4. Restart the host.
5. Use `rmdev -dl` to delete any hdisks in the `Defined` state.
6. Install PowerPath.
7. Connect the remaining physical connections between the host and the storage system. If you are using a switch, update the zone definitions to the new configuration.
8. Configure an hdisk for each path.
9. Run `powermt config`.
10. Run `pprootdev on` to set up multipathing to the root device.
11. Run `bootlist` to add all active **rootvg** storage devices to the boot list.
12. Restart the host.

Configuring an existing PowerPath installation

This section describes the process for converting a system that has AIX installed on an internal disk to boot from a logical device on a storage system.

This process involves:

- ◆ Transferring a complete copy of the operating system from an internal disk to a logical device on the storage system.
- ◆ Configuring PowerPath so the root volume group takes advantage of multipathing and failover capabilities.

EMC recommends that you use this procedure. In the event of a problem, you can revert operations to the host's internal disks.

Updating the system microcode

To update the system microcode:

1. Run **uname -M** to find the system model name.

Output similar to the following appears:

```
IBM, 7025-F80
```

2. Run **lscfg -vp | grep alterable** to find the system microcode level.

Output similar to the following appears:

```
ROM Level. (alterable) . . . . . M2P020806
ROM Level. (alterable) . . . . . CL020807
```

3. Download the latest available microcode from the IBM website.
4. After the system restarts, run **lscfg -vp | grep alterable** to confirm the microcode level.
5. Run **bootinfo -B hdiskx** to confirm that the external disks are now bootable.
 - Return code 1 indicates the disk is bootable.
 - Return code 0 indicates the disk is not bootable.

Installing PowerPath

To install PowerPath:

1. Verify that all device connections to the storage system are established.
2. Verify that all hdisks are configured properly, as described in [“Before you install” on page 12](#).
3. Install PowerPath as described in [“Installation methods” on page 17](#).
4. Run **powermt config** to configure the PowerPath driver.

Cloning the boot device

To clone the boot device:

1. Ensure that the AIX disk includes a copy of the Alternate Disk Installation, `bos.alt_disk_install` by running:

```
lslpp -aL | grep bos.alt
```

If the disk does not include `bos.alt_disk_install`, install it from the AIX installation CD by running:

```
installp -agXd -d/dev/cd0 bos.alt_disk_install
```

2. Locate drives with adequate space, by running:

```
bootinfo -s hdiskx
```

For example, assume hdisks 132-134 are adequate with 8 GB total space.

3. Run **powermt display dev=hdisk132** to determine which `hdiskpower` device contains `hdisk132` (the first `hdisk` identified in [step 2](#) as well as all the path `hdisks` for that `hdiskpower`).

Output similar to the following is displayed:

```
Pseudo name=hdiskpower38
Symmetrix ID=000100006216
Logical device ID=006C
state=alive; policy=SymmOpt; queued-IOS=0
=====
----- Host ----- - Stor - -- I/O Path - -- Stats ---
### HW Path          I/O Paths   Interf.  Mode   State  Q-IOS Errors
=====
  0 fscsi0             hdisk132   FA 14bA  active alive    0    0
  1 fscsi1             hdisk223   FA 14bB  active alive    0    0
  1 fscsi1             hdisk314   FA 14bA  active alive    0    0
  0 fscsi0             hdisk41    FA 14bB  active alive    0    0
```

`hdiskpower38` contains `hdisk132` and that the path `hdisks` for `hdiskpower38` are `hdisk132`, `hdisk223`, `hdisk314`, and `hdisk41`.

4. Record the `hdisk` device to receive a copy of `rootvg`.
5. Run `powermt config`.
6. If the new boot device is on a VNX OE and CLARiiON storage system, run `powermt restore` to ensure that all LUNs are active on the default SP.
7. Move all the `hdiskpower` devices from Available to Defined state, by running:

```
lsdev -Ct power | awk '{print $1}' | xargs -n1 rmdev -l
```

8. Clear PVID of all the native disks (`hdisks` identified in [step 3](#)) under the chosen PowerPath pseudo device by running the following command on each of them.

```
chdev -l hdisk# -a pv=clear
```

9. Create a copy of the operating system on the storage system `hdisks` identified in [step 2](#):

```
alt_disk_install -C hdisk132
```

Note: In this example, multiple devices may be used to offer increased capacity.

Output similar to the following is displayed:

```
Calling mkszfile to create new /image.data file.
Checking disk sizes.
Creating cloned rootvg volume group and associated logical
volumes.
. . .
Bootlist is set to the boot disk: hdisk132
```

10. Restart the host, by running **shutdown -Fr**.
11. Specify that all path hdisks identified in [step 3](#) are included in the bootlist, by running:

```
bootlist -m normal hdisk132 hdisk223 hdisk314 hdisk41
```

Note: If you are booting from a VNX series and CLARiiON storage system, the bootlist for the AIX host should include all hdisks that correspond to active and passive paths of the boot LUN on all associated SPs. [“Considerations for VNX OE and CLARiiON storage” on page 38](#) provides more information.

12. Run `pprootdev on` to enable PowerPath boot support.
13. Restart the host, by running **shutdown -Fr**. When the system starts up, rootvg is using hdiskpower devices.
14. Run `pprootdev fix`.
15. Run `bosboot -ad /dev/ipldevice`
16. Reboot the host.
17. Run `bootlist -m normal -o` to check the boot device (hd5).

Troubleshooting

If ghost devices are created:

Note: Run these commands after the first boot with PowerPath. You do not need to reboot after these commands run.

1. Run `powermt check` and, when prompted, select a to remove all dead paths.
2. Run `powermt restore`.
3. Run `pprootdev fix`.
4. Run `rmdev -dl hdisk#` to delete all the Defined ghost devices.
5. Run `savebase`.
6. Run `bosboot -ad /dev/ipldevice`.

Create a rootvg mirror on AIX

Install PowerPath and configure a boot device

1. Install PowerPath on the host.
2. Start the Boot from SAN procedure on the host installed with PowerPath.
3. Run **pprootdev on** and then reboot to configure the boot device under PowerPath control.

Create a Rootvg mirror

1. Identify one or more PowerPath pseudo devices which could be used for mirroring the existing root device. Extend the root volume group (rootvg) to include these devices. The new device should have same size of the source / existing device in rootvg.

```
# extendvg rootvg <new_pseudo_device>
```

2. (optional) Run **lsvg -p rootvg** to verify rootvg extended successfully. You should see both the physical volumes. If rootvg did not extend successfully, then check for a device issue.
3. Mirror rootvg to the newly added device and then verify mirroring has been completed with 100% sync.

```
# mirrorvg -S rootvg <new_pseudo_device>
# lsvg -l rootvg
```

4. Run **pprootdev fix**. This enables bosboot to function correctly after booting from a PowerPath device.
5. Run **bosboot**.

```
# bosboot -ad /dev/ipldevice
```

PVID mismatch and Usage of 'pprootdev fixback'

1. In case of rootvg device failed it could be identified in **lsvg** with the missing physical volume identifier (PVID).

```
#lsvg -p rootvg
```

Note: powermt display should show all the paths to this device as dead in case device failed and **lsvg -p rootvg** should show PV missing.

2. When the failed device is online there would be PVID mismatch while running **varyonvg**. This should fail with multiple/identical PVID errors.

```
# varyonvg rootvg
```

3. Run **pprootdev fixback**. Fixback enables varyonvg of rootvg in a mirrored rootvg environment as this will transfer the PVIDs onto the hdiskpower device in the ODM.
4. Run **varyonvg**. varyonvg rootvg should succeed without any error.

Adding paths

To add one or more paths while booting off the array:

Note: If you fail to follow this procedure, the host may hang or fail to boot to the PowerPath pseudo device.

1. Run **pprootdev fix**.
2. Add one or more paths.
3. Run **cfgmgr**.
4. Run **powermt config**.
5. Run **savebase**.
6. Run **bosboot -ad /dev/ipldevice**.
7. Modify the bootlist to include the additional devices.
8. Restart the host.

Disabling PowerPath on a storage system boot device

To disable PowerPath load balancing and failover functionality on a storage system boot device:

1. Run **pprootdev off** to turn off multipathing to the root device.
2. Restart the host.

Booting an AIX host from a SAN device

To boot an AIX hosts from a Clariion LUN device:

1. Run **powermt display dev=hdiskpower2 (20Gig)**.

Output similar to the following is displayed:

```
Pseudo name=hdiskpower2
CLARiiON ID=APM00084202720 [libk029_11:12]
Logical device ID=60060160C1502200589C243A15C2DD11 [LUN 150]
state=alive; policy=BasicFailover; priority=0; queued-IOs=0
Owner: default=SP B, current=SP B          Array failover mode: 4
=====
----- Host ----- - Stor - -- I/O Path - -- Stats ---
### HW Path          I/O Paths  Interf.  Mode   State  Q-IOs  Errors
=====
  0 fscsi0            hdisk102  SP A0    active  alive   0       0
  0 fscsi0            hdisk125  SP B1    active  alive   0       0
  1 fscsi1            hdisk148  SP A1    unlic   alive   0       0
  1 fscsi1            hdisk171  SP B0    unlic   alive   0       0
```

2. Run **lsdev -Ct power | awk '{print \$1}' | xargs -n1 rmdev -l** to remove PowerPath control.
3. Run the following commands:
chdev -l hdisk102 -a pv=clear

```
chdev -l hdisk125 -a pv=clear
```

```
chdev -l hdisk148 -a pv=clear
```

```
chdev -l hdisk171 -a pv=clear
```

4. Run **alt_disk_install -C hdisk102**.
5. Run **shutdown -Fr**.
6. Run **bootlist -m normal hdisk102 hdisk148 hdisk171 hdisk125** to add any native disks.
7. Run **pprootdev on** to enable PowerPath boot support.
8. Reboot the host.
9. Run **pprootdev fix**.
10. Run **bosboot -ad /dev/ipldevice**.
11. Reboot the host.
12. Run **bootlist -m normal -o** to check the boot device (hd5).
13. Run **lslv -l hd5** to list the current physical boot device.

```
ipl_varyon -i
```

Output similar to the following is displayed:

```
-bash-4.0# powermt display dev=hdisk359
Pseudo name=hdiskpower40
Symmetrix ID=000192601710
Logical device ID=4AE0
state=alive; policy=SymmOpt; priority=0; queued-I/Os=0;
=====
----- Host ----- - Stor - -- I/O Path -- -- Stats ---
### HW Path          I/O Paths  Interf.  Mode   State  Q-I/Os  Errors
=====
  1 fscsi0             hdisk359  FA  9gA   active  alive    0    0
  1 fscsi0             hdisk461  FA  7gA   active  alive    0    0
  0 fscsi1             hdisk617  FA  9gA   active  alive    0    0
  0 fscsi1             hdisk619  FA  7gA   active  alive    0    0
```


CHAPTER 4

Configuring a PowerPath Logical Device

This chapter includes the following sections:

- ◆ Introduction 48
- ◆ Adding new paths to a logical device..... 48
- ◆ Adding new logical devices 50
- ◆ Removing paths or logical devices..... 51

Introduction

You must reconfigure PowerPath after making configuration changes that affect host-to-storage-system connectivity or logical device identification.

For example:

- ◆ Fibre Channel switch zone changes
- ◆ Adding or removing Fibre Channel switches
- ◆ Adding or removing HBAs or storage-system ports
- ◆ Adding or removing logical devices
- ◆ Redeploying existing logical devices

Note: If you do not reconfigure PowerPath after making configuration changes, many of these changes are treated as unintentional, and PowerPath tries to prevent them from affecting applications.

Many platform-specific and some PowerPath reconfiguration operations will fail if the affected logical devices are in use (either just marked alive or with I/O in progress).

After any reconfiguration, you must monitor the outcome of individual reconfiguration steps and confirm that the resulting changes are as expected, before relying on the new configuration. Otherwise, some paths may not be as expected. For example, on platforms that support native devices, it is possible to inadvertently write to the wrong logical device (thinking that the native device is associated with a different logical device).

Adding new paths to a logical device

This procedure adds new paths to a logical device already configured (with at least one path) in PowerPath and to managed storage system classes only. It does not apply to any storage system class that has been excluded from PowerPath control through the **powermt unmanage** command. The *EMC PowerPath Family CLI and Messages Guide* contains additional information on the **powermt unmanage** command.

This procedure can be done without interruption to running applications on AIX hosts. Once the reconfiguration is completed successfully, new native devices can be used like any other native devices.

⚠ CAUTION

All operations must succeed for the reconfiguration to be successful. If any step fails, resolve that issue before proceeding. Do not use the new configuration until the entire procedure completes successfully.

Path limit

Do not configure more than 32 paths per logical device.

Procedure

To add paths to a PowerPath logical device:

1. Run **powermt display** to confirm the current configuration, and run **powermt display dev=all** to confirm the configuration of the logical devices to which new paths will be added. Ensure that the number of logical devices, hardware paths, and I/O paths are as expected. The path state should be alive for known good paths and dead for known bad paths.

If there is a problem, correct it before proceeding.

2. Make physical path additions as required:
 - a. Map the logical device to additional storage-system ports.
 - b. Add new HBAs. Your vendor documentation will provide more details.
 - c. Attach cables.
 - d. Rezone Fibre Channel switches.
3. If using SAN Manager, Volume Logix, or Access Logix™, make new paths available to the host using those tools.
4. In order for the operating system to recognize new paths so that PowerPath can then recognize the new paths, run the **cfgmgr** .
[“emc_cfgmgr script” on page 68](#) provides more information on the emc_cfgmgr script.

5. Run **powermt config**.
6. Inspect the new PowerPath configuration:
 - a. Run **powermt display dev=all**.
 The new paths should be displayed with a state of alive.
 - b. Run **powermt restore** to test all paths.
 - c. Scan operating system error logs to ensure no errors are logged against the new paths.
7. Correct any issues detected.
8. Run **powermt save** to save the new configuration.

When configuring new paths, if there is a conflict in the **reserve_policy** attribute value for existing native paths and the **reserve_policy** attribute value for the newly configured native paths, an error message is logged.

Example:

```
May 21 11:00:58 libk034 user:warn|warning syslog: cfgpowerdisk::
  hdiskpower16: reserve_policy and PR_key_value attributes for
  hdisk116 and hdisk87 do not match ... (none,none)
  (single_path,no_reserve)
```

This must be corrected using the procedure below before the next reboot.

To change the attribute value:

1. Run **chdev -l hdiskpower# -a reserve_policy=<attr_value>** . This will reconfigure the device with the new value.

For example, if the attribute of the new native path is **single path** but the existing native path has **no_reserve** , then specify **no_reserve** in **<attr_value>**. This will configure the newly added paths with **no_reserve** attribute.

Note: PowerPath reconfiguration operations will fail if the affected devices are in use (example, VG is active on the PP device or I/Os are running).

If affected devices are in use:

1. If the reconfiguration fails because affected devices are in use, run the **chdev** command with the **-P** flag.
2. Reboot the host for the attribute change to come into effect.

Adding new logical devices

This procedure adds new logical devices (with one or more paths) to be managed by PowerPath and applies to managed storage system classes only. It does not apply to any storage system class that has been excluded from PowerPath control through the **powermt unmanage** command.

On AIX hosts, logical devices can be added without interruption of service, since no existing application can be using a logical device that is not yet available. However, the procedure to use a new logical device (regardless of whether it is managed by PowerPath) is platform-specific and may require service interruption.

Once the reconfiguration is completed successfully, new pseudo devices and native devices can be used.



All operations must succeed for the reconfiguration to be successful. If any step fails, resolve that issue before proceeding. Do not use the new configuration until the entire procedure completes successfully.

Procedure

To add logical devices to the PowerPath configuration:

1. Run **powermt display** to confirm the current configuration. Ensure that the number of logical devices, hardware paths, and I/O paths are as expected. The path state should be alive for known good paths and dead for known bad paths.
2. Make logical device and physical path changes as required:
 - a. Create new logical devices.
 - b. Map logical devices to one or more storage-system ports.
 - c. Add new HBAs. Your vendor documentation will provide more details.
 - d. Attach cables.
 - e. Rezone Fibre Channel switches.
3. If using SAN Manager, Volume Logix, or Access Logix, make new paths available to the host using those tools.
4. In order for the operating system to recognize a new logical device so that PowerPath can then recognize a new logical device, run the **emc_cfgmgr** script.

“[emc_cfgmgr script](#)” on page 68 provides more information on the `emc_cfgmgr` script.

5. Run **powermt config**.
6. Inspect the new PowerPath configuration:
 - a. Run **powermt display dev=all**.
All paths associated with the new logical devices should be displayed with a state of alive.
 - b. Run **powermt restore** to test all paths to the new logical device.
 - c. Scan operating system error logs to ensure no errors are logged against the new paths and logical device.
7. Correct any issues detected, before saving the PowerPath configuration or using the new logical device.
8. Set PowerPath-specific options for the new logical devices, such as load-balancing and failover policy, path modes, write throttle enablement, and priority.
9. Run **powermt config**.
10. Run **powermt save** to save the new configuration.
11. Where EMC ControlCenter is installed, run the command that refreshes the ControlCenter database of device information. The documentation for your version of EMC ControlCenter provides information.

Removing paths or logical devices

This section describes how to remove the following from a PowerPath configuration:
Specified paths to logical devices

Note: Before physically removing a device, ensure that the device is removed from PowerPath and AIX configurations. Otherwise, PowerPath may mark all paths dead for other devices and bring them alive again or data loss may occur.

Once a device is removed from the PowerPath configuration, you can remove it from the AIX configuration, and then replace hardware as needed. PowerPath maintains static information about the characteristics of an `hdiskpower` device in several places in the ODM and also within the `powerpath0` driver. This static information allows PowerPath to restore a failed path and to maintain persistent `hdiskpower` numbers across host reboots.

You should also follow this procedure to:

- ◆ Delete a PowerPath configuration, for example, as part of an ODM cleanup.
- ◆ Change logical unit numbers on storage devices.
- ◆ Restore a `mksysb` image from one host to another when the image contains preconfigured PowerPath devices.

CAUTION

Failure to follow this procedure could cause unexpected behavior when you later try to add devices to PowerPath.

Note: To reconfigure for PowerPath control any device you have removed from the PowerPath configuration, you must run `powermt config`.

Procedure

To remove paths or logical devices:

1. Run **powermt display dev=all** to:
 - Confirm the configuration of the logical device(s) from which paths will be removed. Check the number of existing paths. The path state should be **alive** for known good paths and **dead** for known bad paths. If there is a problem, correct it before proceeding.
 - Identify the PowerPath HBA number associated with the paths to be removed. In complex topologies, there can be multiple paths on an HBA.
2. As necessary, identify the physical paths to be removed or zoned out, and confirm that there are other paths to the affected logical devices. (Otherwise, applications using those logical devices could experience I/O errors when you proceed.)
3. Run **powermt remove** for each device you want to unmap. This updates the information in the powerpath0 driver. Specify on the command line:
 - The HBA—to remove the entire HBA.
 - The device—to remove all paths to the specified logical device.
 - Both HBA and device—to remove a single path to the specified logical device.
4. Run **rmdev -dl** for each PowerPath device or HBA you want removed from the host. This command deletes old entries and characteristics from the ODM.
5. Run **rmdev -dl** on the underlying hdisks associated with the devices you removed in the previous step.
6. Run **savebase -v** to update phase 1 device attributes in the boot logical volume.
7. Inspect the new PowerPath configuration.
 - Run **powermt display**. The output should show fewer total paths than before. All paths should have a state of optimal.
 - Run **powermt display dev=all**. All remaining paths associated with the affected logical devices should be displayed with a state of alive.

Correct any issues detected above before saving the PowerPath configuration or using the new logical devices.
8. Run **powermt save** to save the new configuration.

Cleaning up incorrect path or logical volume removal

1. Run **powermt display** to identify the removed volume and its paths. In the output these should all display as **dead**.
2. Run **powermt check dev =<dev>** for device in question.

CHAPTER 5

PowerPath in a Cluster

This chapter includes the following sections:

- ◆ [PowerPath in a PowerHA/HACMP Cluster.....](#) 54
- ◆ [PowerPath in a Veritas Cluster Server cluster](#) 57

PowerPath in a PowerHA/HACMP Cluster

`emcpowerreset` is a binary that is required for any HACMP installation on AIX when PowerPath is also installed. You can download this binary from the EMC Online Support site.

Note: The upgrade procedure of PowerPath with PowerHA/HACMP is identical to the procedures described for other cluster environments.

Installing PowerPath and PowerHA/HACMP on new hosts

To install and configure PowerPath and PowerHA/HACMP when neither PowerPath nor HACMP is installed:

1. On all hosts, prepare the cluster hardware, making the necessary networking and disk connections between the hosts and the storage system. The relevant AIX HACMP documentation provides information. For Symmetrix or CLARiiON systems, the *EMC Host Connectivity Guide for IBM AIX* on EMC Online Support or the CLARiiON storage-system support website contains additional information.
2. On one host in the cluster:
 - a. Install PowerPath. [Chapter 1, “Installing PowerPath on an AIX host”](#) contains additional information.
 - b. Identify the LVM volume groups that use PowerPath hdiskpower devices:
 1. First run `powermt display dev=all` to list each hdiskpower device.
 2. Then run `lspv` to identify the volume group to which each hdiskpower device belongs.
 - c. Install PowerHA/HACMP following the instructions in the relevant AIX PowerHA/HACMP documentation. Configure PowerHA/HACMP to use the volume groups identified in step 2b. Make sure they are the same volume groups on each host.

When using PowerPath in a concurrent resource group environment, the `/usr/sbin/cluster/diag/clconraid.dat` file must be updated to recognize `power` devices:

- For any storage system, you can update `clconraid.dat` manually by inserting the word `power` in the first line of the file using a text editor.
- d. Stop all applications that use the volume groups identified in step 2b.

Unmount all file systems that use these volume groups.
Run `smit varyoffvg` to deactivate each volume group.
3. On each remaining host in the cluster:
 - a. Install PowerPath. [Chapter 1, “Installing PowerPath on an AIX host”](#) contains additional information.

If any hdisk attached to the host does not have a PVID or has a different PVID on different hosts, run `rmdev` on that hdisk. Then run the `/usr/lpp/EMC/Symmetrix/bin/emc_cfgmgr` script, followed by `powermt config`, to configure the devices for the host.

The `emc_cfgmgr` script is available at the following EMC FTP site:

ftp://ftp.EMC.com/pub/elab/aix/ODM_DEFINITIONS

Do not define any volume groups. Instead, you will import the volume groups from the host on which you installed PowerPath.

- b. Use the `smit importvg` command to import each volume group identified in step 2b.
 - c. Use the `smit chvg` command to change the auto-activation status of each volume group that you imported in step 3b from yes to no.
 - d. Install PowerHA/HACMP, following the instructions in the relevant AIX HACMP documentation. Configure HACMP to use the volume groups imported in step 3b.
4. On all hosts, start cluster services, using the `smit clstart` command.

The volume groups and the underlying PowerPath `hdiskpower` devices are now under the control of the HACMP software.

5. Implement the `emcpowerreset` binary. `emcpowerreset` is required for any HACMP installation on AIX 5.x when PowerPath is installed. You can download the `emcpowerreset` binary from <http://support.emc.com> (search for PowerReset). EMC Knowledgebase Solutions `emc69100` and `emc104555` contain more information.

When using PowerPath with HACMP in a CLARiiON environment, “[cfgscsi_id binary in a CLARiiON environment](#)” on page 56 provides information.

Integrating PowerHA/HACMP in a PowerPath environment

When PowerPath, but not HACMP, is installed on the hosts to be part of the cluster, follow these steps:

On each host in the cluster, one host at a time:

1. Prepare the cluster hardware, making the necessary networking and disk connections among the hosts and the storage system. The relevant AIX HACMP documentation provides information. For Symmetrix or CLARiiON systems, the *EMC Host Connectivity Guide for IBM AIX* on EMC Online Support or the CLARiiON storage-system support website contains additional information.
2. Identify the LVM volume groups that use PowerPath `hdiskpower` devices:
 - a. First run `powermt display dev=all` to list each `hdiskpower` device.
 - b. Then run `lspv` to identify the volume group to which each `hdiskpower` device belongs.
3. Install PowerHA/HACMP, following the instructions in the relevant AIX HACMP documentation. Configure HACMP to use the volume groups identified in step 2. Ensure they are the same volume groups on each host.

When using PowerPath in a concurrent resource group environment, the `/usr/sbin/cluster/diag/clconraid.dat` file must be updated to recognize `power` devices:

- For any storage system, you can update `clconraid.dat` manually by inserting the word `power` in the first line of the file using a text editor.

4. Stop all applications that use the volume groups identified in step 2. Unmount all file systems that use these volume groups. Use the `smit varyoffvg` command to deactivate each volume group.
5. Start cluster services, using the `smit clstart` command.

The volume groups and the underlying PowerPath hdiskpower devices are now under the control of the HACMP software.
6. Implement the `emcpowerreset` binary. `emcpowerreset` is required for any HACMP installation on AIX 5.x when PowerPath is installed. You can download the `emcpowerreset` binary from <http://support.emc.com> (search for PowerReset). The EMC Knowledgebase Solutions emc69100 and emc104555 provide more information.

cfgscsi_id binary in a CLARiiON environment

In a CLARiiON environment the `cfgscsi_id` binary is required in addition to the `emcpowerreset` binary. When you install PowerPath, the `cfgscsi_id` binary is unpacked to the `/usr/sbin` directory.

Note: The synchronization for the cluster will fail if `cfgscsi_id` is not used.

Synchronizing volume group information between cluster nodes

If a shared volume group must be extended, use the following procedure:

1. Run the `/usr/lpp/EMC/Symmetrix/bin/emc_cfgmgr` script to configure any new devices to the active host.

The `emc_cfgmgr` script is available at the following EMC FTP site:

ftp://ftp.EMC.com/pub/elab/aix/ODM_DEFINITIONS

2. Run `powermt config`.
3. Run `chdev -l hdiskpowerx -a pv=yes` to create a PVID on the new devices.
4. Configure the new devices to all other nodes in the cluster, following steps 1 and 2.

Once this is complete, you can extend your volume group on the active node.

5. Import the shared volume group to the standby node, using the HACMP CSPOC utilities. On the standby node, use the `smit fastpath smitty cl_vg` command.
6. Select **Import a Shared Volume Group**.
7. Select the volume group that was just extended on the active node.
8. Select the node and physical hdiskpower device on which the volume group was extended.

The node's volume group is updated with the new extended volume.

hdiskpower devices in a PowerHA/HACMP environment

As a safety feature, PowerPath does not put a PVID on an hdiskpower entry unless it can read the information from the disk itself. If the disk is reserved (varied ON) by another AIX node, no other AIX node can access the disk. This includes PowerPath commands run on the other AIX nodes to try to read the PVID information from the disk.

hdiskpower devices may lose PVID and VG information after becoming Defined in an AIX shared disk environment. If this happens, an HACMP cluster does not fail over correctly. Use the following online procedure to correct the situation:

1. Run `lspv` to confirm that all PVID and VG information is correct on all the underlying hdisk entries. (If it is not, EMC Knowledgebase Solution emc74494 provides additional information.)
2. Run `rmdev -l hdiskpowerx` to put all the hdiskpower entries back into a Defined state.
3. Run `lsdev -Cc disk` to confirm all hdiskpower entries are Defined.
4. Move (fail over) all the HACMP resource groups to this AIX node. With the hdiskpower entries in a Defined state, an HACMP failover will now work because the correct break SCSI reservation utility will be called.
5. Run `powermt config` to move the PVID and VG information from the hdisk to the hdiskpower entries and make them **Available**. Once all the HACMP resources are owned by this node, PowerPath commands can access the disk.

Note: If you reboot a host when the hdiskpower device is Available and the PVID is on the underlying hdisk, `lspv` output loses all PVID and VG information. If a failover is attempted at this point, it will fail. [“Troubleshooting” on page 78](#) provides information to help you solve this problem.

EMC Knowledgebase Solution emc74494 provides more information.

PowerPath in a Veritas Cluster Server cluster

This section describes how to:

- ◆ Install PowerPath and VCS in a new cluster, that is, where neither PowerPath nor VCS software is installed on any host to be included in the cluster.
- ◆ Integrate PowerPath into an existing VCS cluster.

Installing PowerPath in a new VCS cluster

To install PowerPath and VCS when neither PowerPath nor VCS is installed on any host:

1. On each host to be included in the cluster:
 - a. Prepare the cluster hardware, making the necessary networking and disk connections among the hosts and the storage system. Refer to the relevant VCS documentation.
 - b. Verify that all storage system devices are seen by each host.

- c. Install PowerPath as described in [Chapter 1, “Installing PowerPath on an AIX host.”](#) Verify that PowerPath can see all the devices.
 - d. Install any applications.
 - e. Install the VCS software, following the installation procedure described in the relevant Veritas Cluster Server documentation. Initially configure VCS to run without a service group.
2. On each node in the cluster:
 - a. Define the resources (for example, VxVM volumes) that make up the service group. (You will configure the service group in step 3.)
 - b. Identify the LVM volume groups that use PowerPath hdiskpower devices.
 - First, run `powermt display dev=all` to list each hdiskpower device.
 - Then, run `lspv` to identify the volume group to which each hdiskpower device belongs.
 3. On one host in the cluster:
 - a. Configure the service group by adding the resources you defined in step 2 to the `/etc/VRTSvcs/conf/config/main.cf` file. The disk or logical device resources should use PowerPath pseudo (hdiskpower) devices.

Note: Using a disk for service group heartbeat instead of a network is subject to restrictions. Not all disks can be used. Consult the Veritas documentation.

 - b. Configure VCS to use the volume groups identified in step 2b. Make sure they are the same volume groups on each host.
 - c. Start cluster services on the host.
 4. Start cluster services on each remaining node in the cluster. These hosts rebuild their local configuration files from the `main.cf` file you edited in step 3a.
 5. On each node in the cluster:
 - a. Verify that the service group is up and running, and use either the VCS GUI (`hagui`) or the `hagrp -list` command to verify that the service group can successfully fail over to all hosts in the cluster.
 - b. Add other service groups as needed.

Integrating PowerPath into a VCS cluster

To integrate PowerPath into an existing VCS cluster, follow these steps on each cluster node, one node at a time:

1. Run `hastop -local -evacuate` to stop cluster services on the node.
2. Install or upgrade PowerPath. [Chapter 1, “Installing PowerPath on an AIX host”](#) contains additional information.
3. Run `hastart` to start cluster services on the node.

Wait for each node to be fully integrated into the cluster before running `hastart` on the next node.

Changing the naming scheme for Veritas volumes

Use this procedure to change the namingscheme for Veritas volumes when PowerPath is installed.

1. To display the attributes of a specified enclosure like enclosure type, enclosure serial number, status, array type, and number of LUNs, run:

```
# vxdmpadm listenclosure all
```

Output similar to the following is displayed:

```
ENCLR_NAME ENCLR_TYPE ENCLR_SNO STATUS ARRAY_TYPE LUN_COUNT
=====
pp_emc_clariion0 PP_EMC_CLARiion APM00120202000 CONNECTED A/A 44
pp_emc0 PP_EMC 000195700879 CONNECTED A/A 45
ibm_vscsi0 IBM_VSCSI VSCSI CONNECTED VSCSI 2
```

2. Change the device naming scheme for TPD-controlled enclosures, run:

```
# vxdmpadm setattr enclosure enclosure_name tpdmode=native
```

Example: # vxdmpadm setattr enclosure pp_emc0 tpdmode=native

3. Select the operating system-based naming by running the command:

```
# vxddladm set namingscheme=osn
```

4. To display the current disk-naming scheme and the mode of operations, run the command:

```
# vxddladm get namingscheme
```

Output similar to the following is displayed:

```
NAMING_SCHEME      PERSISTENCE      LOWERCASE      USE_AVID
=====
OS Native          No                Yes            Yes
-bash-3.00# vxdisk list
hdiskpower124 auto:none - - online invalid
hdiskpower125 auto:none - - online invalid
hdiskpower126 auto:none - - online invalid
hdiskpower127 auto:none - - online invalid
hdiskpower128 auto:LVM - - LVM
hdisk1 auto:LVM - - LVM
hdisk6 auto:LVM - - LVM
```

5. Scan the disks by running the command:

```
# vxdisk scandisks
```

6. Display information on all the disks that are known to VxVM, by running:

```
# vxddladm set namingscheme=ebn
```

Output similar to the following is displayed:

```
# vxdisk list
DEVICE          TYPE          DISK          GROUP          STATUS
pp_emc_clariion0_0 auto:LVM      -             -             LVM
pp_emc_clariion0_1 auto:none     -             -             online
invalid
pp_emc_clariion0_2 auto:none     -             -             online
invalid
```

```

pp_emc_clariion0_3 auto:none      -      -      online
invalid
pp_emc_clariion0_4 auto:LVM      -      -      LVM
pp_emc_clariion0_5 auto:LVM      -      -      LVM
pp_emc_clariion0_6 auto:LVM      -      -      LVM
pp_emc0_0      auto:none      -      -      online invalid
pp_emc0_1      auto:cdsdisk    -      -      online invalid
pp_emc0_2      auto:LVM      -      -      LVM
pp_emc0_3      auto:none      -      -      online invalid
pp_emc0_4      auto:none      -      -      online invalid
pp_emc0_5      auto:none      -      -      online invalid
pp_emc0_6      auto:LVM      -      -      LVM
pp_emc0_7      auto:none      -      -      online invalid

```

CHAPTER 6

Removing PowerPath from an AIX Host

This chapter includes the following sections:

- ◆ Before removing PowerPath 62
- ◆ Removing PowerPath..... 62
- ◆ After removing PowerPath 64
- ◆ When a storage system device is the boot device 64

Before removing PowerPath

Before you remove PowerPath from the host:

- ❑ Disable VxVM volumes on pseudo devices. Ensure no pseudo devices are in use and unmount any mounted file systems contained on pseudo devices.
 1. Unmount any mounted filesystem on pseudo devices.
 2. Close VxVM volumes, run:


```
Vxvol -g <DG> stopall
```
 3. Deport DG, run:


```
Vxdg deport <DG>
```
 4. Remove pseudo devices from VxVM control, run:


```
Vxdisk rm <hdiskpowerx>
```
- ❑ Deactivate all LVM volume groups that use PowerPath hdiskpower devices. Otherwise, the uninstall procedure will fail.
- ❑ Close any application that is using an hdiskpower device.

If an application is using an hdiskpower device when you try to remove PowerPath, the uninstall procedure terminates with the following message:

```
Error removing PowerPath devices. Unable to remove PowerPath.
```

If this error occurs, close the application that is using the hdiskpower device and repeat the uninstall.
- ❑ If EMC Solutions Enabler daemons, such as **storwatchd** and **storapid**, are running, type **stordaemon shutdown <daemon>** to stop the daemons before you uninstall PowerPath.
- ❑ If EMC ControlCenter is running on the host, stop the ControlCenter agents before you uninstall PowerPath.
- ❑ In a CLARiiON environment, if the Navisphere Host Agent is running, type **/etc/rc.agent stop** to stop the agent before uninstalling PowerPath.
- ❑ If you are *not* reinstalling PowerPath after completing the removal procedure, disconnect all duplicate physical connections between the host and the storage system except one cable, leaving a single path. In addition, reconfigure any switches so devices appear only once.

Removing PowerPath

To remove PowerPath, you can use either command line entries or the SMIT utility.

Note: Do not attempt to run `powermt config` while using `installp` or SMIT to add, remove, or update PowerPath software. Doing so can result in a system crash.

Using command line entries

To remove PowerPath by using command line entries:

1. Log in as root.

2. Remove the PowerPath software, by running:

installp -u EMCpower

Output similar to the following appears:

```
+-----+
|                                     Summaries:                                     |
+-----+

Installation Summary
-----
Name                                Level             Part              Event             Result
-----
EMCpower.MgmtComponent             6.x.x.0          USR               DEINSTALL        SUCCESS
EMCpower.migration_enabler         6.x.x.0          USR               DEINSTALL        SUCCESS
EMCpower.mpx                        6.x.x.0          USR               DEINSTALL        SUCCESS
EMCpower.base                      6.x.x.0          USR               DEINSTALL        SUCCESS
```

Where 6.x.x is the PowerPath for AIX release version; for example, PowerPath 6.0.0. PowerPath is now removed from the host. “[After removing PowerPath](#)” on page 64 contains additional information.

Note: The **installp -up EMCpower** command removes all hdiskpower devices and the powerpath0 driver, but leaves the product installed.

Using SMIT

Note: The SMIT procedure described in this section assumes you run the X Windows version of SMIT. You can use the **tty** version of SMIT if you substitute the appropriate **tty SMIT** commands in the following steps.

To remove PowerPath by using SMIT:

1. Log in as root.
2. Run **smit** to open SMIT.
3. Select **Software Installation and Maintenance**, and then select **Software Maintenance and Utilities**.
4. Select **Remove Installed Software**.
 The **Remove Installed Software** dialog box opens.
5. Click **list** to open the **Multi-select** list of installed software.
6. Select entries starting with **EMCpower**, and then click **OK**.
7. Change the **Preview Only?** field to **No**, and click **OK**.
8. When prompted, confirm you want to remove the software.

Output similar to the following appears:

```
+-----+
|                                     Summaries:                                     |
+-----+

Installation Summary
-----
```

Name	Level	Part	Event	Result
EMCpower.MgmtComponent	6.x.x.0	USR	DEINSTALL	SUCCESS
EMCpower.migration_enabler	6.x.x.0	USR	DEINSTALL	SUCCESS
EMCpower.mpx	6.x.x.0	USR	DEINSTALL	SUCCESS
EMCpower.base	6.x.x.0	USR	DEINSTALL	SUCCESS

Where 6.x.x is the PowerPath for AIX release version; for example, PowerPath 6.0.0.

9. From the **EXIT** menu, select **EXIT SMIT**.

PowerPath is now removed from the host. [“After removing PowerPath” on page 64](#), contains additional information.

Note: SMIT preview of uninstall removes all hdiskpower devices and the powerpath0 driver but leaves the product installed.

After removing PowerPath

If you are not reinstalling PowerPath after completing the removal, run `rm /etc/emcp_registration` to remove the `/etc/emcp_registration` file from your system.

When a storage system device is the boot device

If your PowerPath installation uses a storage system device as the boot device, follow these steps to remove PowerPath:

1. Run `pprootdev off` to disable multipathing to the root device.
2. Restart the host.
3. Remove PowerPath as described in [“Removing PowerPath” on page 62](#).
4. Run `bosboot -ad /dev/ipldevice` to create a boot image.

CHAPTER 7

PowerPath Administration on AIX

This chapter includes the following sections:

◆ Management daemon using SNMP.....	66
◆ emc_cfgmgr script.....	68
◆ PowerPath hdiskpower devices	69
◆ PowerPath and the lsvg command.....	71
◆ PowerPath and the iostat command	72
◆ BCVs and cfgmgr operations	72
◆ Bringing hdiskpower-based BCV Symmetrix logical devices online	72
◆ Importing an LVM volume group from a remote host.....	73
◆ Changing the target/LUN address of a storage system logical device	73
◆ Adding new devices to an existing configuration	74
◆ Managing and unmanaging EMC Celerra iSCSI devices	74
◆ Manage IBM XIV devices under PowerPath	74
◆ Replacing an HBA that PowerPath is using online	75
◆ Moving LUNs in and out of a storage group.....	75
◆ Eliminating ghost devices	75
◆ Reconfiguring PowerPath devices online	76
◆ Failover in switched environments	76
◆ SMIT screens.....	77
◆ Audit and error messages	77
◆ Troubleshooting	78

Management daemon using SNMP

The **PowerPath management daemon** monitors specific PowerPath events and sends an SNMP (Simple Network Management Protocol) trap when access to devices is disrupted. The configuration file for the event monitoring daemon specifies the events to monitor and a filter parameter for each event. This way, only those events needing intervention from an administrator are displayed.

An SNMP trap is issued for events listed in the configuration file, `/etc/emc/emcp_mon.d.conf`. The traps are sent to a master SNMP agent running locally. The master agent forwards the trap to a remote network manager where an administrator can view them and take corrective action as necessary.

Monitored events

The PowerPath management daemon monitors and sends an SNMP trap when either of the following events occur:

Note: Traps are only generated for Path Dead and All Paths Dead events when there is I/O running.

- ◆ A path is dead for one minute.
- ◆ All paths to a volume are dead for five minutes.

When a monitored event occurs on a path, a timer is started. After a given delay time elapses, a trap is generated to report the event. If a cancelling event occurs on the path before the delay time elapses, a trap is not generated. In this way the delay time is used to filter out transient path events.

[Table 1 on page 66](#) shows the events, event descriptions, and event IDs and time thresholds. It also shows whether for each case an SNMP trap is issued.

Table 1 Management daemon monitored events

Monitored event	Event cause	Event ID	SNMP trap
Path is Dead	Path is dead for one minute.	0x2,0x4,60	Yes
	Path is dead and then recovers within one minute.	0x2,0x4,60	No
All Paths Dead	All paths are dead for five minutes.	0x20,0x40,300	Yes
	All paths are dead and then recover within five minutes.	0x20,0x40,300	No

The time threshold for each event can be modified as described in [“Editing the event” on page 68](#).

The configuration file for the event monitoring daemon contains comment lines, beginning with a hash mark, and event lines, containing the ID of each event to monitor, the ID of its cancelling event, and the delay time in seconds for the event. These three event parameters are delimited by commas. For example, the line for the path dead event is `0x2,0x4,60`.

Configuring the management daemon

After PowerPath is installed on the host being managed, configure the event monitoring daemon to send traps to the SNMP manager. Configure each host on which path monitoring is desired.

To configure `emcp_mond_edit` related settings

1. Run `/etc/emc/bin/emcp_mond_edit`

Item	Value
====	=====
1 Change remote management service status	Enabled
2 Change management service heartbeat interval	10 minutes
3 Change management service port	9083
4 Single path gatekeeper devices	Allow
5 Change SNMP trap generation status	Enabled
6 Change "Path Dead" event	Enabled, delay 0 sec
7 Change "Path Alive" event	Enabled, delay 0 sec
8 Change "All Paths Dead" event	Enabled, delay 0 sec
9 Change "Not All Paths Dead" event	Enabled, delay 0 sec
10 Change "Volume Dead" event	Enabled, delay 0 sec
11 Change "Max. Path Latency Exceeded" event	Enabled, delay 0 sec
12 Write configuration file	
13 Exit	

2. Select 12 and 13 from this list.

To configure management daemon to send the traps to SNMP TrapReceiver

1. Edit the IP address to include the destination IP where the trap should be sent.

- a. Edit the `/etc/snmpd.conf` file:

```
< trap public <Trap_Destination_IP> <Unique_Object_Identifier> fe
```

where:

- `<Trap_Destination_IP>` is the IP address of the host where the SNMP manager is installed.
- `<Unique_Object_Identifier>` is 1.3.6.1.4.1.1139.12.1. This is the Object identifier for EMC error log trap.
- `< trap public <IP> 1.3.6.1.4.1.1139.12.1 fe`

- b. Edit the `/etc/snmpdv3.conf` `TARGET_ADDRESS` to match the destination IP address:

```
<TARGET_ADDRESS TargetX UDP <Trap_Destination_IP> traptag
trapparms1---
```

where `TargetX` is the next Target number available.

- c. Modify the `/etc/snmpdv3.conf` file to see the trap notifications in the trap receiver.

To receive SNMP traps

1. Edit the `/etc/snmpdv3.conf` file.
2. Uncomment `# VACM_VIEW defaultView internet - included` to ensure the TrapReceiver receives Trap notifications.

```
Example: <TARGET_ADDRESS Target1 UDP < IP> traptag
trapparms1---
```

3. Restart the SNMP daemon (snmpd), and then verify the status:

- a. Run **stopsrc -s snmpd**.
- b. Run **startsrc -s snmpd**.
- c. Run **lssrc -s snmpd**.

Output in the following format is displayed:

```
Subsystem Group PID Status
snmpd tcpip 32258 active
```

2. Start the **emcp_mond** daemon, run:

```
/usr/sbin/emcp_mond
```

Editing the event

The time threshold before an alert is generated can be modified by editing the `emcp_mond.conf` file.

To edit the time threshold:

1. Open the `/etc/emc/emcp_mond.conf` file.
2. To change the time threshold, for example, from 60 to 90 seconds for the **Path is dead** event, edit the following line:

```
0x2,0x4,60
```

```
to:
```

```
0x2,0x4,90
```

3. Close the `emcp_mond.conf` file.
4. Restart the daemon by running:
 - a. Run **stopsrc -s snmpd**.
 - b. Run **startsrc -s snmpd**.
 - c. Run **/usr/sbin/emcp_mond**.

emc_cfgmgr script

PowerPath requires that an hdisk be configured for each logical path it will use to access a storage system logical device. However under certain circumstances, AIX does not configure an hdisk for each logical path to a storage system logical device.

For example, you attach four new SCSI cables to an AIX host. Each cable addresses the same four storage system logical devices, and each of those devices at one time was part of a volume group and is configured with a PVID (which is written on the disk). You then reboot the host. When AIX boots, it does the device discovery on those new SCSI busses in one step. When it sees two or more devices with the *same* PVID, AIX creates only one hdisk. As a result, there are only 4 new hdisks, even though you attached 16 new devices.

⚠ WARNING

Do not share a PowerPath pseudo device used for paging with another host or cluster node. The `hdiskpower` device must be visible to only one host or cluster node.

To ensure hdisks are configured correctly for PowerPath, PowerPath for AIX provides the script `/usr/lpp/EMC/Symmetrix/bin/emc_cfgmgr`. The `emc_cfgmgr` script invokes the AIX `cfgmgr` tool to probe each HBA separately, so the configuration program restarts before it gets confused by disks that appear to be the same. After `emc_cfgmgr` runs, a storage system hdisk is configured for each device on each path.

The `emc_cfgmgr` script is available at the following EMC FTP site:

ftp://ftp.EMC.com/pub/elab/aix/ODM_DEFINITIONS

Changing the major number

To change the major number:

1. Determine a free major number common to all nodes. On each node in the cluster, type:

```
lvlstmajor
```

Compare the free ranges on each node to find a major number that is unused on any node.

2. Run `rmdev -l powerpath0`
3. Change the major number of the `powerpath0` device. On each node in the cluster, type:

```
chdev -l powerpath0 -a pp_major=x
```

where *x* is the major number you found in step 1.

4. Verify the change. On each node in the cluster, type:

```
lsattr -El powerpath0
```

The major number will take effect after you run the `powermt config` command on each node in the cluster.

When all the `hdiskpower` devices are removed and when the `powerpath0` driver is in the **available** state, `pp_major` is not taken into account when configuring new `hdiskpower` devices. However, if the `powerpath0` driver is in the **defined** state or if a `hdiskpower` device already exists, the `pp_major` attribute is correctly taken into account when configuring a `hdiskpower` device.

PowerPath hdiskpower devices

PowerPath for AIX supports only **pseudo devices**, not native devices. To get the benefits of PowerPath load balancing and path failover in an AIX environment, **you must use PowerPath pseudo devices.**

PowerPath pseudo devices on AIX use this naming convention:

- ◆ Block device — `/dev/hdiskpowerx`

- ◆ Raw device — `/dev/rhdiskpowerx`

where `x` is the disk number.

During installation, PowerPath creates an `hdiskpower` device for every logical device configured for the AIX host. After PowerPath is installed, both `hdisk` and `hdiskpower` devices exist on the host. The `hdiskpower` devices reside on top of the `hdisk` devices.

You can run `powermt display dev=all` to determine the correspondence between PowerPath `hdiskpower` devices and AIX `hdisk` devices.

Once PowerPath is installed, applications should direct I/O to `hdiskpower` devices. Using `hdiskpower` devices provides the PowerPath load-balancing and path failover functionality. PowerPath then selects the best path (`hdisk`) to handle the I/O.

During installation, PowerPath migrates existing LVM volume groups that use storage system `hdisks` to PowerPath `hdiskpower` devices. You need not reconfigure existing volume groups after installing PowerPath. If you have an application that accesses AIX `hdisks` directly rather than through a volume group (a DBMS, for example), you must reconfigure that application to use PowerPath `hdiskpower` devices if you want PowerPath load-balancing and path failover functionality.

If an application does not access an `hdisk` directly, you need not reconfigure the application for PowerPath.

When defining new volume groups, use PowerPath `hdiskpower` devices, not AIX `hdisk` devices. If you add an application to your system that typically would access `hdisks` directly, configure the application to use `hdiskpower` devices instead.

Although the underlying `hdisk` devices remain after PowerPath is installed, EMC recommends you not use them for normal I/O because they might interfere with one another. It might not be possible to open `hdisk` devices if the parent `hdiskpower` device is open. Device reservations on the `hdisk` can interfere with device reservations on `hdiskpower` devices. Applications that use the SymmAPI™ cannot use both `hdisk` and `hdiskpower` devices.

PVIDs

A physical volume identifier (PVID) is a unique number written on the first block of the device. The AIX LVM uses this number to identify specific disks. When a volume group is created, the member devices of the group are simply a list of PVIDs. The LVM does not read each device when searching for member devices of a volume group; instead, it expects the PVIDs to be saved in the ODM, and it uses the ODM attribute when determining which device to open.

The PVID for each device is stored in the ODM when the device is configured. When a device is made Available (including device creation and when the device begins in the Defined state), the configuration program tries to read the first block of the device. If it succeeds and the first block contains a valid PVID, the PVID value is saved as an attribute in the ODM for that device. Once the PVID is set in the ODM, it can be seen in the output of the `lspv` command. In a configuration with multiple paths to the same logical devices, multiple `hdisks` show the same PVID in the output of `lspv`. When the LVM needs to open a device, it selects the first `hdisk` in the list with the matching PVID.

hdiskpower and PVID

The PVID for an hdiskpower device is set essentially the same way as an hdisk, but with an extra step or two. When an hdiskpower device is made Available, the configuration program tries to open the device and read the first block. Several conditions can prevent this read from succeeding, including these:

- ◆ There is a SCSI reservation on the device. This usually is caused by an active volume group using one of the hdisk paths on the local machine or varied on from a remote host.
- ◆ hdisk paths to the hdiskpower are marked dead because of a deleted hdisk device. This can prevent the configuration program from opening the device and reading the first block.

These failure conditions happen primarily when PowerPath is being configured long after system boot, and other programs are using hdisk devices on the local machine.

If the hdiskpower configuration program cannot read the first block on the device, it cannot determine the PVID and cannot store it in the ODM for the hdiskpower device.

When the configuration program for the hdiskpower device reads and stores the PVID for the hdiskpower device, it also removes the PVID from the ODM for the corresponding hdisk devices. This is done so the LVM will use the hdiskpower devices instead of the hdisks and take advantage of PowerPath functionality.

When configuring PowerPath devices, keep in mind that:

- ◆ Deleting all hdiskpower devices does not erase the PowerPath knowledge of which hdisks correspond to paths to logical devices. To cause PowerPath to completely rebuild its configuration, you must unconfigure the powerpath0 device.
- ◆ hdisks need not be deleted to make them redo their PVID processing. They can be unconfigured by running `rmdev -l hdisk#`, and reconfigured by rerunning `cfgmgr` on the bus or running `mkdev -l hdisk#`.
- ◆ To have PVIDs on hdiskpower devices, you need only put the hdisks into the Available state. You do not need to delete them, and you do not need to first get the PVID to appear in `lspv` output. You do, however, need to ensure the associated path hdisks are not in use and the device is not reserved.

PowerPath and the lsvg command

The AIX `lsvg` command, when used with the `-p` flag, displays devices in use by the specified volume group. This command, however, is not designed to operate with PowerPath or with storage system logical devices that are addressable as different hdisk devices. In general, the output of `lsvg -p vgroupname` shows correct information, but several administrative tasks change the ODM and could cause `lsvg` to show misleading information. These tasks include:

- ◆ Use of the `pprootdev` tool. This tool changes the ODM and is intended to be used when you expect to reboot the system soon after using `pprootdev`. The `lsvg` command shows misleading device information when run after `pprootdev`. This is not an indication that something is wrong. A reboot corrects the `lsvg` output, but reboot is not required.

- ◆ Use of `cfgmgr` to create new hdisk devices after PowerPath is already configured. Always run `powermt config` after adding new devices to include them in the PowerPath configuration.

PowerPath and the `iostat` command

The `iostat -a` command output may show most I/O going **through a single HBA**, when in fact I/O is balanced across all HBAs. This section explains why.

PowerPath assigns the parent device of a pseudo device based on the last native device retrieved from the ODM database. Thus, if the last native device for every pseudo device happens to be on the same fcs device, all pseudo devices will have that device as parent. For example, if the last native device for every pseudo device is on fcs0, all pseudo devices will have fcs0 as parent.

The `iostat -a` command combines data for adapters based on the ODM parent attribute for each disk. Thus, in this example, all I/O to all hdiskpower devices will show up on the fcs0 parent. It will appear that fcs0 is getting most of the activity, when in fact multipathing is occurring below the pseudo device.

BCVs and `cfgmgr` operations

If you run the `cfgmgr` command, hdisks that correspond to business continuance volumes (BCVs) will be changed to the `Defined` state. If those hdisks are current members of the PowerPath configuration, PowerPath may mark them `Failed` during periodic testing. After the hdisk paths are made `Available`, run `powermt restore`. This causes PowerPath to test the hdisks and change their state to `Alive`.

Note: EMC suggests that you run `cfgmgr` only at nonpeak hours when the effects of performance issues are minimized.

Bringing hdiskpower-based BCV Symmetrix logical devices online

Note: If the hdisks for the BCVs already exist, do not delete them. Do not specify the `-d` flag in the `rmdev` command. It is appropriate to delete and re-create the hdisk devices only when the device mapping changes so the Symmetrix serial number for the hdisk is different than when the hdisks were last in the `Available` state. [“Changing the target/LUN address of a storage system logical device” on page 73](#) provides additional information.

To bring hdiskpower-based BCV Symmetrix logical devices online:

1. Use the EMC management tool of your choice to split the BCV (and make it ready).
2. Use `/usr/lpp/EMC/Symmetrix/bin/mkbcv` to bring the BCV hdisks to the `Available` state.
3. Run `powermt config`.
4. Run `powermt restore`.

If errors are reported, PowerPath's configuration was changed. Verify that all paths are functioning, and run `powermt check` to remove all dead hdisks. Rerun `powermt config`. You should now be able to run `powermt restore` without errors.

5. Verify that expected PVIDs are assigned to `hdiskpower` devices in `lspv` output.

If they are not, ensure corresponding hdisks are not in use or reserved (locally or remotely). Then, unconfigure the corresponding `hdiskpower` devices (`rmdev -l hdiskpower#`) and reconfigure them (`mkdev -l hdiskpower#`). If the expected PVID is not set, the device cannot be accessed due to path failures or a conflict on the device.

Importing an LVM volume group from a remote host

To import an LVM volume group:

1. On the remote host, vary off the volume group to be imported.
2. Verify that all path hdisks for the volume group to be imported are configured on the local host. If they are not, configure them with the `/usr/lpp/EMC/Symmetrix/bin/emc_cfgmgr` script. The `emc_cfgmgr` script is available at the following EMC FTP site:

ftp://ftp.EMC.com/pub/elab/aix/ODM_DEFINITION

3. If the `hdiskpower` devices for the volume group already exist, unconfigure them with `rmdev -l hdiskpower#`.
4. Verify that the devices are not in use on any host.
5. Run `powermt config` to reconfigure or create the `hdiskpower` devices.
6. Verify that expected PVIDs are assigned to `hdiskpower` devices in `lspv` output.

If they are not, ensure that corresponding hdisks are not in use or reserved (locally or remotely). If the expected PVID is not set, the device cannot be accessed due to path failures or a conflict on the device.

Changing the target/LUN address of a storage system logical device

If you need to change device mapping on the storage system, where a target or LUN address will be changed to access a different storage system logical device with a different serial number, that logical device must be deleted from the PowerPath configuration before the change is made. Otherwise, PowerPath will begin marking `hdisk` paths **dead** because they no longer point to the storage system logical device where they should point.

To change the target/LUN address:

1. Run `powermt display` to determine which `hdiskpower` devices need to be removed.
2. Run `powermt remove dev=<device>` for each `hdiskpower` device corresponding to a storage system logical device whose address is changing.

3. Change the storage system addressing configuration.
4. Ensure that all required hdisks are configured for the new storage system configuration and are in the Available state.
5. Run `powermt config`.

Adding new devices to an existing configuration

To add devices to an existing configuration:

1. Verify that all hdisks for new storage system volumes are configured.
2. Verify that all **BCV hdisks** are in the **Available** state.

PowerPath will not configure hdisks in the **Defined** state. If you are adding BCV hdisk devices to the PowerPath configuration, they must be split and **Available** before proceeding.

3. Run `powermt config`.

Managing and unmanaging EMC Celerra iSCSI devices

PowerPath treats EMC Celerra iSCSI devices as generic devices. Enable generic Loadable Array Module (LAM) support so that PowerPath can recognize and manage and unmanage Celerra iSCSI devices.

Manage EMC Celerra iSCSI devices under PowerPath

1. Add the following line to the `/etc/emc/mpaa.lams` file.

```
managed:generic:EMC:Celerra
```

2. Run `powermt init`.
3. Reboot the host.

Unmanage EMC Celerra iSCSI devices under PowerPath

1. Remove the following line to the `/etc/emc/mpaa.lams` file.

```
managed:generic:EMC:Celerra
```

2. Run `powermt init`.
3. Reboot the host.

Manage IBM XIV devices under PowerPath

To manage IBM XIV devices under PowerPath:

1. Install the EMC TPA ODM kit.

```
bash-4.0# ls1pp -l | grep TPA
EMC.AIX.TPA.ODM 1.3.0.2 COMMITTED EMC TPA ODM kit
```

2. Present IBM XIV devices under the host (by default the devices will come under MPIIO control).

```
bash-4.0# lsdev -Cc disk | grep XIV
hdisk47      Available 00-08-02 MPIIO 2810 XIV Disk
hdisk48      Available 00-08-02 MPIIO 2810 XIV Disk
```

3. Remove devices from MPIIO control.

```
rmdev -Rdl hdisk47
rmdev -Rdl hdsik48
```

4. Run the manage command to manage IBM XIV under PowerPath.

```
# manage_disk_drivers -d 2810XIV -o EMCPP_XIV
```

5. Reboot the host.

Replacing an HBA that PowerPath is using online

To replace an HBA online:

1. Run `powermt display` to determine the HBA number.
2. Run `powermt remove hba=<hba#>`

Where `<hba#>` is the adapter number identified in step 1.

3. Use `rmdev` to remove all hdisks associated with the HBA that you removed in step 2, as well as the parent HBA.

Moving LUNs in and out of a storage group

In a CLARiiON environment, if you remove LUNs from a storage group, and then return them to the storage group with **different HLU numbers, `powermt restore` fails,** complaining that native devices are dead. This is because the LUNs now have different SCSI addresses and are treated as new devices by the OS.

To move LUNs in and out of a storage group:

1. Run `/usr/lpp/EMC/Symmetrix/bin/emc_cfgmgr` to configure new native devices.
2. Run `powermt config` to configure new PowerPath pseudo devices.
3. Run `powermt check force` to remove the dead paths.

Eliminating ghost devices

Ghost devices are created when the **system continues to recognize hardware** or a service after it has been disconnected from the system. If you reboot the host when a cable is detached, you must reconfigure the devices using both `emc_cfgmgr` and `powermt config`.

To eliminate ghost devices:

1. Reconnect the cable.

2. Run `/usr/lpp/EMC/Symmetrix/bin/emc_cfgmgr`.
3. Run `powermt config`.
4. Run `powermt restore`.

Reconfiguring PowerPath devices online

Whenever the physical configuration of the storage system or the host changes, you must reconfigure the PowerPath devices to reflect the new configuration. Configuration changes that require you to reconfigure PowerPath devices include the following:

- ◆ 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

To reconfigure PowerPath devices:

1. Make sure all physical device connections are connected.
2. Run the `/usr/lpp/EMC/Symmetrix/bin/emc_cfgmgr` script to ensure hdisks are configured for each path. This script invokes the AIX `cfgmgr` tool to probe each adapter bus separately. After it runs, there should be a storage system hdisk configured for each device on each path.

The `emc_cfgmgr` script is available at the following EMC FTP site:

ftp://ftp.EMC.com/pub/elab/aix/ODM_DEFINITIONS

3. Run `powermt restore` to test all configured paths.
4. Run `powermt check` to remove any hdiskpower devices that are no longer available.
5. Run `powermt config` to configure new devices and paths that were added to the system configuration.
6. Optionally, run `powermt save` to save the new PowerPath configuration.

Failover in switched environments

PowerPath includes a configurable control, `QueueDepthAdj`, that can improve path failover performance in certain limited situations.

`QueueDepthAdj` affects failovers only:

- ◆ In switched environments.
- ◆ When the failover is due to an outage between the switch and the storage system.

Extensive testing indicates that the default setting is best for virtually all installations. If you experience slow failover in a switched environment, contact EMC Customer Support for information on setting `QueueDepthAdj`.

SMIT screens

PowerPath for AIX provides a set of System Management Interface Tools (SMIT) screens that implement `powermt` functionality. Using a SMIT screen relieves you of the burden of having to know PowerPath command syntax.

To access the PowerPath for AIX SMIT screens:

1. Type `smit`.
2. Press **Enter**.
3. Select **Devices > PowerPath Disk**.

The **PowerPath Disk** SMIT screen opens. Select the desired option.

Audit and error messages

PowerPath reports any errors, diagnostic messages, and failover recovery messages through the `syslog` file that is specified by the administrator (for example, `/usr/safe.log`).

The *EMC PowerPath Family CLI and System Messages Reference* contains a complete list of PowerPath error messages.

Error log messages

Error log messages appear along with audit messages in the log file when you configure the AIX operating system for common logging as described in [“Enabling logging on an AIX host” on page 77](#). The error-log messages capture unexpected events that occur. Some error-log messages convey information that appears onscreen when a command fails.

PowerPath provides error notification through the AIX `errlog/errpt` facility. The `powermt` utility reports errors to standard error (`stderr`).

The AIX documentation contains information on the AIX `errlog/errpt` facility.

Enabling logging on an AIX host

If logging is not enabled on the AIX host, you need to enable it in order to log PowerPath messages.

Note: EMC recommends using `/tmp/emcpsyslog.log` to ensure the support utility EMC `grab` can collect the log files.

To log messages in `/tmp/emcpsyslog.log`:

1. Add the following entry to the `/etc/syslog.conf` file:

```
*.info /tmp/emcpsyslog.log
```

Optionally, configure `/etc/syslog.conf` to rotate the log file. For example, to rotate the file daily, keep one week's worth of files, and compress files to save space, add the following entry:

```
*.info /tmp/emcpsyslog.log rotate files 7 time 1d compress
```

2. Create the `/tmp/emcpsyslog.log` file. For example:

```
touch /tmp/emcpsyslog.log
```

3. Enable logging by running:

```
refresh -s syslogd
```

Troubleshooting

This section describes problems you might encounter and suggests how to resolve them.

Problem

You see the following error message:

```
A device is already configured at this location
```

Cause

You cannot configure a defined hdisk if it has the same connection string (in `lsdev` output) as the corresponding hdiskpower and the hdiskpower device is in the Available state.

Solution

Run `mkdev -l hdiskpower#` for the corresponding hdiskpower device. This either changes the connection string for the hdiskpower device or unconfigures the hdiskpower device to allow the hdisk to be configured. If the condition exists for multiple hdisks, you can run `powermt config` instead.

Problem

hdisk paths are marked as failed.

Cause

If you delete an hdisk (running `rmdev -dl hdisk#`) before removing it from PowerPath's configuration, PowerPath marks the hdisk paths as failed because it can no longer access the hdisk it expects to find.

In some cases, an hdisk is present, but it points to the wrong storage system logical device.

Solution

1. Run `powermt restore` to test and mark dead all paths that are missing or point to the wrong logical device.
2. Run `powermt check`. When prompted to remove a dead path, respond with **a** to remove all dead paths.
3. Run `powermt config` to configure all hdisks that might be pointing to storage system logical devices different from the devices PowerPath is aware of.

Problem

`powermt display dev=all` shows all paths as dead or unknown.

Cause

Deleting and remaking hdisk devices while the powerpath0 device is in the Available state can put PowerPath in a state where it has incorrect path information for hdiskpower devices. The `powermt restore` command cannot restore these paths, because they no longer refer to the correct storage system logical device.

Solution

1. Run `powermt restore`.
2. Run `powermt check`. When prompted to remove a dead path, respond with a to remove all dead paths.
3. Run `powermt config`.
4. Verify that an hdisk is configured for each connection and device. [“Before you install” on page 12](#) contains more information. If an hdisk is not configured, complete the procedure to correct the hdisk configuration, and then run `powermt config` again.

Problem

`lspv` output loses all Physical Volume ID (PVID) and Volume Group (VG) information. This will cause a failover attempt to fail.

Cause

If you reboot a host when the hdiskpower device is Available and the PVID is on the underlying hdisk, `lspv` output loses all PVID and VG information. If a failover is attempted at this point, it will fail. Before you proceed, complete the procedure that follows.

Solution

1. Run `rmdev -l hdiskpowerx` to put all the hdiskpower entries back into a Defined state.
2. Run `rmdev -dl hdiskall_ghost_X` to remove all the new “ghost” devices created during the reboot.
3. Run `mkdev -l hdiskall_original_X` to make all the original hdisk entries, now Defined, Available again.
4. Start the HACMP Cluster Services on this node.
5. Once the HACMP Cluster Services are up, run the `lspv` command. Confirm the PVID and VG information is back on the original hdisks, and then follow the procedure in the preceding section.

Problem

I/O for a particular application appears to be failing.

Cause

If I/O for a particular application appears to be failing, it is possible that the correct flags were not set. The application would recognize it specifically as pass-through I/O, though the user may not be able to distinguish this particular I/O type.

Solution

Ensure that the `SC_SIMPLE_Q` flag is set for applications that use pass-through SCSI commands with devices handling I/O. Such applications must set the `SC_SIMPLE_Q` flag to indicate command tag queuing. If this flag is not set, the pass-through SCSI commands could fail. The user application is responsible for handling this condition.

APPENDIX A

Files Changed by PowerPath Installation

This appendix includes the following sections:

- ◆ Files added by installation 82
- ◆ Files modified by installation 83

Files added by installation

The following files are added when PowerPath is installed on AIX:

```

/usr/lib/methods/power.cat
/usr/lib/methods/powerpath
/usr/lib/methods/powerpath/cfgudid
/usr/lib/methods/powerpath/cfgmpx
/usr/lib/methods/powerpath/cfgdm
/usr/lib/methods/cfgpower
/usr/lib/methods/chgpowerdisk
/usr/lib/methods/ucfgpower
/usr/lib/methods/cfgpowerdisk
/usr/lib/methods/ucfgpowerdisk
/usr/lib/methods/undpowerdisk
/usr/lib/methods/undpower
/usr/lib/nls/msg/en_US/power.cat
/usr/lib/drivers/powerdd
/usr/lib/drivers/powerdiskd
/usr/lib/drivers/mpxext
/usr/lib/drivers/dmext
/usr/lib/drivers/gpxext
/usr/lib/boot/protoext/disk.proto.ext.scsi.pseudo.power
/usr/lib/libemcp_sockcom.a
/usr/lib/libemcp_xml.a
/usr/lib/libemcp_pp_util.a
/usr/lib/libemcp_pp_util.so
/usr/lib/libemcp_mpapi.a
/usr/lib/libemcp_mpapi_rtl.so
/usr/lib/libcg.so
/usr/lib/libcg.a
/usr/lib/libemcp.so
/usr/lib/libemcp.a
/usr/lib/libemcp_core.so
/usr/lib/libemcp_core.a
/usr/lib/libemcp_lam.so
/usr/lib/libemcp_lam.a
/usr/lib/libemcp_mp_rtl.so
/usr/lib/libemcp_mp.a
/usr/lib/libmp.a
/usr/lib/libemcp_power.a
/usr/lib/libemcp_shlib.a
/usr/lib/libemcp_lic_rtl.so
/usr/lib/libpn.so
/usr/lib/libpn.a
/usr/lib/libpower.a
/usr/sbin/cfgscsi_id
/usr/sbin/powercf
/usr/sbin/powermt
/usr/sbin/emcp_mond
/usr/sbin/emcpmgr
/usr/sbin/emcpreg
/usr/sbin/emcpminor
/usr/sbin/emcpupgrade
/usr/sbin/emcpadm
/usr/sbin/pprootdev
/usr/sbin/emcpassive2active
/usr/sbin/powermig
/usr/sbin/emcpmigd
/usr/share/man/man1/emcpminor.1
/usr/share/man/man1/emcpreg.1
/usr/share/man/man1/emcpadm.1
/usr/share/man/man1/powermt.1
/usr/share/man/man1/powermig.1
/usr/lpp/EMCpower/liblpp.a

```

```
/etc/rc.emcpower  
/etc/emc/diag/powermt_display_dead  
/etc/emc/bin/emcp_mond_edit  
/etc/emc/emcp_mond.conf  
/etc/opt/emcpower/server.pem  
/etc/rc.emcp_mond  
/etc/rc.powermig
```

Note: The location of the `/etc/disk.proto.ext.scsi.pseudo.power` file is part of the list of files added as part of the installation process, by default. However, after the boot of a SAN setup, the location of the file is

`/usr/lib/boot/protoext/disk.proto.ext.scsi.pseudo.power.`

Files modified by installation

When you install PowerPath on AIX, the PowerPath template for error logging is updated. In addition, the following files are modified as follows:

- ◆ `/etc/trcfmt` is updated with the PowerPath trace format file.
- ◆ `/etc/inittab` is modified with entries added for `/etc/rc.emcpower` and `/etc/rc.powermig`.

INDEX

A

- adding devices to existing configuration 74
- adding logical devices 50
- adding paths, while booting off the array 45

B

- BCV devices, hdiskpower-based 72
- boot device
 - configuring 38
 - disabling 45
- booting off the array, adding paths while 45
- bosboot tool 39

C

- CLARiiON storage systems
 - booting on 38
 - setting up 12
 - trespass 39
- cluster. *See also* HACMP cluster, EMC Legato AAM cluster, VCS cluster
- configuring, boot device 38

D

- device mapping, changing 73
- device naming conventions 69
- devices, adding online 76
- disabling multipathing 30
- DMP (Dynamic Multipathing) 36
- documentation, related 7

E

- ECC. *See* EMC ControlCenter
- EMC ControlCenter 51
- emc_cfgmgr script 69
- emcp_mond.conf 66
- emcpassive2active command 39
- error messages 77

G

- ghost devices, removing 43

H

- HACMP cluster
 - failover problems 57
 - integrating with PowerPath 55
- HBA, replacing 75
- hdisk devices, configuring for PowerPath 69
- hdisk, configuring, troubleshooting 78
- hdisk, deleting, troubleshooting 78
- hdiskpower devices 70
 - initializing 22

I

- installing PowerPath
 - command line 18
 - download from EMC Online Support 19
 - mksysb master image, as part of a 16
 - postinstallation instructions 21
 - preinstallation instructions 12
 - service packs 20
 - SMIT (System Maintenance Interface Tool) procedure 19
- iostat command output 72

L

- licensing PowerPath 21
- logical devices, adding 50
- logical devices, removing from PowerPath
 - configuration 51
- LPP (Licensed Program Products), required 13
- lsvg command 71

M

- management daemon
 - configuring 67
 - definition 66
 - restarting 68

- mksysb master image
 - installing PowerPath as part of 16
- monitored events, SNMP 66
- multipathing, disabling 30

N

- naming conventions 69
- nimadm command 31

P

- patches. *See* service packs
- physical volume identifier (PVID) mismatch 44
- powermig transition -startup 25
- PowerPath Configuration Checker (PPCC) 25
- PowerPath devices
 - hdiskpower 69
 - reconfiguring online 76
- PowerPath error messages 77
- pprootdev tool 30, 39, 45
- pseudo devices 69
- PVID (physical volume identifier) 70

R

- reconfiguring PowerPath devices online 76
- registering PowerPath 21
- removing PowerPath 62
- removing PowerPath boot device 64
- rootvg mirror, creating 44

S

- service packs, installing 20
- SMIT (System Maintenance Interface Tool)
 - accessing screens 77
 - installing by using 19
- SNMP management daemon 66
- Symmetrix storage arrays 38

T

- tar file, installing 19
- time threshold, editing 68
- troubleshooting 78
 - removing ghost devices 43
 - trespass occurrence 39
 - upgrade, PowerPath 29

U

- uninstalling PowerPath 62
- upgrading PowerPath 25

V

- varyonvg vgname command 23
- VCS cluster
 - installing with PowerPath 57
 - integrating PowerPath into 58
- volume groups, LVM
 - importing 73
 - varying on 23
- VxVM (Veritas Volume Manager) and hdiskpower devices 36