

EMC® PowerPath® for AIX
Version 5.3

Installation and Administration Guide
P/N 300-008-341
REV A08

EMC Corporation
Corporate Headquarters:
Hopkinton, MA 01748-9103
1-508-435-1000
www.EMC.com

Copyright © 1997 - 2011 EMC Corporation. All rights reserved.

Published November, 2011

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.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com.

All other trademarks used herein are the property of their respective owners.

Preface

Chapter 1

Installing PowerPath on an AIX host

Before you install	14
Obtain up-to-date information	14
Locate your license key	15
Prepare the host and storage system	15
Installation procedure	19
Mounting the installation CD	19
Installing the tar file	22
Installing PowerPath as part of a mksysb master image	23
After you install	25
Register PowerPath on the host.....	25
Install a PowerPath service pack or patch	26
Choose a major number	27
Initialize PowerPath hdiskpower devices.....	27
Remove the CD	28
Configure the management daemon	28
Vary on volume groups	28
Reconfigure applications	29
Upgrading PowerPath	30
Before upgrading PowerPath.....	30
During an upgrade	30
After an upgrade.....	30
Upgrade procedure	31
Upgrading PowerPath in a VIO server environment.....	32
Troubleshooting the upgrade	34
If a PowerPath device is the boot device.....	35
Upgrading the AIX operating system.....	36

	Coexistence with third-party path management software	37
Chapter 2	Installing PowerPath and Veritas Volume Manager	
	Installing with Veritas Volume Manager on new hosts	40
	Converting to PowerPath from DMP	41
	Adding a new path.....	41
Chapter 3	Configuring a PowerPath Boot Device on AIX	
	Setting up a PowerPath boot device	44
	Considerations for Symmetrix storage.....	45
	Considerations for VNX OE and CLARiiON storage	45
	The pprootdev tool	46
	Syntax	46
	PowerPath and the bosboot tool.....	47
	VNX OE and CLARiiON trespasses	47
	Configuring a new PowerPath installation.....	48
	Configuring an existing PowerPath installation.....	50
	Updating the system microcode.....	50
	Installing PowerPath.....	51
	Cloning the boot device.....	51
	Troubleshooting	53
	Adding paths.....	55
	Disabling PowerPath on a storage system boot device	56
Chapter 4	Configuring a PowerPath Logical Device	
	Introduction	58
	Adding new paths to a logical device	59
	Adding new logical devices	60
	Removing paths or logical devices.....	63
Chapter 5	PowerPath in a Cluster	
	PowerPath in an HACMP Cluster	66
	Installing PowerPath and HACMP on new hosts	66
	Upgrading PowerPath in an HACMP Cluster.....	68
	Integrating HACMP in a PowerPath environment.....	69
	cfgscsi_id binary in a VNX OE and CLARiiON environment.	70
	Synchronizing volume group information between	
	cluster nodes.....	72
	hdiskpower devices in an HACMP environment	73

PowerPath in an EMC Legato AAM cluster	75
Installing PowerPath in a New AAM Cluster	75
Integrating PowerPath into an AAM cluster	76
PowerPath in a Veritas Cluster Server cluster	77
Installing PowerPath in a new VCS cluster.....	77
Integrating PowerPath into a VCS cluster.....	78

Chapter 6 Removing PowerPath from an AIX Host

Before Removing PowerPath.....	80
Removing PowerPath.....	81
Using command line entries.....	81
Using SMIT	82
After removing PowerPath.....	83
When a storage system device is the boot device.....	83

Chapter 7 PowerPath Administration on AIX

Management daemon using SNMP	86
Monitored events	86
Configuring the management daemon.....	87
Editing the event	88
emc_cfgmgr script.....	90
PowerPath hdiskpower devices.....	90
PVIDs	92
hdiskpower and PVID.....	92
PowerPath and the lsvg command.....	94
PowerPath and the iostat command	95
BCVs and cfgmgr operations.....	95
Bringing hdiskpower-based BCV Symmetrix logical devices online.....	95
Importing an LVM volume group from a remote host.....	96
Changing the target/LUN address of a storage system logical device.....	97
Adding new devices to an existing configuration.....	97
Replacing an HBA that PowerPath is using online.....	98
Moving LUNs in and out of a storage group	98
Eliminating ghost devices	98
Reconfiguring PowerPath devices online.....	99
Removing paths or logical devices from a PowerPath configuration.....	100
Failover in switched environments	102
SMIT screens	102

Audit and error messages.....	103
Audit log messages	103
Error log messages	104
Troubleshooting	106

Appendix A Files Changed by PowerPath Installation

Files added by installation.....	110
Files modified by installation.....	114
ODM modifications caused by PowerPath installation.....	114

Index

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 Product Guide*
- ◆ *EMC PowerPath Family CLI and System Messages Reference Guide*
- ◆ *EMC PowerPath for AIX Installation and Administration Guide*
- ◆ *EMC PowerPath for AIX Release Notes*

These documents are updated periodically. Electronic versions of the updated manuals are available on the Powerlink® website:

<http://Powerlink.EMC.com>

The following table provides a list of reference documents for specific storage environments. All documents are available on the Powerlink website.

If your storage environment includes	Refer to
Symmetrix storage system	<i>EMC Host Connectivity Guide for IBM AIX</i>
VNX OE and CLARiiON storage system	• <i>EMC Host Connectivity Guide for Linux</i> • <i>CLARiiON Storage-System Support website</i> (www.emc.com/clariionsupport) • <i>VNX OE support website</i> (www.emc.com/vnxsupport)
Invista system	• <i>EMC Invista Installation and Setup Guide</i> • <i>EMC Element Manager Administration Guide</i>
Other vendors' storage system	The appropriate documentation from your vendor.

Another related EMC enterprise storage publication is the E-Lab Interoperability Navigator, available on the Powerlink website.

A subset PowerPath functions are available through the Unisphere™ application for VNX OE systems. Refer to the VNX OE Storage System Support website (www.emc.com/vnxsupport).

A subset of PowerPath functions are available through the Navisphere® and the Unisphere applications for CLARiiON systems. Refer to CLARiiON Storage-System Support website (www.emc.com/clariionsupport).

Revision history The following table presents the revision history of this document

Revision	Date	Description
A08	November 16, 2011	As of this release, this document now includes a revision table. Added <i>patch</i> to the service pack installation instructions since full package patch installations on PowerPath for AIX are the same as service pack installations. See “Install a PowerPath service pack or patch” on page 26 .

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.

IMPORTANT

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

Typographical conventions

EMC uses the following type style conventions in this guide:

Normal font

In running text:

- 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	• 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>	• Book titles • New terms in text • Emphasis in text
Courier	• Prompts • System output • Filenames • Pathnames • URLs • Syntax when shown in command line or other examples
Courier bold	• User entry • Options in command-line syntax
<i>Courier italic</i>	• 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 for information about EMC products, licensing, and service, go to the EMC Powerlink website (registration required) at:

<http://Powerlink.EMC.com>

Technical support — For technical support, go to EMC WebSupport on Powerlink. To open a case on EMC WebSupport, you must be a WebSupport customer. Information about your site configuration and the circumstances under which the problem occurred is required.

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:

SSGDocumentation@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.

Installing PowerPath on an AIX host

This chapter includes the following sections:

- ◆ Before you install..... 14
- ◆ Installing PowerPath as part of a mksysb master image 23
- ◆ Installation procedure 19
- ◆ After you install..... 25
- ◆ Upgrading PowerPath..... 30
- ◆ Upgrading the AIX operating system..... 36
- ◆ Coexistence with third-party path management software 37

Before you install

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

EMC Symmetrix, Hitachi Lightning, Hitachi TagmaStore, HP StorageWorks XP, HP StorageWorks EVA, and IBM ESS systems

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

EMC VNX OE and CLARiiON® storage systems

PowerPath installation is an integral part of an EMC VNX OE and CLARiiON setup and configuration procedure. The VNX OE Support home page (www.emc.com/vnxsupport), and EMC CLARiiON Support home page (www.emc.com/clariionsupport) provide detailed support information. There you will find step-by-step instructions for installing PowerPath and other CLARiiON software.

To obtain the latest versions of VNX OE and CLARiiON documentation, log in to the EMC Powerlink website, select **Support > Technical Documentation**.

Note: First-time PowerPath installations — Remove any version of EMC Navisphere® Application Transparent Failover (ATF) installed on the host.

The procedure for migrating from ATF or Common Desktop Environment (CDE) to PowerPath is not straightforward and could result in data loss if not performed correctly. We strongly recommend that EMC Professional Services perform the migration. If you nevertheless decide to perform the migration yourself, *Removing ATF or CDE Software Before Installing Other Failover Software* provides additional information. This document is available only on the EMC Powerlink website.

Obtain up-to-date information

Check the Powerlink website for current information:

- ☐ PowerPath documentation — EMC updates PowerPath documentation, including this installation guide, when it releases new features with a service pack or when documentation errors are reported. To obtain updated PowerPath documentation on Powerlink, select **Support > Technical Documentation**.

- ❑ Service packs and upgrades — You can download PowerPath service pack software from Powerlink. Determine which service packs (if any) to install after PowerPath, and whether those service packs have any additional installation prerequisites. To obtain service packs on Powerlink, select **Support > Downloads**.

Locate your license key

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

Note: If you are upgrading from an earlier version of PowerPath, you do not need to reregister. PowerPath will use your old key.

Prepare the host and storage system

- ❑ Go to Powerlink and verify your environment meets the requirements found in the *EMC PowerPath for AIX Release Notes* and the E-Lab™ Interoperability Navigator by running the PowerPath Configuration Checker (PPCC). For PPCC installation information, the *EMC PowerPath Configuration Checker User Guide*, available on Powerlink, provides additional information.
- ❑ All AIX hosts with EMC Invista® systems must select **Type 4** when registering an HBA with Invista 1.0 SP2 Patch 1 and later. **Type 4** is the supported HBA initiator type for AIX on Invista 1.0 SP2 Patch 1 and later. The *EMC Invista Element Manager Administration Guide* provides additional information on registering an HBA.
- ❑ Use the AIX **varyoffvg** command to take offline all AIX Logical Volume Manager (LVM) volume groups that use storage-system hdisk devices, except the root volume group (**rootvg**). If a file system or application uses these volume groups, unmount the file system or stop the application before taking the volume groups offline.
- ❑ Ensure any required AIX Licensed Program Products (LPPs) are installed.

Table 1 on page 16 lists and describes the required AIX LPPs.

Table 1 Required AIX LPPs

LPP	Description	Required for
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 driver kit	Fibre Channel devices
EMC.Symmetrix.fcp.rte EMC.CLARiiON.fcp.rte EMC.INVISTA.fcp.rte	EMC driver kit	Fibre Channel devices
EMC.Symmetrix.ha.rte EMC.CLARiiON.ha.rte	EMC driver kit Symmetrix or VNX OE and CLARiiON support for HACMP	HACMP clusters
Hitachi.aix.support.rte	Hitachi AIX support for xp disks	All Hitachi Lightning and Hitachi Tagmstore configurations
HP.aix.support.rte	HP AIX support for disks	All HP StorageWorks xp configurations
ibm2105.rte	IBM 2105 disk device	All IBM ESS configurations
PC1000.driver.obj	Cambex PC1000 or PC2000 HBA support	All HP StorageWorks EVA-hsv configurations
hsv100.rte or hsv110.rte	HP AIX support for disks	All HP StorageWorks EVA-hsv configurations

The Symmetrix, VNX OE, CLARiiON, and Invista LPPs are included in the EMC 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 105](#) provides more information.
- ❑ Ensure that the AIX hdisk devices are configured correctly. Each logical path PowerPath will use 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 and CLARiiON devices, type:

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

To remove hdisks corresponding to Hitachi Lightning and Hitachi TagmaStore devices, 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 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.

- ❑ Once all storage system hdisks are removed, run the **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. After the script runs, there should be a storage-system hdisk configured for each device on each path. You can run **lsdev -Cc disk** to confirm that this is the case.

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.

CAUTION

Be sure to follow HBA driver configuration guidelines outlined in the E-Lab Interoperability Navigator and product documentation. Using improper settings can cause erratic failover behavior, such as greatly increased I/O delays.

- ❑ With AIX version 5.2 ML-02 (APAR IY48488) or later, EMC recommends setting certain Object Data Manager (ODM) attributes to enable enhanced error recovery during certain failover scenarios.

Note: This does not apply to Cambex cards.

To configure the ODM attributes:

1. For Fibre Channel switch environments only. Enable the adapter driver **fast_fail** feature for each adapter to be managed by PowerPath:

```
chdev -l fscsi $x$  -a fc_err_recov=fast_fail -P
```

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

2. Restart the host.

The *EMC Host Connectivity Guide for IBM AIX*, available on Powerlink, provides more information on ODM attributes in general, and the **fast_fail** feature in particular.

- ❑ 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.

Installation procedure

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.

You can install PowerPath as a full package installation by using the installation CD, the package file, or the `mksysb` master image.

- ◆ [“Mounting the installation CD” on page 19](#)
- ◆ [“Installing the tar file” on page 22](#)
- ◆ [“Installing PowerPath as part of a `mksysb` master image” on page 23](#)

If you are installing a service pack of an already installed PowerPath release, follow the instructions in [“Install a PowerPath service pack or patch” on page 26](#).

If you are upgrading from an earlier release of PowerPath, follow the instructions in [“Upgrading PowerPath” on page 30](#) before you begin the installation.

Note: 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 [“Install a PowerPath service pack or patch” on page 26](#).

Mounting the installation CD

To mount the PowerPath installation CD:

1. Log in as root.
2. Create the directory `/cdrom` to be the mount point for the CD:


```
mkdir /cdrom
```
3. Insert the PowerPath installation CD into the CD-ROM drive.
4. Mount the CD on `/cdrom`:

```
mount -v cdrfs -p -r /dev/cd0 /cdrom
```

5. Change to the /cdrom/UNIX/AIX directory:

```
cd /cdrom/UNIX/AIX
```

6. Do one of the following:

- To install by using the command line, proceed to [step 2](#).
- To install by using SMIT, proceed to [step 2](#).

**Installing by using the
command line**

Note: You must enable the adapter driver fast_fail feature. To do this, type `chdev -l fscsiX -a fc_err_recov=fast_fail -P`, then reboot.

To install PowerPath by using command line entries:

1. Mount the CD in the CD-ROM drive and change to the /cdrom/UNIX/AIX directory as described in [“Mounting the installation CD” on page 19](#).
2. Install the software. Type:

```
installp -agXd . EMCpower
```

Note: The man page for the **installp** command describes optional flags for the command.

Output similar to the following appears:

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

Installation Summary

Name	Level	Part	Event	Result
EMCpower.base	5.3.0.0	USR	APPLY	SUCCESS
EMCpower.migration_enabler	5.3.0.0	USR	APPLY	SUCCESS
EMCpower.mpx	5.3.0.0	USR	APPLY	SUCCESS
EMCpower.consistency_grp	5.3.0.0	USR	APPLY	SUCCESS
EMCpower.encryption	5.3.0.0	USR	APPLY	SUCCESS

PowerPath is installed on the host, but before PowerPath can run, you must perform some administrative tasks (including registering PowerPath). [“After you install” on page 25](#) provides postinstallation information.

Installing by using SMIT

The SMIT (System Maintenance Interface Tool) procedure described in this section assumes you are running the X Window System version of SMIT. You can use the tty version of SMIT, provided you substitute the appropriate tty SMIT procedures.

Note: You must enable the adapter driver fast_fail feature. To do this, type `chdev -l fscsiX -a fc_err_recov=fast_fail -P`, and then reboot.

To install PowerPath by using SMIT:

1. Mount the CD in the CD-ROM drive and change to the /cdrom/UNIX/AIX directory as described in “Mounting the installation CD” on page 19.
2. Run **smit** to open SMIT.
3. On the main **SMIT** window, click **Software Installation and Maintenance**.
4. On the **Software Installation and Maintenance** window, click **Install and Update Software**.
5. On the **Install and Update Software** window, click **Install Software**.
6. For **INPUT device / directory for software**, type the following:
`/cdrom/UNIX/AIX`
7. Click **List** to open the list of software to install.
8. Select the first line in the list, **EMCpower ALL**, and press **Enter**.
The **Install Software** dialog box opens.
9. Review the installation options, make any necessary changes, and click **OK**.
10. When prompted, click **OK**.

Output similar to the following appears:

Summaries:				
Installation Summary				
Name	Level	Part	Event	Result
EMCpower.base	5.3.0.0	USR	APPLY	SUCCESS
EMCpower.migration_enabler	5.3.0.0	USR	APPLY	SUCCESS
EMCpower.mpx	5.3.0.0	USR	APPLY	SUCCESS

EMCpower.consistency_grp	5.3.0.0	USR	APPLY	SUCCESS
EMCpower.encryption	5.3.0.0	USR	APPLY	SUCCESS

11. Select **Exit SMIT** from the **Exit** menu to close SMIT.

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 25](#) provides postinstallation information.

Installing the tar file

To download from Powerlink and unzip the compressed tar file:

1. On Powerlink, select **Support > Downloads**.
2. Download the tar file to **/usr**.
3. Unzip the installation package. Type:


```
gunzip EMCpower.AIX.5.3.0.GA.tar.gz
```
4. Extract EMCpower_install. Type:


```
tar -xvf EMCpower.AIX.5.3.0.GA.tar.gz
```
5. Perform one of the following:
 - For an existing PowerPath installation, run **rm.toc** to delete the .toc file.
 - For new PowerPath installations, skip to [step 6](#).
6. Install the software. Type:


```
installp -agXd . EMCpower
```
7. Run **powermt config**.

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 25](#) provides postinstallation information.

Installing PowerPath as part of a mksysb master image

If you are creating a master site **mksysb** backup image for installing complete host images (including PowerPath software on target systems), you must completely remove all hdiskpower devices and the powerpath0 device from the source system from which the **mksysb** image will be created.

CAUTION

Failure to perform the following steps prior to creating the **mksysb** 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.

Before you begin

If the source host used to create the **mksysb** image is booted off 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 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 -Ccdisk -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 has been booted up.

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

After you install

After installing the PowerPath software:

- ◆ “Register PowerPath on the host” on page 25 (first-time installation only).
- ◆ “Install a PowerPath service pack or patch” on page 26.
- ◆ “Choose a major number” on page 27.
- ◆ “Initialize PowerPath hdiskpower devices” on page 27.
- ◆ “Remove the CD” on page 28.
- ◆ “Configure the management daemon” on page 28.
- ◆ “Vary on volume groups” on page 28.
- ◆ “Reconfigure applications” on page 29.

Register PowerPath on the host

Note: 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.

To register the PowerPath software:

1. Type the following:

```
emcpreg -install
```

Output similar to the following appears:

```
===== 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 appears:

```
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**. You can also add keys for other technologies, such as PowerPath Migration Enabler or PowerPath Encryption with RSA.

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

```
1 key(s) successfully added.
Key successfully installed.
Key (Enter if done):
```

If you type an invalid registration key, the screen displays 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 appears:

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

Install a PowerPath service pack or patch

You can install related Service Pack or Patch releases to add fixes and new features to an already installed PowerPath release. Usually a service pack or patch installation is a full package installation.

To install a PowerPath service pack or a patch:

1. On Powerlink, select **Support > Downloads**.
2. Download the tar file to **/usr**.
3. Unzip the installation package. For example, for a service pack, type:


```
gunzip EMCpower.AIX.5.3.SPx.bxxx.tar.gz
```
4. Extract EMCpower_install. For example, for a service pack, type:


```
tar -xvf EMCpower.AIX.5.3.SPx.bxxx.tar.gz
```
5. Perform one of the following:
 - For an existing PowerPath installation, run **rm.toc** to delete the .toc file.
 - For new PowerPath installations, skip to [step 6](#).

6. Install the software. Type:


```
installp -agXd . EMCpower
```
7. Run **powermt config**.

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 27](#), which follows, provides additional information.

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.

To initialize devices at the command line, type:

```
powermt config
```

To initialize devices by using SMIT:

1. Open SMIT by typing **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

To remove the CD:

1. Unmount the CD, by typing:

```
cd /umount /cdrom
```
2. Remove the CD from the CD-ROM drive.

Configure the management daemon

After the version of PowerPath that supports the management daemon is installed on the host, 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 87](#) 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. Run the command **varyonvg vgname**.
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 hdisks to PowerPath hdiskpower devices.

Note, however, that if you failed to vary off a volume group before installing PowerPath, this migration will fail. You will be able to vary off the volume group, but any 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 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. 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.

Upgrading PowerPath

You can upgrade PowerPath from versions 4.5.x or later to 5.3 SP1. The upgrading process requires knowledge of the PowerPath environment and its unique upgrading process.

Based on your environment and requirements, you can upgrade PowerPath either by using one of the following methods:

- ◆ Full package installation: You can directly upgrade PowerPath from a previous version with the later version of PowerPath package, or
- ◆ Uninstall the previous version and install the later version

Before upgrading PowerPath

- ❑ Download the latest version of EMC Grabs, available on Powerlink, and then run the PowerPath Configuration Checker (PPCC).
- ❑ Check the Powerlink website, for the most current release notes and service packs. Determine which service packs (if any) you want to install.
- ❑ You do not need to remove the previous version of PowerPath before upgrading to PowerPath 5.3.

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 done the following:
You have closed all applications that use PowerPath devices before you install PowerPath 5.3.
- ❑ 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

The environment in which PowerPath runs affects the upgrading process.

Note: When upgrading in a dual VIO server environment, follow the procedure described in [“Upgrading PowerPath in a VIO server environment.”](#)

To upgrade from PowerPath 4.5.x or later to PowerPath 5.3 SP1:

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 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, stop the daemons. Type:
`stordaeomon shutdown <daemon>`
- If EMC ControlCenter is running on the host, stop the ControlCenter agents.
- In a VNX OE and CLARiiON environment, if the Navisphere Host Agent is running, type:
`/etc/rc.agent stop`

2. (Optional) Run **powermt save** in PowerPath 4.x to save the changes made in the configuration file.
3. Install PowerPath as described in [“Installation procedure” on page 19](#).
4. Run **powermt config**.

5. (Optional) Run **powermt load** to load the previously saved configuration file.

When upgrading from PowerPath 4.x to PowerPath 5.3, an error message is displayed after running **powermt load**, due to differences in the PowerPath architecture. This is an expected result and the error message can be ignored.

Even if the command succeeds in updating the saved configuration, the following error message is displayed by running **powermt load**:

```
host1a 5300-08-01-0819:/ #powermt load Error loading auto-restore value
Warning:Error occurred loading saved driver state from file /etc/powermt.custom
...
Loading continues...
Error loading auto-restore value
```

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. The *EMC PowerPath Family CLI and System Messages Reference Guide* contains additional information on the **powermt** command.

Note that VNX OE devices are managed under the **clariion** storage class.

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 an earlier version and install PowerPath Version 5.3 SP1, you need to follow a different procedure:

1. Before uninstalling the earlier version of PowerPath, back up the mapping information first.
2. Unmap the devices.
3. Uninstall the existing PowerPath.

4. Install the later version of PowerPath.
5. Use the stored mapping information and map the devices.

Note: The VIOS Map Manager tool is not applicable when upgrading to PowerPath Version 5.3 SP1. However, you can use the tool when upgrading from PowerPath Version 5.3 SP1 to version 5.5 or later.

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 displays the mapping between physical, logical, and virtual devices.

Output similar to the following appears:

```
$ lsmap -all
SVSA
-----
vhost1

VTD          vtscsil
Status       Available
LUN          0x8100000000000000
Backing device hdiskpower5
Physloc      U789C.001.DQD0564-P1-C2-T1-L67
```

Physloc	Client Partition ID
U8203.E4A.10B9141-V1-C30	0x00000000

2. Log in on the same VIO server as the **padmin** user.
3. Unconfigure the PowerPath pseudo devices listed in step 1 by running:

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

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

For example, **rmdev -dev vtscsil -ucfg**

The VTD status changes to **Defined**.

Note: Run **rmdev -dev <VTD> -ucfg** for all VTDs displayed in step 1.

4. Upgrade PowerPath as described in [“Upgrade procedure” on page 31](#).
5. Run **powermt config**.
6. Log in as the **padmin** user and then configure the VTD unconfigured from step 3 by running:

```
cfgdev -dev <VTD>
```

Where <VTD> is the virtual target device.

For example, `cfgdev -dev vtscsil`

The VTD status changes to **Available**.

Note: Run `cfgdev -dev <VTD>` for all VTDs unconfigured in step 3

7. Run `lspath -h` on all clients to verify all paths are **Available**.
8. Perform steps 1 through 7 on the second VIO server.

Troubleshooting the upgrade

Upgrading from PowerPath 4.5.x or later could fail for the following reasons:

- ◆ Not all applications that use PowerPath devices were closed before starting the upgrade procedure.
- ◆ The PowerPath 4.5.x configuration files are corrupted.

If you failed to close applications

If you failed to close all applications that use PowerPath devices before starting the upgrade, you 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 start the upgrade procedure again.

If files are corrupted

If the PowerPath 4.5.x configuration files are corrupted, a message similar to the following appears during installation of PowerPath 5.3:

```
Upgrade fails and PowerPath installation continues
```

In this case, the installation continues:

- ◆ PowerPath 5.3 is installed on the host, but your custom configuration is no longer available.
- ◆ All PowerPath configuration parameters are set to their default values.
- ◆ You must reset any custom settings such as policies, priorities, and so on.

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.

To disable multipathing to the root device:

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

After upgrading, reenable multipathing to the root device:

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

[“The pprootdev tool” on page 46](#) contains more information about the **pprootdev** command.

Upgrading the AIX operating system

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.
3. Reinstall PowerPath, as described in [“Installation procedure” on page 19.](#)

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 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 37](#) 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 array type is disabled in PowerPath
HDLM	Hitachi Lightning, Hitachi TagmaStore
SDD	IBM ESS
HP StorageWorks Secure Path	HP StorageWorks EVA, EMA, MA
HP StorageWorks AutoPath XP Virtual Array	HP StorageWorks XP

Since PowerPath is not supported in third-party path management environments when the corresponding storage system class type is in the managed state, multipathing behavior for any attached arrays in the corresponding class is undefined if you subsequently run **powermt manage class=<class>**.

Note that VNX OE devices are managed under the **clariion** storage class.

[Table 3 on page 38](#) 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
HDLM	powermt manage class=hitachi
SDD	powermt manage class=ess
HP StorageWorks Secure Path	powermt manage class=hphsx
HP StorageWorks AutoPath XP Virtual Array	powermt manage class=hpxp

Since PowerPath and the third-party software cannot comanage devices, do not initiate comanagement by executing this command on third-party array class machines. As long as this command is 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>**.

Note that VNX OE devices are managed under the **clariion** storage class.

[Table 4 on page 38](#) 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
HDLM	powermt unmanage class=hitachi
SDD	powermt unmanage class=ess
HP StorageWorks Secure Path	powermt unmanage class=hphsx
HP StorageWorks AutoPath XP Virtual Array	powermt unmanage class=hpxp

Installing PowerPath and Veritas Volume Manager

This chapter includes the following sections:

Note: PowerPath 5.3 for AIX requires Veritas Volume Manager 4.0 or later.

- ◆ [Installing with Veritas Volume Manager on new hosts.....](#) 40
- ◆ [Converting to PowerPath from DMP.....](#) 41

Installing with Veritas Volume Manager on new hosts

When installing PowerPath 5.3 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.

Veritas Volume Manager versions 4.0 with MP1 or later do not require `powervxvm` to discover `hdiskpower` devices. You must run the `vxdisk list <DMP_name>` 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 <DMP_name>` output as follows:

```
lcla232 5300-04:/ #vxdisk list EMC0_5
Device:      EMC0_5
devicetag:   EMC0_5
type:        auto
info:        format=none
flags:       online ready private autoconfig invalid
pubpaths:    block=/dev/vx/dmp/EMC0_5
char=/dev/vx/rdmp/EMC0_5
guid:        -
udid:        EMC%5FSYMMETRIX%5F910018%5F180E8000
site:        -
Multipathing information:
numpaths:    1
hdiskpower7      state=enabled
```

Prior to MP1, the native `hdisks` were listed.

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** or **/usr/lpp/Symmetrix/bin/emc_cfgmgr**.
2. Run **powermt config**.
3. Run **vxdctl enable**.

Configuring a PowerPath Boot Device on AIX

This chapter includes the following sections:

- ◆ Setting up a PowerPath boot device 44
- ◆ The pprootdev tool 46
- ◆ PowerPath and the bosboot tool 47
- ◆ Configuring a new PowerPath installation 48
- ◆ Configuring an existing PowerPath installation 50
- ◆ Adding paths 55
- ◆ Disabling PowerPath on a storage system boot device 56

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 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 45](#) 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 OE and CLARiiON array, but such a configuration has some functional limitations and extra configuration steps.

The primary limitation to 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 OE or CLARiiON environment, the bootlist for the AIX host should contain 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 **lsdev** command. Passive hdisks will show a PVID of None. If PowerPath is configured, use the **powermt display dev=*n*** 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.

The pprootdev tool

The **pprootdev** tool:

- ◆ Enables and disables PowerPath load balancing and failover functionality on a storage system boot device.
- ◆ Enables the **bosboot** tool to function correctly with PowerPath.

Syntax

pprootdev has the following syntax:

```
pprootdev { on | off | fix }
```

on

Enables PowerPath pseudo boot. This command changes AIX configuration rules and updates the boot image so that AIX uses hdiskpower devices to vary on the rootvg in PHASE1 the next time the system boots.

off

Disables PowerPath pseudo boot.

fix

Enables bosboot to function correctly after booting from a PowerPath device (until the next system boot).

Note:

—The **pprootdev fix** and **pprootdev off** commands may display the following warning message:

```
0301-168 bosboot: The current boot logical volume,
/dev/hd5, does not exist on /dev/hdiskX.
```

You can ignore this message; it does not reflect an actual problem.

—You need invoke each of these **pprootdev** commands only *once*.

The **pprootdev** tool cannot change the state of rootvg on a running system. It does, however, modify ODM data that other tools use to determine what devices rootvg is using. For this reason, some commands report information that may appear to be incorrect if they are run after pprootdev and before a restart.

Note: If the /tmp directory is full and the disk containing the OS does not have free partitions available, the **pprootdev** command will fail. Ensure that there is enough available space before running the command.

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, as described in [“The pprootdev tool” on page 46](#), 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 a system restarts that uses PowerPath. **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 50](#) 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.

To configure a new PowerPath installation:

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 by a fibre adapter.
3. Install the current storage system drivers.
4. Restart the host.
5. [List and Delete](#) any hdisks in the **Defined** state. Type:

```
lsdev -Ccdisk -Sdefined -rname | xargs -n1 rmdev -dl
```
6. Install PowerPath as described in [Chapter 1, “Installing PowerPath on an AIX host.”](#)
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. [“Before you install” on page 14](#) contains additional information.
9. Run **powermt config**.
10. Run **pprootdev on** to set up multipathing to the root device.
[“The pprootdev tool” on page 46](#) contains additional information.

11. Run **bootlist -m** and include all paths to the boot device making sure that the paths to the active SP are included before the paths to the passive SP.

Note: If you are booting from a VNX OE and CLARiiON storage system, 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 45](#) contains more information.

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. Find the system model name, by running:

```
uname -M
```

Output similar to the following appears:

```
IBM, 7025-F80
```

2. Find the system microcode level, by running:

```
lscfg -vp | grep alterable
```

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, confirm the microcode level by running:

```
lscfg -vp | grep alterable
```

5. Confirm that the external disks are now bootable, by running:

```
bootinfo -B hdiskx
```

- 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 14](#).
3. Install PowerPath as described in [“Installation procedure” on page 19](#).
4. ~~Run `powermt config` to~~ ~~C~~configure the PowerPath driver, ~~by~~ ~~running:~~
~~`powermt config.`~~

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:

```
lsllpp -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 hev=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 appears:

```
Pseudo name=hdiskpower38
Symmetrix ID=000100006216
Logical device ID=006C
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
=====
```

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 -c disk -F name | xargs -n1 rmdev -l
```
8. ~~Move the powerpath0 device from Available to Defined state, by running:~~ Run **rmdev -l powerpath0** to remove the powerpath0 device.
~~**rmdev -l powerpath0**~~
9. Clear the PVID of the hdisk devices identified in step 3, by running:


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

Question: ***[ERIC -- note... Host Connectivity has “-apv” and not “-a pv”. Which is it? confirm]***

10. Create a copy of the operating system on the storage system hdisks identified in step 2:
 - For AIX 5.2:


```
alt_disk_install -C "hdisk132"
```

Question: Host connectivity guide uses only one command: “alt_install_disk -C hdisk_list”. Also is it “alt_disk_install” or “alt_install_disk”?

- For AIX 5.3:

```
alt_disk_copy -d "hdisk132"
```

Question: Same for other AIX releases supported on PP53 AIX?

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

Output similar to the following appears:

```
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
```

11. Restart the host, by running **shutdown -Fr**.
12. 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 OE and CLARiiON storage system, 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 45](#) provides more information.

13. Run **pprootdev on**.
[“The pprootdev tool” on page 46](#) contains more information.

14. Restart the host, by running **shutdown -Fr**.

When the system starts up, rootvg is using hdiskpower devices.

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**.

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.
[“The pprootdev tool” on page 46](#) contains more information.
2. Restart the host.

This chapter includes the following sections:

- ◆ Introduction 58
- ◆ Adding new paths to a logical device..... 59
- ◆ Adding new logical devices 60
- ◆ Removing paths or logical devices..... 63

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.



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 **emc_cfgmgr** script.
[“emc_cfgmgr script” on page 90](#) provides more information on the **emc_cfgmgr** script.
5. Reconfigure PowerPath.
6. Run **powermt config**.
7. 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.
8. Correct any issues detected.
9. Run **powermt save** to save the new configuration.

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. The *EMC PowerPath Family CLI and Messages Guide* provides additional information on 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.

⚠ 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.

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 90](#) 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.

Note that VNX OE devices are managed under the `clariion` storage class.

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
- ◆ Logical devices
- ◆ Entire HBAs

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 mkysb image from one host to another when the image contains preconfigured PowerPath devices.



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.

PowerPath in a Cluster

This chapter includes the following sections:

- ◆ PowerPath in an HACMP Cluster 66
- ◆ PowerPath in an EMC Legato AAM cluster 75
- ◆ PowerPath in a Veritas Cluster Server cluster 77

PowerPath in an 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 Powerlink website.

Installing PowerPath and HACMP on new hosts

To install and configure PowerPath and 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 VNX OE and CLARiiON systems, the *EMC Host Connectivity Guide for IBM AIX* or the VNX OE and 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.

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.

- c. Install HACMP, following the instructions in the relevant AIX HACMP documentation. Configure 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.

- For a Symmetrix or VNX OE and CLARiiON storage system, you can update the file automatically by running the `/usr/lpp/EMC/Symmetrix/bin/symcurrent` or `/usr/lpp/EMC/CLARiiON/bin/clarcurent` script.

The Readme file (`/usr/lpp/EMC/README.5100`) provides details.

- 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/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 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://Powerlink.EMC.com> (search for PowerReset). EMC Knowledgebase Solutions emc69100 and emc104555 contain more information.

When using PowerPath with HACMP in a VNX OE and CLARiiON environment, “[cfgscsi_id binary in a VNX OE and CLARiiON environment](#)” on page 70 provides information.

Upgrading PowerPath in an HACMP Cluster

You can upgrade PowerPath from Version 4.5.x or later to 5.3 SP1 in an HACMP cluster where two or more hosts run PowerPath.

1. Select the node where you want to upgrade PowerPath.
2. Fail over all the resource groups of the selected node to another node.
3. On the selected node, stop HACMP.
4. Stop all applications on the selected node.
5. Vary off all the volume groups on the selected node, except the root volume group (**rootvg**).
6. Ensure that the **naviagent** and **stordaemons** are stopped.
7. Using the **rmdev -Rdl fcsN** command, remove the SAN disk from all the host bus adapters.

In the command, *N* identifies the HBA number.

8. Install PowerPath 5.3 SP1 by using either **smit update_all** or **installp**.
9. Reboot the selected node where PowerPath 5.3 SP1 is installed.
10. Start HACMP on the selected node.
11. Failback the resources from the other node.
12. Repeat the same procedure for all other nodes where PowerPath needs to be upgraded.

Integrating 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 VNX OE and CLARiiON systems, the *EMC Host Connectivity Guide for IBM AIX* or the VNX OE and CLARiiON storage-system support website contains additional information.
2. 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. Install HACMP, following the instructions in the relevant AIX HACMP documentation. Configure HACMP to use the volume groups identified in step 2. 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.
- For a Symmetrix or VNX OE and CLARiiON storage system, you can update the file automatically by running the **/usr/lpp/EMC/Symmetrix/bin/symcurrent** or **/usr/lpp/EMC/CLARiiON/bin/clarcurent** script.

The readme file (**/usr/lpp/EMC/README.5100**) provides details.

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://Powerlink.EMC.com> (search for PowerReset). The EMC Knowledgebase Solutions emc69100 and emc104555 provide more information.

When using PowerPath with HACMP in a VNX OE and CLARiiON environment, “[cfgscsi_id binary in a VNX OE and CLARiiON environment](#)” on page 70 provides information.

cfgscsi_id binary in a VNX OE and CLARiiON environment

In a VNX OE and CLARiiON environment the cfgscsi_id binary is required in addition to the emcpowerreset binary. When you install PowerPath version 5.3, 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.

Adding the `cfgscsi_id` binary

To set up an HACMP cluster in a VNX OE and CLARiiON environment:

Note: EMC Knowledgebase Solution emc143075 provides more information on the `cfgscsi_id` binary.

1. Add the custom cluster event to your configuration. This event is the name given to the script, which you later add to select predefined HACMP events. Type:

Note: This is a single command.

```
/usr/es/sbin/cluster/utilities/claddcustom -t event -n'cfgscsi_id' -I'Set correct
scsi id on EMC CLARiiON pseudo devices.' -v'/usr/sbin/cfgscsi_id'
```

2. Verify your custom cluster event was added by running:

```
odmget HACMPcustom
```

Output similar to the following appears:

```
HACMPcustom:
name = "cfgscsi_id"
type = "event"
description = "Set correct scsi id on EMC CLARiiON pseudo devices."
value = "/usr/sbin/cfgscsi_id"
relation = ""
status = 0
```

3. Modify the predefined HACMP event by giving the event command your custom cluster event as a pre-event command. Type:

```
/usr/es/sbin/cluster/utilities/clchevent -O'node_up'
-s'/usr/es/sbin/cluster/events/node_up' -b 'cfgscsi_id' -c '0'

/usr/es/sbin/cluster/utilities/clchevent -O'node_down'
-s'/usr/es/sbin/cluster/events/node_down' -b 'cfgscsi_id' -c '0'
```

4. Verify the `node_up` event was properly modified, by running:

```
odmget -q name=node_up HACMPevent
```

Output similar to the following appears:

```
HACMPevent:
name = "node_up"
desc = "Script run when a node is attempting to join the cluster."
setno = 101
msgno = 7
```

```

catalog = "events.cat"
cmd = "/usr/es/sbin/cluster/events/node_up"
notify = ""
pre = "cfgscsi_id"
post = ""
recv = ""
count = 0
event_duration = 0

```

5. Verify that the node_down event was properly modified, by running:

```
odmget -q name=node_down HACMEvent
```

Output similar to the following appears:

```

HACMEvent:
  name = "node_down"
  desc = "Script run when a node is attempting to leave the cluster."
  setno = 101
  msgno = 8
  catalog = "events.cat"
  cmd = "/usr/es/sbin/cluster/events/node_down"
  notify = ""
  pre = "cfgscsi_id"
  post = ""
  recv = ""
  count = 0
  event_duration = 0

```

6. Synchronize your cluster and ensure there are no errors resulting from the addition of the pre-event scripts.

Note: The modified script must reside in the same location on all nodes in the cluster, and the script must be made executable on all nodes. Otherwise, synchronization for the cluster fails.

Synchronizing volume group information between cluster nodes

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

1. Run the `/usr/lpp/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. Create a PVID on the new devices by running:

```
chdev -l hdiskpowerx -a pv=yes
```
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 an 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.

⚠ CAUTION

With PowerPath 5.3, 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 106](#) provides information to help you solve this problem.

EMC Knowledgebase Solution emc74494 provides more information.

PowerPath in an EMC Legato AAM cluster

This section describes how to:

- ◆ Install PowerPath and EMC Legato® Automated Availability Manager (AAM) in a new cluster, that is, where neither PowerPath nor the AAM software is installed on any host to be included in the cluster.
- ◆ Integrate PowerPath into an existing AAM cluster.

Installing PowerPath in a New AAM Cluster

To install and configure PowerPath and AAM when neither PowerPath nor AAM is installed:

1. Prepare the cluster environment. Refer to the relevant AAM documentation.
2. Verify that all storage system devices are seen by each host.
3. Install PowerPath on all nodes. [Chapter 1, “Installing PowerPath on an AIX host”](#) provides more information.
4. Identify the LVM volume groups that use PowerPath hdiskpower devices:
 - a. Run **powermt display dev=all** to list each hdiskpower device.
 - b. Run **lspv** to identify the volume group to which each hdiskpower device belongs.
5. Install AAM on all nodes. Refer to the relevant AAM documentation. Configure AAM to use the volume groups identified in step 4. Make sure they are the same volume groups on each host.
6. If the AAM agent is not already started, start the agent on each node in the cluster, using either the **ft_startup command** or the Management Console.

Integrating PowerPath into an AAM cluster

To integrate PowerPath into an existing AAM 5.3 (or later) cluster, follow these steps on each cluster node, one node at a time:

1. Stop cluster services on the node using either the **ft_shutdown** command or the Management Console.
2. Install or upgrade PowerPath on the node. [Chapter 1, “Installing PowerPath on an AIX host”](#) provides additional information.
3. Start cluster services on the node using either the **ft_startup** command or the Management Console, and wait for the node to be fully reintegrated into the cluster.

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.

Removing PowerPath from an AIX Host

This chapter includes the following sections:

- ◆ Before Removing PowerPath 80
- ◆ Removing PowerPath..... 81
- ◆ After removing PowerPath..... 83
- ◆ When a storage system device is the boot device 83

Before Removing PowerPath

Before you remove PowerPath from the host:

- ❑ Check the Powerlink website for the most current information. EMC updates PowerPath release notes periodically and posts them on the Powerlink website.
- ❑ 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.

- ❑ Vary off all LVM volume groups that use PowerPath hdiskpower devices. Otherwise, the uninstall procedure will fail.
- ❑ If EMC Solutions Enabler daemons, such as **storwatchd** and **storapid**, are running, type **stordaeomon 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 VNX OE and 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.

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

⚠ 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.

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.base	5.3.0.0	USR	DEINSTALL	SUCCESS
EMCpower.migration_enabler	5.3.0.0	USR	DEINSTALL	SUCCESS
EMCpower.mpx	5.3.0.0	USR	DEINSTALL	SUCCESS
EMCpower.consistency_grp	5.3.0.0	USR	DEINSTALL	SUCCESS
EMCpower.encryption	5.3.0.0	USR	DEINSTALL	SUCCESS

PowerPath is now removed from the host. “After removing PowerPath” on page 83 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. Open SMIT, by running:
`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.base	5.3.0.0	USR	DEINSTALL	SUCCESS
EMCpower.migration_enabler	5.3.0.0	USR	DEINSTALL	SUCCESS
EMCpower.mpx	5.3.0.0	USR	DEINSTALL	SUCCESS
EMCpower.consistency_grp	5.3.0.0	USR	DEINSTALL	SUCCESS
EMCpower.encryption	5.3.0.0	USR	DEINSTALL	SUCCESS

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

PowerPath is now removed from the host. [“After removing PowerPath” on page 83](#), 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**.
“[The pprootdev tool](#)” on [page 46](#) provides information on disabling multipathing to the root device.
2. Restart the host.
3. Remove PowerPath as described in “[Removing PowerPath](#)” on [page 81](#).
4. Run **bosboot -ad /dev/ipldevice** to create a boot image.

This chapter includes the following sections:

◆ Management daemon using SNMP	86
◆ emc_cfgmgr script.....	90
◆ PowerPath hdiskpower devices.....	90
◆ PowerPath and the lsvg command.....	94
◆ PowerPath and the iostat command	95
◆ BCVs and cfgmgr operations	95
◆ Bringing hdiskpower-based BCV Symmetrix logical devices online	95
◆ Importing an LVM volume group from a remote host.....	96
◆ Changing the target/LUN address of a storage system logical device	97
◆ Adding new devices to an existing configuration	97
◆ Replacing an HBA that PowerPath is using online	98
◆ Moving LUNs in and out of a storage group.....	98
◆ Eliminating ghost devices.....	98
◆ Reconfiguring PowerPath devices online	99
◆ Removing paths or logical devices from a PowerPath configuration.....	100
◆ Failover in switched environments	102
◆ SMIT screens	102
◆ Audit and error messages.....	103
◆ Troubleshooting.....	106

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_mond.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 is 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 87](#) 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	Any/all paths are dead for five minutes.	0x20,0x40,300	Yes
	Any/all paths are dead and then recovers within five minutes.	0x20,0x40,300	No

The time threshold for each event can be modified as described in [“Editing the event” on page 88](#).

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 Version 5.3 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.

Procedure To configure the management daemon:

1. Edit the IP address to 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.

For example,

```
< trap    public    10.x.xx.xxx  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.

For example,

```
<TARGET_ADDRESS Target1 UDP 10.x.xx.xxx    traptag trapparms1---
```

2. Restart the SNMP daemon (snmpd), and then check the status:
 - a. Run **stopsrc -s snmpd**.
 - b. Run **startsrc -s snmpd**.
 - c. Run **lssrc -s snmpd**.

Output similar to the following appears:

Subsystem	Group	PID	Status
snmpd	tcpip	32258	active

3. Start the emcp_mond daemon, by running:
`/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.

Note: Events #3 and #4 are examples and are not supported in PowerPath 5.3.

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. Under certain circumstances, however, AIX does not configure an hdisk for each logical path to a storage system logical device.

Suppose, 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.



[OPT_353776][Affected Docs: IAG of 5.3 and 5.5, RN of 5.3 and 5.5] 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/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

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. The *EMC PowerPath Family CLI and System Messages Reference Guide* provides additional details.

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 stores 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 vgname** 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. [“The pprootdev tool” on page 46](#) provides additional information. 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 97](#) 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 could not 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/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 could not 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/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 PowerPath thinks 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**.

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 VNX OE and 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/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/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/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.

Removing paths or logical devices from a PowerPath configuration

This section describes how to remove the following from a PowerPath configuration:

- ◆ Specified paths to logical devices
- ◆ Logical devices
- ◆ Entire HBAs

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**.

To remove paths or logical devices from the PowerPath configuration:

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.

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 nevertheless 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 Guide* contains a complete list of PowerPath error messages.

Audit log messages

The powermig commands for PowerPath Migration Enabler, most powermt commands and emcpadm commands generate audit messages that allow you to track the commands run on a particular host. The following emcpadm and powermt commands generate audit messages:

- ◆ emcpadm rename
- ◆ powermt check
- ◆ powermt config
- ◆ powermt disable hba
- ◆ powermt enable hba
- ◆ powermt load
- ◆ powermt manage
- ◆ powermt remove
- ◆ powermt restore
- ◆ powermt save
- ◆ powermt set mode
- ◆ powermt set path_latency_monitor
- ◆ powermt set path_latency_threshold
- ◆ powermt set periodic_autorestore
- ◆ powermt set policy
- ◆ powermt set port_disable
- ◆ powermt set priority
- ◆ powermt set transparent
- ◆ powermt set write_throttle
- ◆ powermt set write_throttle_queue

- ◆ powermt unmanage
- ◆ powermt update lun_names

The *EMC PowerPath Family CLI and System Messages Reference Guide*, available on the Powerlink website, contains more information on these commands.

Message format and description

Audit log messages have the following format:

```
<timestamp> <hostname> EMCPP: emcpAudit: Info: cmd=<message> (<user_ID>)
```

For example,

```
Oct 28 13:06:57 Host1a local0:info EMCPP: emcpAudit: Info: cmd=powermt:manage
class=clariion(user ID real=0 effective=0)
```

where:

- <timestamp> is the date and time the event occurred.
- <hostname> is the name of the host where PowerPath resides.
- EMCPP: emcpAudit: is the program performing the auditing function. This is the tag used for audit messages for PowerPath.
- cmd is the utility that was used.
- <message> is the audit message.
- <user_ID> is the identification of the user.

Error log messages

Error log messages appear along with audit messages in the log file when you configure the AIX operating system for common logging as described in [“Enabling logging on an AIX host” on page 105](#). 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 *EMC PowerPath Family CLI and System Messages Reference Guide* contains a complete list of PowerPath error messages. The AIX documentation contains information on the AIX errlog/errpt facility.

Message format and description

Error log messages have the following format:

```
<timestamp> <hostname> kernel: <level>:<module>:<message>
```

For example,

```
Sep  7 10:20:45 Host1a kernel: Error:Mpx:Path Bus 5 Tgt 0 Lun 0 to FCNBD054405535
is dead.
```

where:

- *<timestamp>* is the date and time the event occurred.
- *<hostname>* is the name of the host where PowerPath resides.
- *<level>* is the message type (**Info**, **Error**, or **Warning**).
- *<module>* is the program or module that is reporting the message.
- *<message>* is the error message. The *EMC PowerPath Family CLI and System Messages Reference Guide*, available on the Powerlink website, contains the error message descriptions.

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 14](#) 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

With PowerPath 5.3, 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.

Files Changed by PowerPath Installation

This appendix includes the following sections:

- ◆ Files added by installation..... 110
- ◆ Files modified by installation..... 114
- ◆ ODM modifications caused by PowerPath installation 114

Files added by installation

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

```

/etc/emc/emcp_mond.conf
/etc/emc/rsa/cst/xml/Config-Template.xml
/etc/emc/rsa/cst/xml/Config.xml
/etc/emc/rsa/cst/xml/CSP.xml
/etc/emc/rsa/cst/xml/CSP-Admin.xml
/etc/emc/rsa/cst/xml/csp-app.xml
/etc/emc/rsa/cst/xml/CSP-Authn.xml
/etc/emc/rsa/cst/xml/CSP-CatalogPosix.xml
/etc/emc/rsa/cst/xml/CSP-CatalogService.xml
/etc/emc/rsa/cst/xml/CSP-CatalogWindows.xml
/etc/emc/rsa/cst/xml/CSP-Config.xml
/etc/emc/rsa/cst/xml/CSP-DestinationAlert.xml
/etc/emc/rsa/cst/xml/CSP-DestinationEventLog.xml
/etc/emc/rsa/cst/xml/CSP-DestinationFile.xml
/etc/emc/rsa/cst/xml/CSP-DestinationODBC.xml
/etc/emc/rsa/cst/xml/CSP-DestinationSyslog.xml
/etc/emc/rsa/cst/xml/CSP-LayoutFlat.xml
/etc/emc/rsa/cst/xml/CSP-LayoutRaw.xml
/etc/emc/rsa/cst/xml/CSP-LayoutXML.xml
/etc/emc/rsa/cst/xml/CSP-LDAP.xml
/etc/emc/rsa/cst/xml/CSP-LocalDirectory.xml
/etc/emc/rsa/cst/xml/CSP-LocalizationService.xml
/etc/emc/rsa/cst/xml/CSP-Logging.xml
/etc/emc/rsa/cst/xml/CSP-OSLogin.xml
/etc/emc/rsa/cst/xml/CSP-PositionalFormatter.xml
/etc/emc/rsa/cst/xml/CSP-RoleManagement.xml
/etc/emc/rsa/cst/xml/CSP-SecurID.xml
/etc/emc/rsa/cst/xml/CSP-SM.xml
/etc/emc/rsa/cst/xml/CSP-SprintfFormatter.xml
/etc/emc/rsa/cst/xml/CSP-SSC.xml
/etc/emc/rsa/cst/xml/CSP-SSC-FOB.xml
/etc/emc/rsa/cst/xml/CSP-TrustedUser.xml
/etc/emc/rsa/cst/xml/CSP-UserManagement.xml
/etc/emc/rsa/rkm_client/config/rkm_keyclass.conf.tmpl
/etc/emc/rsa/rkm_client/config/rkm_init.conf.tmpl
/etc/emc/rsa/rkm_client/config/rkm_svc.conf.tmpl
/etc/rc.emcp_mond
/etc/rc.emcp_xcryptd
/etc/rc.emcpower
/etc/rc.powermig
/etc/set_scsi_id
/usr/lib/boot/protoext/disk.proto.ext.scsi.pseudo.power
/usr/lib/drivers/cgext

```

```
/usr/lib/drivers/dmext
/usr/lib/drivers/gpxext
/usr/lib/drivers/mpxext
/usr/lib/drivers/powerdd
/usr/lib/drivers/powerdiskdd
/usr/lib/drivers/sapiext
/usr/lib/drivers/vlumext
/usr/lib/drivers/xcryptext
/usr/lib/libaceclnt.so
/usr/lib/libbccme.so
/usr/lib/libbccme_base.so
/usr/lib/libbccme_ecc.so
/usr/lib/libbccme_eccaccel.so
/usr/lib/libcg.a
/usr/lib/libcg.so
/usr/lib/libcryptocme2.so
/usr/lib/libcryptocme2.sig
/usr/lib/libCSP.so
/usr/lib/libCSP-Admin.so
/usr/lib/libCSP-AuthenticationService.so
/usr/lib/libCSP-C.so
/usr/lib/libCSP-CatalogPosix.so
/usr/lib/libCSP-CatalogService.so
/usr/lib/libCSP-CatalogWindows.so
/usr/lib/libCSP-Configuration.so
/usr/lib/libCSP-DestinationAlert.so
/usr/lib/libCSP-DestinationEventLog.so
/usr/lib/libCSP-DestinationFile.so
/usr/lib/libCSP-DestinationODBC.so
/usr/lib/libCSP-DestinationSyslog.so
/usr/lib/libCSP-LayoutFlat.so
/usr/lib/libCSP-LayoutRaw.so
/usr/lib/libCSP-LayoutXML.so
/usr/lib/libCSP-lb.so
/usr/lib/libCSP-LDAP.so
/usr/lib/libCSP-LocalDirectory.so
/usr/lib/libCSP-LocalizationService.so
/usr/lib/libCSP-LogService.so
/usr/lib/libCSP-OSLogin.so
/usr/lib/libCSP-PositionalFormatter.so
/usr/lib/libCSP-RoleManagement.so
/usr/lib/libCSP-RT.so
/usr/lib/libCSP-SecurID.so
/usr/lib/libCSP-SimpleCrypto.so
/usr/lib/libCSP-SM.so
/usr/lib/libCSP-SprintfFormatter.so
/usr/lib/libCSP-SSC.so
```

```
/usr/lib/libCSP-TrustedUser.so
/usr/lib/libCSP-UserManagement.so
/usr/lib/libCSPI.so
/usr/lib/libCSPJNIWrappers.so
/usr/lib/libemcp.a
/usr/lib/libemcp.so
/usr/lib/libemcp_core.a
/usr/lib/libemcp_core.so
/usr/lib/libemcp_event.a
/usr/lib/libemcp_event.so
/usr/lib/libemcp_lam.a
/usr/lib/libemcp_lam.so
/usr/lib/libemcp_lic_rtl.so
/usr/lib/libemcp_mp.a
/usr/lib/libemcp_mp_rtl.so
/usr/lib/libemcp_power.a
/usr/lib/libemcp_shlib.a
/usr/lib/libemcp_vlumd.a
/usr/lib/libemcp_vlumd.so
/usr/lib/libemcp_xcrypt.a
/usr/lib/libemcp_xcrypt.so
/usr/lib/libemcpvt_api.a
/usr/lib/libemcpvt_api.so
/usr/lib/libexpat.so
/usr/lib/libkmclient_shared.so
/usr/lib/libkmcryptolib.so
/usr/lib/libkmsvcshlib.so
/usr/lib/libLBJNIWrappers.so
/usr/lib/liblber.so
/usr/lib/liblber-2.3.so
/usr/lib/liblber-2.3.so.0
/usr/lib/liblber_r-2.3.so.0
/usr/lib/libldap.so
/usr/lib/libldap-2.3.so
/usr/lib/libldap-2.3.so.0
/usr/lib/libldap_r.so
/usr/lib/libldap_r-2.3.so
/usr/lib/libmp.a
/usr/lib/libosysrt.so
/usr/lib/libosysrtxml.so
/usr/lib/libpn.a
/usr/lib/libpn.so
/usr/lib/libpower.a
/usr/lib/methods/cfgpower
/usr/lib/methods/cfgpowerdisk
/usr/lib/methods/chgpowerdisk
/usr/lib/methods/CSP-MessageCatalogen-US.cat
```

```

/usr/lib/methods/power.cat
/usr/lib/methods/powerpath/cfgdm
/usr/lib/methods/powerpath/cfggpx
/usr/lib/methods/powerpath/cfgmpx
/usr/lib/methods/powerpath/cfgudid
/usr/lib/methods/powerpath/cfgvlumd
/usr/lib/methods/powerpath/cfgxcrypt
/usr/lib/methods/ucfgpower
/usr/lib/methods/ucfgpowerdisk
/usr/lib/methods/undpower
/usr/lib/methods/undpowerdisk
/usr/lib/nls/msg/en_US/CSP-MessageCatalogen-US.cat
/usr/lib/nls/msg/en_US/power.cat
/usr/sbin/cfgscsi_id
/usr/sbin/ckmadm
/usr/sbin/cstadmin
/usr/sbin/emcp_mond
/usr/sbin/emcp_setup.csh
/usr/sbin/emcp_setup.sh
/usr/sbin/emcp_xcryptd
/usr/sbin/emcpadm
/usr/sbin/emcpassive2active
/usr/sbin/emcpmgr
/usr/sbin/emcpmigd
/usr/sbin/emcpminor
/usr/sbin/emcpreg
/usr/sbin/emcpupgrade
/usr/sbin/powercf
/usr/sbin/powermig
/usr/sbin/powermt
/usr/sbin/powervt
/usr/sbin/powervxvm
/usr/sbin/pprootdev
/usr/sbin/xcrypt_config
/usr/sbin/xcrypt_upgrade
/usr/share/man/man1/emcpadm.1
/usr/share/man/man1/emcpminor.1
/usr/share/man/man1/emcpreg.1
/usr/share/man/man1/emcpupgrade.1
/usr/share/man/man1/powermig.1
/usr/share/man/man1/powermt.1
/usr/share/man/man1/powervt.1
/usr/share/man/man1/powervxvm.1
/usr/share/man/man5/rkm_init.conf.5
/usr/share/man/man5/rkm_keyclass.conf.5
/usr/share/man/man5/rkm_svc.conf.5
/usr/share/man/man8/ckmadm.8

```

`/usr/share/man/man8/cstadmin.8`
`/usr/share/man/man8/xcrypt_config.8`

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`.

ODM modifications caused by PowerPath installation

When you install PowerPath on AIX, the following ODM modifications occur:

- ◆ Additions to PdDv for hdiskpower and powerpath0 devices
- ◆ Additions to PdAt for attributes for hdiskpower and powerpath0 devices
- ◆ Additions to SMIT menus for PowerPath controls
- ◆ Addition of Config_Rule to configure PowerPath after reboot
- ◆ Updating of PdCn so the **lsparent** command works for PowerPath devices
- ◆ Addition to CuDv for a powerpath0 device

The following ODM modifications occur during PowerPath device configuration (powermt config):

- ◆ Additions to CuVPD for hdiskpower devices
- ◆ Additions to CuAt for hdiskpower and powerpath0 devices
- ◆ Additions to CuDv for hdiskpower devices
- ◆ Additions to CuDvDr for hdiskpower and powerpath0 devices

A

AAM. *See* EMC Legato AAM cluster
adding devices to existing configuration 97
adding logical devices 60
adding paths, while booting off the array 55

B

BCV devices, hdiskpower-based 96
boot device
 configuring 44
 disabling 56
booting off the array, adding paths while 55
bosboot tool 47

C

CLARiiON storage systems
 booting on 45
 setting up 14
 trespass 47
cluster. *See also* HACMP cluster, EMC Legato
 AAM cluster, VCS cluster
configuring, boot device 44

D

device mapping, changing 97
device naming conventions 90
devices, adding online 99
disabling multipathing 35
DMP (Dynamic Multipathing) 41
documentation, related 7, 8

E

ECC. *See* EMC ControlCenter
EMC ControlCenter 62
EMC Legato AAM cluster
 installing with PowerPath 75
 integrating PowerPath into 76
emc_cfgmgr script 18, 90
emcp_mond.conf 86
emcpassive2active command 47
error messages 104

G

ghost devices, removing 53

H

HACMP cluster
 failover problems 73
 integrating with PowerPath 69
HBA, replacing 98
hdisk devices, configuring for PowerPath 90
hdisk, configuring, troubleshooting 106
hdisk, deleting, troubleshooting 106
hdiskpower devices 91
 initializing 27

I

installing PowerPath
 command line 20
 download from Powerlink 22
 mksysb master image, as part of a 23
 postinstallation instructions 25
 preinstallation instructions 14

- service packs 26
- SMIT (System Maintenance Interface Tool)
 - procedure 21
- iostat command output 95

L

- licensing PowerPath 25
- logical devices, adding 60
- logical devices, removing from PowerPath
 - configuration 63, 100
- LPP (Licensed Program Products), required 15
- lsvg command 94

M

- management daemon
 - configuring 87
 - definition 86
 - restarting 89
- mksysb master image
 - installing PowerPath as part of 23
- monitored events, SNMP 86
- multipathing, disabling 35

N

- naming conventions 90

P

- patches. *See* service packs
- powermig transition -startup 31
- PowerPath boot device
 - pprootdev command 46
- PowerPath Configuration Checker (PPCC) 15, 30
- PowerPath devices
 - hdiskpower 90
 - reconfiguring online 99
- PowerPath error messages 104
- pprootdev tool 35, 46, 47, 56
- pseudo devices 90
- PVID (physical volume identifier) 92

R

- reconfiguring PowerPath devices online 99
- registering PowerPath 25

- removing logical devices from PowerPath
 - configuration 63, 100
- removing PowerPath 80
- removing PowerPath boot device 83

S

- service packs, installing 26
- SMIT (System Maintenance Interface Tool)
 - accessing screens 102
 - installing by using 21
- SNMP management daemon 86
- Symmetrix storage arrays 45

T

- tar file, installing 22
- time threshold, editing 88
- trap destination IP 87
- troubleshooting 106
 - removing ghost devices 53
 - trespass occurrence 47
 - upgrade, PowerPath 34

U

- uninstalling PowerPath 80
- unique object identifier 87
- upgrading PowerPath 30

V

- varyonvg vname command 28
- VCS cluster
 - installing with PowerPath 77
 - integrating PowerPath into 78
- VNX OE storage systems, setting up 14
- VNX storage systems, booting on 45
- VNX storage systems, trespass 47
- volume groups, LVM
 - importing 96
 - varying off 15
 - varying on 28
- VxVM (Veritas Volume Manager) and
 - hdiskpower devices 40