



EMC® PowerPath® for Linux
Version 5.5

Installation and Administration Guide

P/N 300-011-249
Rev A04

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

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

Published September, 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	7
Chapter 1	Installing PowerPath on a Linux host
Before you install PowerPath.....	12
Obtaining up-to-date information.....	12
License keys.....	12
Check storage systems	13
Check hosts.....	14
Check PowerPath major number	15
Install Linux updates.....	16
Select language for installation.....	16
IA64 systems only.....	18
Installing PowerPath.....	19
Installing PowerPath in a Xen environment.....	22
Installing PowerPath if Veritas Storage Foundation is already installed.....	25
Registering PowerPath.....	26
After installing PowerPath	28
Configure the management daemon	28
Verify your PowerPath capabilities are correct.....	28
Reconfiguring device mappings.....	28
Verify extensions are loaded.....	29
Edit the boot.vxvm file.....	31
Upgrading PowerPath	32
Before you upgrade PowerPath.....	32
Upgrading from a CD	33
Upgrading the Linux operating system	35
Upgrading PowerPath from one major release to next major release	36

Upgrading PowerPath in a SteelEye LifeKeeper cluster environment.....	36
Upgrading PowerPath Encryption	37
Coexistence with third-party path management software	38
PowerPath behavior	38

Chapter 2 Configuring a PowerPath root device on Linux

Introduction.....	42
Root device considerations.....	42
Configuring PowerPath in a boot-from-SAN setup (SLES 10 SP4)	43
Configuring PowerPath in a boot-from-SAN setup (SLES11 SP1)	45
Configuring PowerPath in a boot-from-SAN setup (RHEL 5 and updates)	47
Upgrading the Linux kernel in a boot-from-SAN setup.....	50
Upgrading the Linux kernel in a boot-from-SAN setup (SLES11SP1)	51
Upgrading PowerPath in a boot-from-SAN setup	53
Removing PowerPath in a boot-from-SAN setup.....	54
Removing PowerPath in a boot-from-SAN setup (SLES11SP1)	55

Chapter 3 Configuring a PowerPath logical device

Introduction.....	58
Adding new paths to a PowerPath logical device	59
Removing some paths to PowerPath logical devices	61

Chapter 4 Maintaining PowerPath on Linux

Management daemon using SNMP	66
Monitored events.....	66
Configuring the management daemon	68
Editing the event.....	69
Installing QLogic qla2300 HBAs after PowerPath installation	70
Saving new configurations.....	71
Configuring LVM2 support	72
Root file system not mounted on a logical volume	72
Root file system mounted on a logical volume	73
Root file system mounted on mirrored logical volumes	74
Dynamic LUN addition and removal.....	76

Adding a LUN.....	76
Removing a LUN	77
Hot swapping an HBA.....	78
Removing an HBA.....	78
Adding an HBA.....	79
R1/R2 boot failover support	80
R1/R2 supported configurations.....	80
R1/R2 boot procedure.....	81
Configure the network interface on the R1 host.....	81
Setting the naming convention to enclosure-based names.....	83
SCSI-2 reservation support in PowerPath for Linux.....	83
Audit and error messages	83
Chapter 5	Removing PowerPath from a Linux host
Before you remove PowerPath.....	86
Uninstalling PowerPath	87
Chapter 6	Troubleshooting
Recovering from a failed PowerPath installation.....	90
Troubleshooting a PowerPath upgrade.....	91
PowerPath custom settings not preserved	91
Multiple PowerPath entries in the RPM database	93
Missing or deleted files.....	93
Incorrect PowerPath configuration	94
The _netdev option in /etc/fstab	95
Problems booting the host	95
System hangs	95
Unsupported HBA drivers	95
Other causes.....	96
Appendix A	Files changed by PowerPath
Linux files modified by PowerPath installation	98
Files created by PowerPath installation	99
x86_64 hosts only	103
ia64 hosts only	103
x86_64 and ia64 hosts	103
Driver files.....	104
Encryption files	106
Index	113



As part of its effort to continuously improve and enhance the performance and capabilities of the EMC product line, EMC periodically releases new versions of PowerPath. Therefore, some functions described in this manual may not be supported by all versions of PowerPath or the storage system hardware it supports. For the most up-to-date information on product features, refer to your product release notes.

If a PowerPath feature does not function properly or as described in this manual, please contact the EMC Customer Support Center for assistance. Refer to "Where to get help" on page 10 in this Preface for contact information.

This guide describes how to install and remove PowerPath for Linux Version 5.5 and its minor releases.

Audience

This document is intended for storage administrators and other information system professionals responsible for installing and maintaining PowerPath in a Linux environment.

In addition to understanding PowerPath, administrators should be familiar with the host operating system where PowerPath runs and the applications used with PowerPath.

Related documentation

The following set of PowerPath documentation is available on the EMC Powerlink website:

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

◆ *EMC PowerPath Encryption with RSA User Guide*

If your environment includes EMC Symmetrix storage systems, refer also to the EMC host connectivity guides, which are available on the EMC Powerlink website <http://Powerlink.EMC.com>.

If your environment includes EMC CLARiiON storage systems, refer also to the EMC host connectivity guides and the CLARiiON Storage-System Support website (www.EMC.com/clariionsupport).

Revision History

The following table presents the revision history of this document.

Revision	Date	Description
A01	October, 2010	First release of the EMC PowerPath for Linux Version 5.5 Installation and Administration Guide.
A02	November 8, 2010	Second release of the EMC PowerPath for Linux Version 5.5 Installation and Administration Guide.
A03	June 03, 2011	Updated for the SLES 10 SP4 support.
A04	September 12, 2011	Updated Chapter 4, "Maintaining PowerPath on Linux," at page 4-73 .

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

Conventions used in this document

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

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



CAUTION

A caution contains information essential to avoid data loss or damage to the system or equipment. The caution may apply to hardware or software.



IMPORTANT

An important notice contains information essential to operation of the software. The important notice applies only to software.

Typographical conventions

EMC uses the following type style conventions in this document:

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
- Path names, URLs, filenames, directory names, computer names, links, groups, service keys, file systems, environment variables (for example, command line and text), notifications

Bold

In procedures:

- Names of dialog boxes, buttons, icons, menus, fields
- Selections from the user interface, including menu items and field entries
- Key names
- Window names

In running text:

- Command names, daemons, options, programs, processes, notifications, system calls, man pages, services, applications, utilities, kernels

Italic

Used for:

- Full publications titles referenced in text
- Unique word usage in text

Bold Italic

Anything requiring extra emphasis

`Courier`

Used for:

- System output
- Filenames
- Complete paths
- Command-line entries
- URLs

`Courier, bold`

Used for:

- User entry
- Options in command-line syntax

`Courier, italic`

Used for:

- Arguments used in examples of command-line syntax
- Variables in examples of screen or file output
- Variables in path names

`Courier, bold, italic`

Variables used in a command-line sample

< >	Angle brackets enclose parameter or variable values supplied by the user
[]	Square brackets enclose optional values
	Vertical bar indicates alternate selections - the bar means “or”
{ }	Braces indicate content that you must specify (that is, x or y or z)
...	Ellipses indicate nonessential information omitted from the example

Where to get help

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

Product information — For documentation, release notes, software updates, or 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.

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 a Linux host

This chapter includes the following sections:

- ◆ Before you install PowerPath 12
- ◆ Installing PowerPath 20
- ◆ After installing PowerPath 29
- ◆ Upgrading PowerPath..... 33
- ◆ Coexistence with third-party path management software 39

Before you install PowerPath

This section describes how to prepare a Linux host for an EMC® PowerPath® installation.

Obtaining up-to-date information

Check the Powerlink website for current information:

- ❑ PowerPath documentation — EMC updates PowerPath documentation 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**.

License keys

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

Take note of the following regarding licenses:

- ◆ In nonvirtual environments, installing PowerPath on a host connected exclusively to a CLARiiON AX-series array does not require a license. PowerPath provides full support, with or without a PowerPath license, when the host is connected exclusively to a CLARiiON AX-series array.
- ◆ If you are upgrading from an earlier version of PowerPath, you do not need to reregister. PowerPath will use your old key.
- ◆ Separate license keys are required to enable PowerPath Migration Enabler and PowerPath Encryption with RSA®.
- ◆ Type the registration number exactly as it appears on the card or letter. To reduce common typographical errors, the **License Key** field accepts either uppercase or lowercase letters, and certain numbers and letters are interchangeable. Specifically, an entry of the alphabetic letters O, I, S, and B is equivalent to an entry of the numbers, 0, 1, 5, and 8.

Prepare the host and storage system

Go to Powerlink and verify that your environment meets the requirements found in the *EMC PowerPath for Linux Release Notes* and the E-Lab™ Interoperability Navigator.

To verify, download the latest version of the EMC Grab utilities from Powerlink, and then run the PowerPath Configuration Checker (PPCC). For PPCC installation information, the *EMC PowerPath Configuration Checker User Guide*, available on Powerlink, provides additional information.

Check storage systems

The sequence of steps for configuring a storage system and installing PowerPath on a host depends on which storage system you use.

- ❑ **EMC Symmetrix® storage systems** — Install PowerPath *after* you set up and verify that the Symmetrix is working properly. Then, install PowerPath using the instructions in this guide.
- ❑ **EMC CLARiiON® storage systems:**

PowerPath installation is an integral part of a CLARiiON setup and configuration procedure.

- ❑ For information on when to install PowerPath on a host attached to a CLARiiON storage system, refer to the CLARiiON storage-system support website (www.EMC.com/clariionsupport). From the home page, select your storage system type, then use the Attach server link (CX4 systems) or the Install link (all other systems) to generate step-by-step instructions.
- ❑ First-time PowerPath installations — Remove any version of EMC Navisphere® Application Transparent Failover (ATF) installed on a host that supports this application.

Note: The procedure for migrating from ATF or 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, refer to *Removing ATF or CDE Software Before Installing Other Failover Software* for more information. This document is available only on the EMC Powerlink website.

- ❑ **Third-party storage systems** — Install PowerPath *after* you set up and verify that the third-party storage system is working properly. Then, install PowerPath using the instructions in this guide.

Check hosts

Before installing PowerPath on any Linux host, ensure that the host meets the following requirements:

- ❑ The EMC *PowerPath for Linux Release Notes* describes the operating system versions and service packs needed to support PowerPath. Refer also to the E-Lab™ Interoperability Navigator on the Powerlink website, for information on host-storage system interconnection topologies and specific PowerPath requirements.
- ❑ EMC recommends that all hosts have at least 2 GB of memory.
- ❑ If the host was not installed with the default file system layout provided by the operating system vendor, then PowerPath may report "no such file or directory" during system initialization and the Linux host may not boot with PowerPath installed. EMC supports only the default system layout (s) as supplied by the Operating System vendor.
- ❑ Ensure that the storage-system logical devices are configured for PowerPath support. The EMC host connectivity guides for both Symmetrix and CLARiiON systems and the CLARiiON storage-system support website, contains additional information.
- ❑ Ensure that the root partition has at least 120 MB of free space.

Note: If you attempt to install PowerPath on a host that lacks sufficient space, installation fails as expected. RPM, however, returns a misleading message about the amount of space needed to install PowerPath. ["Recovering from a failed PowerPath installation" on page 90](#) contains more information.

- ❑ Configure the HBA BIOS and drivers if you have not already done so.



CAUTION

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

- ❑ If you plan to install PowerPath and Veritas Volume Manager (VxVM), install PowerPath first if possible.

If you install PowerPath on a host on which Veritas Storage Foundation is already installed, "[Installing PowerPath if Veritas Storage Foundation is already installed](#)" on page 26 provides additional information.

- ❑ PowerPath is not compatible with the native Linux device mapper (DM-MPIO). Configuring both products on the same host can cause system instability. EMC recommends that you blacklist all non-local devices on a host on which PowerPath will be installed.
- ❑ On RHEL5.x hosts, dm multipath modules are installed by default and are loaded when the host is started. Make sure that the following is not commented in `/etc/multipath.conf`

```
# Blacklist all devices by default. Remove this to
  enable multipathing
# on the default devices.
blacklist { devnode "*"
}
```

- ❑ On a SuSE host, run `chkconfig | grep boot.multipath` to ensure that the native multipath service is not enabled.

If the output indicates that the multipath service is on, run `chkconfig boot.multipath off` to disable it before installing PowerPath.

Note: If the `chkconfig | grep boot.multipath` command returns no output, then the multipathing service is not installed on the host.

- ❑ Review the Linux kernel updates as described in the *EMC Host Connectivity with QLogic Fibre Channel and iSCSI Host Bus Adapters (HBAs) and Converged Network Adapters (CNAs) for the Linux environment* guide. Review the requirements and do all that apply in your configuration. The E-Lab Interoperability Navigator on the Powerlink website provides information on which HBA or fibre drivers support PowerPath with your host.

Check PowerPath major number

Verify that `/usr/include/linux/major.h` and `/proc/devices` does not include or define a major number of 120. After PowerPath

installation, the `/proc/devices` file shows 120 power2 under the block heading for emcpower devices.

Install Linux updates

The Linux kernel and fibre drivers supported by your version of PowerPath must be installed and running.

Select language for installation

The PowerPath for Linux installation and startup procedure has been localized for the following languages.

Table 1 Supported languages

Language	Locale ID	Encoding
Chinese (simplified)	zh_CN.utf8	UTF8
English	En_US.utf8	
French	fr_FR.utf8	
German	de_DE.utf8	
Italian	it_IT.utf8	
Japanese	ja_JP.utf8	
Korean	ko_KR.utf8	
Portuguese (Brazilian)	pt_BR.utf8	
Spanish	es_ES.utf8	

LANG environment variable

If the LANG environment variable is set to a locale ID listed in [Table 1 on page 17](#), messages from the PowerPath installation program are displayed in the language corresponding to that locale ID. If you want to display PowerPath installation messages in a language that is not the default for the host, set the LANG variable for the terminal session to the locale ID for that language. For example, to display installation messages in simplified Chinese, enter one of the following commands (depending on the shell being used):

```
export LANG=zh_CN.utf8
```

or


```
setenv LANG zh_CN.utf8
```

Note: If you set the LANG variable to an unsupported locale ID, PowerPath continues the installation in English.

gettext utility

The GNU **gettext** utilities must be installed on the host if you intend to use a language other than English. If the gettext command is present in the /usr/bin/gettext directory, the gettext utilities are installed on the host.

Note: If the gettext utilities are not installed on the host, PowerPath returns a warning and continues the installation in English.

IA64 systems only

If you are going to install PowerPath on an IA64 system, then you must install the following packages on that system before installing PowerPath.

Distribution	Packages
RHEL 5	glibc-<xxx>.i686.rpm ^a ia32el-<xxx>.ia64.rpm
SLES 10 SP4	glibc-x86-2.4-31.74.1.ia64.rpm ia32el-7042_7022-0.5.13.ia64.rpm
SLES 11 SP1	glibc-x86-2.11.1-0.17.4 ia32el-7042_7030-0.2.24

- a. Where <xxx> is the version of the package available in RHEL 5.

Note: The kernels listed above are those of or pertaining to the Linux 5.5 and its minor releases. To view the list of kernels pertaining to earlier releases, log in to Powerlink website.

Installing PowerPath

This section describes how to install PowerPath on a Linux host on which PowerPath is not currently installed. ["Upgrading PowerPath" on page 33](#) contains instructions on upgrading PowerPath on a host that has an earlier version of PowerPath.

Take note of the following before installing PowerPath:

- ◆ You do not need to reboot the host after installing PowerPath.
- ◆ ["Recovering from a failed PowerPath installation" on page 90](#) contains information on recovering from a failed PowerPath installation.

Note:

—To install PowerPath in a Xen environment, refer to ["Installing PowerPath in a Xen environment" on page 23](#) before completing this procedure.

—To install PowerPath if Veritas Storage Foundation is already installed, refer to ["Installing PowerPath if Veritas Storage Foundation is already installed" on page 26](#).

To install PowerPath software:

1. Log in as root.
2. Copy the RPM package into a temporary folder and type the following:

```
rpm -ivh EMCpower.LINUX-5.5.0.01.00-xxx.yyyy.zzzz.rpm
```

where:

xxx = build number

yyyy = platform, such as RHEL 5, SLES10, or SLES 11 SP1

zzzz = i386, ia64, or x86_64

Refer to ["Package names" on page 22](#).

Note: Support for Linux kernels is limited for IBM PowerPC hosts. Refer to the Release Notes or to the E-Lab Interoperability Navigator to verify kernel support.

Installing from a CD

To install from a CD:

1. Insert the PowerPath installation CD into the CD-ROM drive.
2. Create the directory `/cdrom` to be the mount point for the CD-ROM. Type:

```
mkdir /cdrom
```

3. Mount the PowerPath CD on `/cdrom`. Type:

```
mount -o ro /dev/cdrom /cdrom
```

4. Change directories to the directory for your operating system.

- For SLES, type:

```
cd /cdrom/LINUX/2.6/pp5.5.x/sles
```

- For RHEL and Asianux, type:

```
cd /cdrom/LINUX/2.6/pp5.5.x/rhel
```

5. Proceed to ["Installing PowerPath" on page 20](#).

Installing from a compressed archive

To install from a compressed archive on Powerlink:

1. Download the PowerPath archive from Powerlink. From the **Support** menu, select **Software Downloads and Licensing > Downloads P-R > PowerPath for Linux**.
2. Untar the PowerPath archive. Type:

```
tar -xzf EMCpower.LINUX-5.5.x.xx.xx-<build>.tar.gz
```

3. Proceed to ["Installing PowerPath" on page 20](#).

Table 2 Package names (1 of 2)

Package	Description
RHEL5.i386	PowerPath for RHEL5, OEL5.x (x=1,2) and Asianux 3 i386 platforms.
RHEL5.ia64	PowerPath for RHEL5, OEL5 and Asianux 3 IA64 platforms.
RHEL5.x86_64	PowerPath for RHEL5, OEL5 and Asianux 3 x86_64 platforms.
RHEL5.ppc64	PowerPath for RHEL5, OEL5 and Asianux 3 ppc64 platforms.

Table 2 Package names (2 of 2)

Package	Description
SLES 10sp3.i386	PowerPath for SLES 10 SP3 i386 platform.
SLES 10sp3.ia64	PowerPath on SLES 10 SP3 IA64 platforms.
SLES 10sp3.x86_64	PowerPath for SLES 10 SP3 x86_64 platforms.
SLES 10sp3.ppc64	PowerPath on SLES 10 SP4 ppc64 platforms.
SLES 10sp4.i386	PowerPath for SLES 10 SP4 i386 platform.
SLES 10sp4.ia64	PowerPath on SLES 10 SP4 IA64 platforms.
SLES 10sp4.x86_64	PowerPath for SLES 10 SP3 x86_64 platforms.
SLES 11sp1.i386	PowerPath for SLES 11 SP1 i386 platform.
SLES 11sp1.ia64	PowerPath on SLES 11 SP1 IA64 platforms.
SLES 11sp1.x86_64	PowerPath for SLES 11 SP1 x86_64 platforms.
SLES 11sp1.ppc64	PowerPath on SLES 11 SP1 ppc64 platforms.

Installing PowerPath in a Xen environment

To install PowerPath in a Xen environment:

Note: Xen is supported with PowerPath 5.5 on RHEL 5, SLES 10, and SLES 11 SP1. Xen is not supported on RHEL 5 IA64 kernels. Refer to the release notes or to the E-Lab Interoperability Navigator to verify kernel support.

1. Verify if the following Xen tools and kernel RPMs are installed on the host. If not present, you must install them before installing PowerPath.

RHEL 5

```
xen-libs-3.0.3-25.e15
xen-3.0.3-25.e15
kernel-xen-2.6.18-8.e15
bridge-utils-1.1-2
python-2.4.3-19.e15
```

RHEL 5.1

```
xen-libs-3.0.3-41.e15
xen-3.0.3-41.e15
kernel-xen-2.6.18-53.e15
```

```
bridge-utils-1.1-2
python-2.4.3-19.el5
```

RHEL 5.2

```
xen-libs-3.0.3-64.el5
xen-3.0.3-64.el5
kernel-xen-2.6.18-92.el5
bridge-utils-1.1-2
python-2.4.3-21.el5
```

RHEL 5.3

```
xen-libs-3.0.3-80.el5
xen-3.0.3-80.el5
kernel-xen-2.6.18-128.el5
bridge-utils-1.1-2
python-2.4.3-24.el5
```

RHEL 5.4

```
xen-libs-3.0.3-94.el5
xen-3.0.3-94.el5
kernel-xen-2.6.18-164.el5
bridge-utils-1.1-2
python-2.4.3-27.el5
```

RHEL 5.5

```
xen-libs-3.0.3-105.el5
xen-3.0.3-105.el5
kernel-xen-2.6.18-194.el5
bridge-utils-1.1-2
python-2.4.3-27.el5
```

SLES 10 SP3

```
xen-3.2.3_17040_16-0.4.9
xen-doc-pdf-3.2.3_17040_16-0.4.9
xen-kmp-smp-3.2.3_17040_16_2.6.16.60_0.54.5-0.4.9
or
xen-kmp-bigsmpt-3.2.3_17040_16_2.6.16.60_0.54.5-0.4.9
xen-libs-32bit-3.2.3_17040_16-0.4.9
xen-libs-3.2.3_17040_16-0.4.9
xen-tools-ioemu-3.2.3_17040_16-0.4.9
xen-tools-3.2.3_17040_16-0.4.9
```

SLES 10 SP4

```
xen-3.2.3_17040_16-0.4.9
xen-doc-pdf-3.2.3_17040_16-0.4.9
xen-kmp-smp-3.2.3_17040_16_2.6.16.60_0.85.1
or
xen-kmp-bigsmpt-3.2.3_17040_16_2.6.16.60_0.85.1
xen-libs-32bit-3.2.3_17040_16-0.4.9
```

```
xen-libs-3.2.3_17040_16-0.4.9
xen-tools-ioemu-3.2.3_17040_16-0.4.9
xen-tools-3.2.3_17040_16-0.4.9
```

SLES 11 SP1

```
bridge-utils-1.4-23.16
openssh-5.1p1-41.31.36
kernel-xen-2.6.32.12-0.7.1
kernel-xen-base-2.6.32.12-0.7.1
sles-xen_en-pdf-11.1-16.19.1
virt-manager-0.8.4-1.4.11
virt-viewer-0.2.0-1.5.65
vm-install-0.4.25-1.3.3
xen-4.0.0_21091_04-0.2.6
xen-doc-html-4.0.0_21091_04-0.2.6
xen-doc-pdf-4.0.0_21091_04-0.2.6
xen-libs-4.0.0_21091_04-0.2.6
xen-tools-4.0.0_21091_04-0.2.6
```

2. If any packages are missing, install them.

For example, use YaST to install the required Xen rpms for SLES 10.

3. Restart the host to the Xen kernel.
4. Type **uname -a** to verify your system is running the Xen kernel.

Output similar to the following appears:

```
Linux lclld0076 2.6.16.21-0.15-xenpae #1 SMP Tue Jul 25 15:28:49 UTC 2006 i686 i686
i386 GNU/Linux
```

5. Install PowerPath. Follow the procedure in "[Installing PowerPath](#)" on page 20.
6. Create a Xen virtual machine. For more information, refer to the vendor documentation for your operating system.
7. Select a pseudo device (for example, /dev/emcpowerd) and open the virtual machine configuration file. Type:

```
vi /etc/xen/vm/<vm_name>
```

where *<vm_name>* is the name of the virtual machine you created in step 4.

8. Locate the following line in the configuration file:

```
disk = [ 'file:/images/lcle055/hda.img,hda,w' ]
```

9. Add an entry for the new disk, separated by a comma.

For example, add the following line:

```
disk = [ 'file:/images/lc1e055/had.img,hda,w', 'phy:</dev/emcpowern>,<hdb>,w' ]
```

where:

- `</dev/emcpowern>` is the actual device being exported to the virtual machine. Because this is a physical device (not an image file device), **phy:** needs to be included.
- `<hdb>` is the name the virtual machine uses to identify the device.
- **w** is read-write mode

10. Type **xm create -c <vm_name>** to boot the virtual machine.

where `<vm_name>` is the name of the virtual machine you created in step 4.

Note: If the VM configuration file resides in the `/etc/xen/vm/` directory, you do not need to include the full path.

11. Verify the new device (`/dev/hdb`) is registered with the operating system by looking at the `cat /proc/partitions` file.

major	minor	#blocks	name
3	0	7168000	hda
3	1	6819561	hda1
3	2	345397	hda2
3	6	919680	hdb

12. As necessary, create partitions on the newly added device using standard procedures.

Installing PowerPath if Veritas Storage Foundation is already installed

Pseudo devices are not recognized by VxVM if PowerPath is installed after Veritas Storage Foundation.

To install PowerPath if Veritas Storage Foundation is already installed:

- ◆ For Veritas Storage Foundation 5.0 MP2 and earlier:
 1. Install PowerPath as described in "Installing PowerPath" on page 20.
 2. Restart the host.

3. Type `rm -f /etc/vx/disk.info*` to remove disk.info files.
4. Run `vxconfigd -k`.
- ◆ For Veritas Storage Foundation 5.0 MP3 and later:
 1. Type the following to exclude EMC arrays from DMP management:
 - For Symmetrix arrays:
`vxddladm excludearray libvxemc.so`
 - For CLARiiON arrays:
`vxddladm excludearray libvxCLARiiON.so`
 2. Run `vxctl enable`.
Verify enclosure EMC0 is **disconnected** and all EMC LUNs are on disk enclosure EMC0.
 3. Install PowerPath as described in ["Installing PowerPath" on page 20](#).
 4. Run `vxctl enable` so PowerPath claims the devices.
Verify enclosure EMC0 has moved to **pp_emc**.
 5. Type the following to include EMC arrays. This allows DMP to coexist with PowerPath:
 - For Symmetrix arrays:
`vxddladm includearray libvxemc.so`
 - For CLARiiON arrays:
`vxddladm includearray libvxCLARiiON.so`

Registering PowerPath

Take note of the following regarding PowerPath registration:

- ◆ Installing PowerPath on a host connected to a CLARiiON AX-series array does not require a license. PowerPath provides full functionality with or without a PowerPath license when the host is connected to a CLARiiON AX-series array. Therefore, you do not need to run the `emcpreg -install` command when installing PowerPath on a host connected to a CLARiiON AX-series array.

- ◆ If you do not register the PowerPath software during PowerPath installation (on a host connected to any storage system other than a CLARiiON AX-series array), and then you run **powermt config** or reboot the host, the load balancing and failover policy is set to Basic Failover. If this happens, you must register PowerPath using the **emcprep -install** command and then run the **powermt set policy** command to reset the policy as appropriate. The *EMC PowerPath Family CLI and System Messages Reference Guide* contains more information on powermt commands.

To register the PowerPath license:

1. Run **emcprep -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**.

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. Enter the 24-character alphanumeric sequence found on the License Key Card delivered with the PowerPath media kit, and press **Enter**.

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

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

If you enter an invalid registration key, the screen displays an error message and prompts you to enter a valid key. The *EMC PowerPath Family CLI and System Messages Reference Guide* contains 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.
```

5. Type **/etc/init.d/PowerPath start** if PowerPath service is not started.

Note:

—Only use the above script to load and unload the PowerPath modules.
—If the LANG variable is set to the locale ID for a language listed in [Table 1 on page 17](#), messages from the PowerPath start command are displayed in that language.

6. Run the PowerPath **powermt set policy** command to reset the policy as appropriate. The *PowerPath Family CLI and System Messages Reference Guide* contains more information on powermt commands.
7. Type **powermt save** to ensure PowerPath configuration settings persist across host restarts.

After installing PowerPath

After installing PowerPath:

- Configure the management daemon.
- Verify that your PowerPath capabilities are correct.
- Reconfigure device mappings.
- If you load the HBA driver as a module, verify that all PowerPath extensions are loaded.

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 68](#) describes how to configure the management daemon.

Verify your PowerPath capabilities are correct

Verify that your PowerPath capabilities match your license. For example, run **powermt display dev=all**, check the output and take any corrective actions to ensure the load balancing and failover policies are correct.

Note: If the load balancing policy does not match your expectations for your license level, check the license registration.

The *EMC PowerPath Family CLI and System Messages Reference Guide* and the `powermt(1)` man page contains more information about the `powermt display` command, and the load balancing and failover policies.

Reconfiguring device mappings

Run **powermt display dev=all** to view current device mappings. When reconfiguring the device mappings on an array, you should pay careful attention to the resulting configuration. Device names on

the host are assigned dynamically during the loading of the HBA driver. Therefore, any changes to the configuration may result in changes in the pre-existing device naming association. For example, this may render some existing mount tables inaccurate if you do not update the mount points to correspond to the new device configuration and its device naming association.

Note: Refer to the `emcpadm` command to change device mappings.

Verify extensions are loaded

If you load the HBA driver as a module (not as a static kernel driver), verify that all extensions are loaded by running `lsmod`.

Provided you built the HBA driver as a module, the HBA driver name appears in the `lsmod` output, below the PowerPath drivers (which begin with `emc`).

For example:

Module	Size	Used by
<code>emcpvlund</code>	35784	0
<code>emcpxcrypt</code>	147656	0
<code>emcpdm</code>	54720	0
<code>emcpgpx</code>	37328	3 <code>emcpvlund, emcpxcrypt, emcpdm</code>
<code>emcpmpx</code>	174504	108
<code>emcp</code>	2442048	6 <code>emcpvlund, emcpxcrypt, emcpdm, emcpgpx, emcpmpx</code>
<code>vfat</code>	34048	0
<code>fat</code>	68512	1 <code>vfat</code>
<code>ext3</code>	149872	0
<code>jbd</code>	95048	1 <code>ext3</code>
<code>usbserial</code>	52472	0
<code>parport_pc</code>	58432	0
<code>lp</code>	29352	0
<code>parport</code>	65292	1 <code>parport_pc, lp</code>
<code>autofs</code>	36736	6
<code>edd</code>	28448	0
<code>floppy</code>	84432	0
<code>freq_table</code>	22528	0
<code>thermal</code>	32012	0
<code>processor</code>	39008	1 <code>thermal</code>
<code>fan</code>	21128	0
<code>button</code>	24096	0
<code>battery</code>	26632	0
<code>ac</code>	22280	0
<code>ipv6</code>	369784	21
<code>joydev</code>	28032	0

```

sg                61112      0
st                60708      0
sr_mod           35748      0
ohci_hcd         39428      0
usbcore         148728      4  usbserial,ohci_hcd
hw_random       22952      0
tpm_atmel       23936      0
tpm             29600      1  tpm_atmel
evdev           27520      0
tg3             137092     0
subfs           25984      2
raw             89760      0
dm_mod          83552      0
dump_blockdev   30204      0
dump_gzip       20648      0
zlib_deflate    40472      1  dump_gzip
dump            82432      2  dump_blockdev,dump_gzip
reiserfs       290512      1
lpfc           256864     108
scsi_transport_fc 24832      1  lpfc
mptspi         27916      4
mptscsih       63492      1  mptspi
mptbase        85984      2  mptspi,mptscsih
sd_mod         38912     113
scsi_mod       162440      9
emcp,sg,st,sr_mod,lpfc,scsi_transport_fc,mptspi,mptscsih,sd_mod

```

The `/etc/modprobe.conf` file should contain the following text:

```

###BEGINPP
include /etc/modprobe.conf.pp
###ENDPP

```

The `/etc/modprobe.conf.pp` file should contain the following text:

```

options emcp
    managedclass=symm,clariion,hitachi,hpxp,ess,invista
options emcpsf SfUnmanagedClasses=hphsx

```

Edit the `boot.vxvm` file

If you have installed PowerPath on a SLES 10 host where you plan to install VxVM follow these steps after installing VxVM:

Note: The following steps are specific to VxVM 4.1.

1. Open the `boot.vxvm` startup script.
2. Edit the file by adding the following to the **Required Start** line:

boot.powerpath

3. Run **insserv**.

Editing this file prevents the VxVM boot script from starting before the PowerPath boot script, in which case VxVM would not recognize PowerPath pseudo devices.

Upgrading PowerPath

This section describes how to upgrade PowerPath.

Before you upgrade PowerPath

Before upgrading PowerPath:

- ◆ Download the latest version of EMC Grab utilities, available on Powerlink, and then run the PowerPath Configuration Checker (PPCC).
- ◆ You can only upgrade to PowerPath 5.5 from PowerPath 4.5 or later.
- ◆ If the host was not installed with the default file system layout provided by the operating system vendor, then PowerPath may report "no such file or directory" during system initialization and the Linux host may not boot with PowerPath installed. EMC supports only the default filesystem layout (s) as supplied by the operating system vendor.
- ◆ Check the Powerlink website for the most current information.
- ◆ The PowerPath for Linux upgrade script saves the current configuration settings before beginning the upgrade. Any previous settings in the `powermt.custom` file will be overwritten.
- ◆ Back up the `/etc/powermt.custom` and `/etc/emcp_registration` files.
- ◆ Document your existing powerdevice mappings, which can be done by saving the output of `powermt display dev=all`.
- ◆ Unmount any file systems mounted on PowerPath devices and deactivate any volume groups using these file systems.
- ◆ Stop `powermt display` if it is running.
- ◆ Do not use the `-i` option to upgrade from an earlier version of PowerPath. Using the `-i` option creates multiple PowerPath entries in the Linux RPM database, which can cause ongoing maintenance problems. If you inadvertently use the `-i` option to upgrade PowerPath, follow the procedure in "[Multiple PowerPath entries in the RPM database](#)" on page 93 to correct the problem.

- ◆ Localized versions of the PowerPath installer are available in Brazilian Portuguese, French, German, Italian, Korean, Japanese, Spanish, and simplified Chinese. Refer to ["Select language for installation" on page 17](#) for information on using a localized version of the PowerPath installer when upgrading PowerPath.
- ◆ **CLARiiON systems only:** Stop the Navisphere agent.
- ◆ Refer to ["Troubleshooting a PowerPath upgrade" on page 91](#) if you encounter any problems while upgrading PowerPath.

Upgrading from a CD

To upgrade from a CD:

1. Insert the PowerPath installation CD into the CD-ROM drive.
2. Create the directory /cdrom as the mount point. Type:


```
mkdir /cdrom
```
3. Mount the PowerPath CD on /cdrom. Type:


```
mount -o ro /dev/cdrom /cdrom
```
4. Change directories to the directory for your operating system.
 - For RHEL type:


```
cd /cdrom/LINUX/2.6/pp5.5.x/rhel
```
 - For SLES type:


```
cd /cdrom/LINUX/2.6/pp5.5.x/sles
```
5. Proceed to ["Upgrading PowerPath" on page 35](#).

Upgrading from a compressed archive

To upgrade from a compressed archive on Powerlink:

1. Download the PowerPath archive from Powerlink. From the **Support** menu, select **Software Downloads and Licensing > Downloads P-R > PowerPath for Linux**.
2. Untar the PowerPath archive. Type:


```
tar -xzf EMCpower.LINUX-5.5.x.xx.xx-<build>.tar.gz
```
3. Proceed to ["Upgrading PowerPath" on page 35](#).

Upgrading PowerPath

To upgrade PowerPath:

1. Log in as root.

- To upgrade to the new PowerPath version, type:

```
rpm -Uhv EMCpower.LINUX-5.5.0-xxx.yyyy.zzzz.rpm
```

where:

xxx = build number

yyyy = platform, such as RHEL5 or SLES10

zzzz = i386, ia64, or x86_64

Refer to "Package names" on page 22.

where *<package_name>* is one of the package names shown in Table 2 on page 22.

- Run `/etc/init.d/PowerPath start` to start PowerPath.
- If the PowerPath configuration includes any custom policies or settings, run `powermt load` to load these settings.

Upgrading the Linux operating system

This section describes how to upgrade PowerPath across different versions of Linux.

Note: Before upgrading to a new version of Linux, check the E-Lab Interoperability Navigator to verify that version of Linux has been qualified with PowerPath.

There are three sets of procedures:

- ◆ Upgrading PowerPath before upgrading SLES Service pack or RHEL Update on Linux
- ◆ Upgrading PowerPath after upgrading RHEL Update on Linux
- ◆ Upgrading PowerPath after upgrading from major Linux release to next major Linux release

Upgrading PowerPath before upgrading SLES Service pack or RHEL Update on Linux

To upgrade PowerPath before upgrading Linux (SLES 10 SP2 to SLES 10 SP4, RHEL 5 U4 to RHEL 5 U5), perform the following steps:

- Upgrade PowerPath to the latest version, as described in "Installing from a CD" on page 21.

Note: Do not start PowerPath service at this time.

2. Upgrade the operating system to a PowerPath supported version as described in the RedHat and Novell documentation.
3. Restart the host. This is the last step on SLES hosts.
4. On RHEL hosts, Run '/etc/init.d/PowerPath start' to start PowerPath.

Upgrading PowerPath after upgrading RHEL Update on Linux

To upgrade PowerPath after upgrading Linux (RHEL5 U4 to RHEL 5 U5), perform the following steps:

Note: Upgrading PowerPath after upgrading Linux is not supported on SLES hosts.

1. Upgrade the operating system to the latest version PowerPath supports, as described in the Red Hat documentation.
2. Restart the host.
3. Upgrade PowerPath to the latest version, as described in ["Upgrading from a CD" on page 34](#).
4. Run '/etc/init.d/PowerPath start' to start PowerPath.

Upgrading PowerPath from one major release to next major release

To upgrading PowerPath after upgrading from major Linux release to next major Linux release (RHEL4 U8 to RHEL5 U5, SLES 9 SP4 to SLES 10 SP4, SLES 10 SP2/SP4 to SLES 11 SP1), perform the following steps:

1. Rename: /etc/init.d/PowerPath to /etc/init.d/PowerPath.old.
2. Upgrade the operating system to a PowerPath supported version as described in the RedHat and Novell documentation.
3. Upgrade PowerPath to the latest version.
4. Run: /etc/init.d/PowerPath start to start PowerPath.
5. Delete: /etc/init.d/PowerPath.old.

Upgrading PowerPath in a SteelEye LifeKeeper cluster environment

To upgrade PowerPath in a SteelEye LifeKeeper cluster environment:

Note: Ensure that the node being upgraded has no active resources or hierarchies

1. Before upgrading PowerPath, type **lkbackup -c --cluster** to backup the LifeKeeper configuration.

Note: Ensure that you have a good backup of the configuration that is stored in a safe location (for example, copy to a system not involved with the cluster).

2. Run **lkstop -r** on the node you are upgrading to stop LifeKeeper.
3. Uninstall PowerPath.
 - a. Type **rpm -qa | grep -i emcpower**
 - b. Type **rpm -ev --nodeps <package_name>**.
4. Install PowerPath as described in ["Installing PowerPath" on page 20](#).
5. Restart the host.
6. Verify the PowerPath configuration:
 - a. Run **powermt display**.
 - b. Run **powermt display dev=all**.
7. Run **lkstart** to start SteelEye LifeKeeper.
8. Repeat steps 1-7 for each node in the cluster.

Note: EMC recommends manually switching over the resources to the previously upgraded node to ensure that the resource hierarchies are working properly on the upgraded node. Manually switching over minimizes any impact to running applications as opposed to using automatic switchover.

Upgrading PowerPath Encryption

To upgrade PowerPath Encryption, refer to "Upgrading PowerPath software" in the *PowerPath 5.5 Encryption with RSA User Guide*.

Coexistence with third-party path management software

PowerPath can coexist with MPIO-based third-party path management software such as:

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

However, PowerPath cannot co-manage devices with third-party path management software.

When PowerPath coexists with other multipathing software, the HBAs must be dedicated to either PowerPath or the third-party software; HBAs cannot be used for both packages.

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

PowerPath behavior

PowerPath co-management is not supported by third-party array path management products; therefore, when you enable PowerPath for a class that is being managed by a third-party product, multipathing behavior is undefined.

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>**.

If this software is present or enabled	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 co-manage devices, do not initiate co-management 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 reboots.

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

Before you install or enable 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

This chapter includes the following sections:

- ◆ Introduction 42
- ◆ Configuring PowerPath in a boot-from-SAN setup (SLES 10 SP4). 43
- ◆ Configuring PowerPath in a boot-from-SAN setup (SLES11 SP1) .. 45
- ◆ Configuring PowerPath in a boot-from-SAN setup (RHEL 5 and updates)47
- ◆ Upgrading the Linux kernel in a boot-from-SAN setup 50
- ◆ Upgrading the Linux kernel in a boot-from-SAN setup (SLES11SP1)51
- ◆ Upgrading PowerPath in a boot-from-SAN setup..... 53
- ◆ Removing PowerPath in a boot-from-SAN setup..... 54
- ◆ Removing PowerPath in a boot-from-SAN setup (SLES11SP1)... 55

Introduction

On Symmetrix and CLARiiON storage systems, you can use a PowerPath pseudo (emcpower) device located on external storage as a root device (the device that contains the startup image). To use a PowerPath pseudo device as the root device, the device must be under LVM control. Once the PowerPath drivers have been loaded, using a PowerPath pseudo device as the root device provides load balancing and path failover for the root device.

The *EMC PowerPath Family for Linux Release Notes* contains the supported root device configurations.

Root device considerations

Note the following before configuring a root device:

- ◆ With PowerPath, EMC recommends the following:
 - Use a PowerPath pseudo (emcpower) device instead of the native sd device in the **/etc/fstab** file.
 - For RHEL 5, use the uuid in the **/etc/fstab** file to mount the /boot file system.
 - For SLES 10 and SLES 11, use a PowerPath pseudo (emcpower) device instead of the default /dev/disk/by-id device in the **/etc/fstab** file.



IMPORTANT

To use a PowerPath pseudo device as a root device, the device must be under LVM control.

- ◆ When booting from the SAN, duplicate ID messages may appear if you run any LVM-related commands. To avoid seeing duplicate messages, filter out all sd devices using the **/etc/lvm/lvm.conf** file. Refer to "[Configuring LVM2 support](#)" on [page 72](#) for information on modifying **/etc/lvm/lvm.conf**.vice

Configuring PowerPath in a boot-from-SAN setup (SLES 10 SP4)

To configure a PowerPath root device using LVM on a SLES 10 and service packs only host:

1. Create a custom partition setup on the target storage device for two partitions:
 - The /boot partition formatted as either an **ext3** or **reiserfs** file system and at least 100 MB in size.
 - A second partition of type **0x8E Linux LVM** using the remaining space on the storage device.

Note: For PowerPC, follow partition recommendations for creating Prep partitions. For IA64, create /boot/efi partition using vfat filesystem with EFI boot label and at least 300 MB in size.

2. Install SLES on the host. Configure a single active path to the boot LUN during the initial installation. You attach additional LUNs and configure additional paths at the end of this procedure.
3. Create a volume group for the LVM partition.
4. In the volume group, create a separate volume for each file system. Format each volume as swap, ext3, or reiserfs and designate the desired mount point.
5. Change the default boot loader from **LILO** to **GRUB**. Ensure that the boot loader location specifies the MBR of the desired device rather than the boot sector of the /boot partition.

Note: This step is not applicable for IBM PowerPC hosts and IA64 hosts.

6. Upgrade the kernel revision to a revision compatible with PowerPath. The *EMC PowerPath Family for Linux Release Notes* contains the current PowerPath requirements.
7. Install and configure PowerPath following the instructions in [Chapter 1, "Installing PowerPath on a Linux host,"](#) and the *EMC Host Connectivity with QLogic Fibre Channel and iSCSI Host Bus Adapters (HBAs) and Converged Network Adapters (CNAs) for the Linux environment* guide.

8. Edit the `/etc/fstab` file to mount the `/boot` partition on a PowerPath pseudo device instead of the default `/dev/disk/by-id` device in the `/etc/fstab` file. Consider the following host where `emcpowera` is the corresponding PowerPath pseudo name for the default `/dev/disk/by-id` device.

The original `/etc/fstab` file will be similar to the following:

Note: This step is not applicable for IBM PowerPC hosts.

```

/dev/system/root      /                reiserfs  acl,user_xattr      1  1
/dev/disk/by-id/scsi-360060160909a1d00eafe889a2941dc11-part1
  /boot              ext2           acl,user_xattr      1  2
/dev/system/swap     swap           swap        defaults             0  0
proc                 /proc         proc        defaults             0  0
sysfs                /sys          sysfs       noauto              0  0
debugfs              /sys/kernel/debug  debugfs     noauto              0  0
usbfs                /proc/bus/usb  usbfs       noauto              0  0
devpts               /dev/pts      devpts      mode=0620,gid=5     0  0
/dev/fd0             /media/floppy  auto        noauto,user, sync   0  0
  
```

Edit the `/etc/fstab` file so that the `/boot` partition is mounted by a PowerPath pseudo device. For example, `/dev/emcpoweru1`. The modified `/etc/fstab` file is shown below.

```

/dev/system/root      /                reiserfs  acl,user_xattr      1  1
/dev/emcpoweru1      /boot           ext2           acl,user_xattr      1  2
/dev/system/swap     swap           swap        defaults             0  0
proc                 /proc         proc        defaults             0  0
sysfs                /sys          sysfs       noauto              0  0
debugfs              /sys/kernel/debug  debugfs     noauto              0  0
usbfs                /proc/bus/usb  usbfs       noauto              0  0
devpts               /dev/pts      devpts      mode=0620,gid=5     0  0
/dev/fd0             /media/floppy  auto        noauto,user, sync   0  0
  
```

9. Configure additional paths to the storage devices. Attach additional LUNs to the host.

Configuring PowerPath in a boot-from-SAN setup (SLES11 SP1)

To configure a PowerPath boot-from-SAN device using the LVM on an SLES 11 and service packs only host:

1. Create a custom partition setup on the target storage device for two partitions:
 - The /boot partition formatted as an ext3 file system and at least 100 MB in size.
 - A second partition of type **0x8E Linux LVM** using the remaining space on the storage device.

Note: For PowerPC, follow partition recommendations for creating Prep partitions. For IA64, create **/boot/efi** partition using vfat filesystem with EFI boot label and at least 300 MB in size.

2. Install SLES on the host. Configure a single active path to the boot LUN during the initial installation. You attach additional LUNs and configure additional paths at the end of this procedure.
3. Create a volume group for the LVM partition.
4. In the volume group, create a separate volume for each file system. Format each volume as swap or ext3, and designate the desired mount point to each partition.
5. Ensure **GRUB** is the boot loader selected and that the boot loader location specifies the MBR of the desired device rather than the boot sector of the /boot partition.

Note: This step is not applicable for IBM PowerPC hosts and IA64 hosts.

6. Install and configure PowerPath following the instructions in [Chapter 1, "Installing PowerPath on a Linux host,"](#) and the *EMC Host Connectivity with QLogic Fibre Channel and iSCSI Host Bus Adapters (HBAs) and Converged Network Adapters (CNAs) for the Linux environment* guide.
7. Add additional paths to the boot LUN and configure them in PowerPath.
8. Edit the **/etc/fstab** file to mount the /boot partition on a PowerPath pseudo device instead of the default **/dev/disk/by-id** or **/dev/disk/by-uuid** device in the **/etc/fstab** file.

The original `/etc/fstab` file will be similar to the following:

```
/dev/sles11/swap      swap                swap                defaults            0 0
/dev/sles11/root      /                  ext3                acl,user_xattr      1 1
/dev/disk/by-id/scsi-3600601605c7019002820bd32b767de11-part1
                        /boot              ext3                acl,user_xattr      1 2
proc                  /proc              proc                defaults            0 0
sysfs                 /sys               sysfs               noauto              0 0
debugfs              /sys/kernel/debug debugfs             noauto              0 0
usbfs                 /proc/bus/usb      usbfs               noauto              0 0
devpts                /dev/pts           devpts              mode=0620,gid=5     0 0
```

Edit the `/etc/fstab` file so that the `/boot` partition is mounted by a PowerPath pseudo device. For example, `/dev/emcpowera1`. The modified `/etc/fstab` file is shown below.

```
/dev/sles11/swap      swap                swap                defaults            0 0
/dev/sles11/root      /                  ext3                acl,user_xattr      1 1
# /dev/disk/by-id/scsi-3600601605c7019002820bd32b767de11-part1
                        /boot              ext3                acl,user_xattr      1 2
/dev/emcpowera1      /boot              ext3                acl,user_xattr      1 2
proc                  /proc              proc                defaults            0 0
sysfs                 /sys               sysfs               noauto              0 0
debugfs              /sys/kernel/debug debugfs             noauto              0 0
usbfs                 /proc/bus/usb      usbfs               noauto              0 0
devpts                /dev/pts           devpts              mode=0620,gid=5     0 0
```

If necessary, save a backup of the `/etc/fstab` file.

Note: For IA64, unmount `/boot/efi` from the native device and run `mount -a` to ensure `/boot/efi` is mounted on PowerPath Pseudo device. For example, `/dev/emcpowera1`.

9. Edit the `/etc/lvm/lvm.conf` file to filter out SCSI device nodes from its internal cache so that LVM2 recognizes a PowerPath pseudo device as the single path to a LUN. Modify the filter field in the `/etc/lvm/lcm.conf` file with the appropriate line for the SLES 11 kernel, as follows:

```
# By default we accept every block device except udev names:
# filter = [ "r|/dev./*/by-path/*|", "r|/dev./*/by-id/*|", "a./*/" ]
filter = [ "r|/dev./*/by-path/*|", "r|/dev./*/by-id/*|",
           "r|/dev/sd.*|", "a./*/" ]
```

If necessary, save a backup of the `/etc/lvm/lvm.conf` file.

10. Run `vgscan -v` to rebuild the LVM2 cache.

11. Run **lvmdiskscan** to ensure that the filter field is working correctly. Verify that the **filtered** device nodes are not listed in the output.
12. Create a new **initrd** image.

```
mkinitrd -i /boot/initrd-pp-2.6.27.19-5-default -k
/boot/vmlinuz-2.6.27.19-5-default -f powerpath
```

13. Add a new boot entry for the new **initrd** image created in step 12 in **/boot/grub/menu.lst**.

Note: For IA64, use **/etc/elilo.conf** and run **/sbin/elilo**.

14. Attach additional LUNs to the host and configure them in PowerPath.
15. Reboot the host so that the new **initrd** image that has PowerPath modules takes effect.

Configuring PowerPath in a boot-from-SAN setup (RHEL 5 and updates)

To configure a PowerPath root device using the LVM on a RHEL host:

1. Create a custom partition setup on the target storage device for two partitions:
 - The **/boot** partition formatted as an ext3 file system and at least 100 MB in size.
 - A second partition of type **0x8E Linux LVM** using the remaining space on the storage device.

Note: For PowerPC, follow partition recommendations for creating Prep partitions. For IA64, create **/boot/efi** partition using vfat filesystem with EFI boot label and at least 300 MB in size.

2. Install RHEL on the host. Configure a single active path to the boot LUN during the initial installation. You attach additional LUNs and configure additional paths at the end of this procedure.
3. Create a volume group for the LVM partition.
4. In the volume group, create a separate volume for each file system. Format each volume as swap or ext3, and designate the desired mount point to each partition.

5. Ensure **GRUB** is the boot loader selected and that the boot loader location specifies the MBR of the desired device rather than the boot sector of the /boot partition.

Note: This step is not applicable for IBM PowerPC hosts and IA64 hosts.

6. Install and configure PowerPath following the instructions in [Chapter 1, "Installing PowerPath on a Linux host,"](#) and the *EMC Host Connectivity with QLogic Fibre Channel and iSCSI Host Bus Adapters (HBAs) and Converged Network Adapters (CNAs) for the Linux environment* guide.
7. If you are working on RHEL 5 or IBM PowerPC hosts, you will need to edit **/etc/fstab** to mount /boot by UUID. By default, the /boot partition is configured to mount by LABEL. The following steps show an example on configuring /etc/fstab to mount /boot by UUID.
 - a. Identify the partition hosting /boot partition:

```
# df -kh /boot
Filesystem                Size      Used Avail Use% Mounted on
/dev/emcpower1            99M       22M   73M  23% /boot
# ls -l /dev/disk/by-uuid/ | grep emcpower1
lrwxrwxrwx 1 root root 16 Jun 15 20:38 a4ef34a8-4634-4a4a-bd8f-c79026b63224 ->
  ../../emcpower1
```

The UUID for the partition is:

/dev/disk/by-uuid/ a4ef34a8-4634-4a4a-bd8f-c79026b63224

- b. Edit the **/etc/fstab**, by replacing LABEL= with the UUID on the line containing /boot

```
/dev/system/LogVol100          /                ext3              defaults 1 1
/dev/disk/by-uuid/a4ef34a8-4634-4a4a-bd8f-c79026b63224 /boot ext3 defaults 1 2
#LABEL=/boot                  /boot           ext3              defaults 1 2
devpts                        /dev/pts        devpts           gid=5,mode=620 0 0
tmpfs                         /dev/shm        tmpfs            defaults 0 0
proc                          /proc           proc              defaults 0 0
sysfs                         /sys            sysfs            defaults 0 0
/dev/system/LogVol101         swap             swap              defaults 0 0
```

Note that the **mount by UUID multipath device** function does not depend on the PowerPath version installed because it is a feature that the underlying operating system either supports or does not support.

8. Configure additional paths to the storage devices. Attach additional LUNs to the host.

Note: The Boot From SAN host may not boot properly if both of the following conditions occur:

- a. Ramdisk image was created after PowerPath has been installed and configured with specific LVM filter conditions. This may happen automatically on driver or kernel updates.
- b. The native device name of the root device has changed because of configuration (fabric/array) changes.

On boot, the ramdisk will be loaded, but the root volume group will not be found; the boot cannot continue leading to a kernel panic.

9. Make a copy of the original ramdisk image found in /boot folder. Please make sure that the LVM filter in the image accepts all devices:

```
[root@libl044 ~]# cp /boot/ initrd-2.6.18-194.el5.img
/boot/initrd-2.6.18-194.el5.img.original
```

10. Create a new entry in **/boot/grub/menu.lst** file with the original ramdisk image created in the step above.

```
title Red Hat Enterprise Linux Server (2.6.18-194.el5)
  root (hd0,0)
  kernel /vmlinuz-2.6.18-194.el5 ro
  root=/dev/VolGroup00/LogVol00 rhgb quiet
  initrd /initrd-2.6.18-194.el5.img
```

```
title Red Hat Enterprise Linux Server (2.6.18-194.el5)
  root (hd0,0)
  kernel /vmlinuz-2.6.18-194.el5 ro
  root=/dev/VolGroup00/LogVol00 rhgb quiet
  initrd /initrd-2.6.18-194.el5.img.original
```

Upgrading the Linux kernel in a boot-from-SAN setup

To upgrade the Linux kernel in a boot-from-SAN setup:

1. Upgrade the kernel, following the steps provided by RedHat and Novell for upgrading the kernel in the host.
2. Before restarting the host, edit the `/etc/fstab` file to comment out entries that refer to the PowerPath pseudo (emcpower) names.

An example `/etc/fstab` file with a commented out entry for the `/boot` partition is shown below.

```
/dev/VolGroup00/LogVol00 / ext3 defaults 1 1
#/dev/emcpowera1 /boot ext3 defaults 1 2
none /dev/pts none devpts gid=5,mode=620 0 0
none /dev/shm none tmpfs defaults 0 0
none /proc none proc defaults 0 0
none /sys none sysfs defaults 0 0
/dev/VolGroup00/LogVol01 swap swap defaults 0 0
/dev/hda /media/cdrom auto auto 0 0
pamconsole,fscontext=system_u:object_r:removable_t,exec,noauto,managed 0 0
/dev/fd0 /media/floppy auto auto 0 0
pamconsole,fscontext=system_u:object_r:removable_t,exec,noauto,managed 0 0
```

3. Restart the host.

On a PowerPath 5.0.1 and later host, PowerPath detects that a new version of Linux has been installed on the host and automatically reinstalls the PowerPath drivers. You do not need to reinstall PowerPath after upgrading Linux.

4. Uncomment all entries in the `/etc/fstab` file that refer to PowerPath pseudo (emcpower) devices. A modified `/etc/fstab` file is shown below:

```
/dev/VolGroup00/LogVol00 / ext3 defaults 1 1
#/dev/emcpowera1 /boot ext3 defaults 1 2
none /dev/pts none devpts gid=5,mode=620 0 0
none /dev/shm none tmpfs defaults 0 0
none /proc none proc defaults 0 0
none /sys none sysfs defaults 0 0
/dev/VolGroup00/LogVol01 swap swap defaults 0 0
/dev/hda /media/cdrom auto auto 0 0
pamconsole,fscontext=system_u:object_r:removable_t,exec,noauto,managed 0 0
/dev/fd0 /media/floppy auto auto 0 0
pamconsole,fscontext=system_u:object_r:removable_t,exec,noauto,managed 0 0
```

5. Run `mount -a` to ensure that all emcpower partitions in the `/etc/fstab` file are mounted.

Upgrading the Linux kernel in a boot-from-SAN setup (SLES11SP1)

To upgrade the Linux kernel in a boot-from-SAN setup on a SLES 11 host:

1. Edit the `/etc/fstab` file to comment out any references to PowerPath pseudo (emcpower) devices and enable (or add) the native device entries.

```
/dev/disk/by-id/scsi-3600601605c7019002820bd32b767de11-part1
/boot                ext3                acl,user_xattr      1 2
#/dev/emcpowera1    /boot                ext3                acl,user_xattr      1 2
```

2. Revert back the LVM filter `/etc/lvm/lvm.conf` to the default setting.

```
filter = [ "r|/dev/.*by-path/.*|", "r|/dev/.*by-id/.*|", "a./.*/" ]
#filter = [ "r|/dev/.*by-path/.*|", "r|/dev/.*by-id/.*|",
            "r|/dev/sd.*|", "a./.*/" ]
```

3. Rebuild the LVM2 cache and verify the filter is working properly:
 - a. Run `vgscan -v` to rebuild the LVM2 cache.
 - b. Run `lvmdiskscan` to ensure that the filter field is working correctly.
4. Upgrade the kernel as described in Novell documentation and reboot the host.
5. Edit the `/etc/fstab` file to mount the `/boot` partition on a PowerPath pseudo device.

```
#/dev/disk/by-id/scsi-3600601605c7019002820bd32b767de11-part1
/boot                ext3                acl,user_xattr      1 2
/dev/emcpowera1    /boot                ext3                acl,user_xattr      1 2
```

6. Set the LVM filter `/etc/lvm/lvm.conf` to reject the native devices.

```
filter = [ "r|/dev/.*by-path/.*|", "r|/dev/.*by-id/.*|",
            "r|/dev/sd.*|", "a./.*/" ]
```

7. Rebuild the LVM2 cache and Verify the filter is working properly:
 - a. Run `vgscan -v` to rebuild the LVM2 cache.
 - b. Run `lvmdiskscan` to ensure that the filter field is working correctly.

8. Create the **initrd** image with PowerPath feature.

```
mkinitrd -i /boot/initrd-pp-2.6.27.23-0.1-default -k  
/boot/vmlinuz-2.6.27.23-0.1-default -f powerpath
```

9. Add a new boot entry which has PowerPath **initrd** image (initrd-pp-2.6.27.23-0.1-default) in **/boot/grub/menu.lst**.
10. Reboot the host.

Upgrading PowerPath in a boot-from-SAN setup

To upgrade PowerPath in a boot-from-SAN setup:

1. Stop the Navisphere agent and any other application that may be holding devices open.

Note: The Navisphere agent is not applicable for IBM PowerPC hosts. Therefore, stopping the Navisphere agent is not necessary for IBM PowerPC hosts.

2. Unmount any file systems mounted on PowerPath devices and deactivate any volume groups using these file systems.
3. Rename `/etc/init.d/PowerPath` to **`/etc/init.d/PowerPath.old`**.
4. Edit the `/etc/fstab` file to remove any references to PowerPath pseudo (emcpower) devices.
5. Restart the host.
6. After the host boots, type:

```
rpm -Uvh <package_name>
```

where `<package_name>` is one of the package names shown in [Table 2 on page 21](#)
7. Type `/etc/init.d/PowerPath start` to start PowerPath.

Removing PowerPath in a boot-from-SAN setup

To remove PowerPath from a boot-from-SAN setup:

1. Move the `/etc/init.d/PowerPath` script to a different location.
2. Edit the `/etc/fstab` file to remove any references to PowerPath pseudo (emcpower) devices.
3. Restart the host.
4. Remove PowerPath following the instructions in [Chapter 5, "Removing PowerPath from a Linux host."](#)

Removing PowerPath in a boot-from-SAN setup (SLES11SP1)

To remove PowerPath from a boot-from-SAN setup on SLES 11:

1. Move the `/etc/init.d/PowerPath` script to a different location.
2. Edit the `/etc/fstab` file to restore it to its original state by replacing references to emcpower device with SCSI devices.
3. Edit the `/etc/lvm/lvm.conf` file to restore the filter and accept SCSI device nodes.
4. Run `vgscan -v` to rebuild the LVM2 cache.
5. Run `lvmdiskscan` to ensure that the filter field is working correctly.
6. Create a new `initrd` image without the PowerPath feature.

```
mkinitrd -i /boot/initrd-2.6.27.19-5-default -k  
/boot/vmlinuz-2.6.27.19-5-default
```

7. Add a new boot entry for the new `initrd` image created in step 6 in `/boot/grub/menu.lst`.
8. Reboot the host and select the boot loader entry with the `initrd` image that is created in step 6.
9. Remove PowerPath following the instructions in [Chapter 5](#), "Removing PowerPath from a Linux host."

Configuring a PowerPath logical device

This chapter includes the following sections:

- ◆ Introduction 58
- ◆ Adding new paths to a PowerPath logical device 59
- ◆ Removing some paths to PowerPath logical devices 61

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

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). This is noted in the appropriate places in the procedures that follow.

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, 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 PowerPath logical device

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

Reconfigurations on a Linux host require you to stop applications. Once the reconfiguration is completed successfully, new native devices can be used like any other native devices.



CAUTION

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

Path limit Do not configure more than 32 paths per logical device.

Procedure To add paths to a PowerPath logical device:

1. Run **powermt display** to confirm the current configuration.
2. Run **powermt display dev=all** to confirm the configuration of the logical devices to which new paths will be added.
3. 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.

The *EMC PowerPath Family Product Guide* contains a description of hardware paths and I/O paths.

If there is a problem, correct it before proceeding.

4. Make physical path additions as required:
 - a. Map the logical device to additional storage-system ports.
 - b. Add new HBAs. For details, refer to your vendor documentation.

- c. Attach cables.
- d. Adjust zoning on your storage array fabric.
5. If using SAN Manager™, Volume Logix, or Access Logix™, make new paths available to the host using those tools.
6. Scan the HBA Bus on which new paths are configured.
echo "---"/>/sys/class/scsi_host/host<#>/scan
7. Run **powermt config** to configure the new paths under PowerPath control.
8. 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.
9. Correct any issues detected.
10. Run **powermt save** to save the new configuration.

Removing some paths to PowerPath logical devices

This procedure removes some (but not all) paths to logical devices configured in PowerPath.

For managed storage array classes, paths not configured in PowerPath are configured automatically when the associated native device is first accessed. If a working path is unconfigured as described in the following procedure, the path is reconfigured automatically the next time it is accessed.



CAUTION

Do not remove the last live path to a logical device unless you plan to remove the logical device entirely, as data access will be interrupted. (If a PowerPath device is open, powermt remove does not remove that device or the last path to that device)



CAUTION

If any application uses a native device associated with any path to be removed, this procedure may fail. Select different paths to remove or stop using that path.

Procedure

To remove some paths to PowerPath logical devices:

1. Run **powermt display dev=all** to confirm the configuration of the logical devices from which paths will be removed.
2. Check the number of existing paths.
The path state should be alive for known good paths and dead for known bad paths.
3. If there is a problem, correct it before proceeding.
4. 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.)
5. Run **powermt display dev=all** to identify the PowerPath HBA number associated with the paths to be removed, and identify the native device associated with the paths to be removed.

In complex topologies, there can be multiple paths on an HBA.

6. Run **pp_inq -btl** to obtain bus, TID and LUN. Save the output of the command.
7. If using character I/O (sg) devices on Linux, run **sg_map -i -x**, and save the output. The **sg** utilities are available at <http://www.torque.net/sg/index.html>.
8. If a native device that points to the path to be removed is in use, modify all applications configured to use that device, so that they use another device.
9. Run **powermt set mode=standby** to set the mode to standby for all paths to be removed. If an HBA is to be removed, specify the HBA (instead of the device) when running **powermt set mode**.
10. Physically disconnect the paths, or logically disconnect them by removing them from active zones.

As a result of this step, the operating system may log some error messages, and the PowerPath Management Daemon will also issue errors when paths are physically disconnected. However, PowerPath prevents these errors from affecting applications.

11. Run **powermt restore**, and confirm that the paths listed as dead are the ones you intended to remove.
 - a. Note down all the native device names that the paths listed as dead are the ones you intended to remove.
12. Run **powermt remove**, specifying both the **hba** and **dev** arguments.



CAUTION

If you do not specify the hba argument, all paths to the specified logical device are removed.



CAUTION

Ensure that no applications will try to use the native device that is being removed. If a native path is accessed by an application or utility after powermt remove is run but before the physical path is removed or the platform-specific unconfigure device operations are completed, PowerPath may reconfigure the path.

13. Remove the stale native devices associated with the paths that were removed (the native device names that are noted in Step 11a).

Note: This step is applicable for SLES 10 and RHEL5 only. SLES 11 SP1 automatically removes the stale native device.



IMPORTANT

It is critical to remove the dead paths from PowerPath as mentioned in [step 12](#) before removing the stale native devices from the kernel. If not, the host can panic.

To remove the stale native devices from the kernel:

```
echo 1 > /sys/block/<Native device>/device/delete
```

14. Inspect the new PowerPath configuration:
 - a. Run **powermt display**.

The output should show fewer total paths than before. All paths should have a state of optimal.
 - b. Run **powermt display dev=all**.

All remaining paths associated with the affected logical devices should be displayed with a state of alive.
15. Correct any issues detected above, before saving the PowerPath configuration or using the new logical devices.
16. Run **powermt save** to save the new configuration.

This chapter includes the following sections:

- ◆ Management daemon using SNMP 66
- ◆ Installing QLogic qla2300 HBAs after PowerPath installation 70
- ◆ Saving new configurations 71
- ◆ Configuring LVM2 support..... 72
- ◆ Dynamic LUN addition and removal 76
- ◆ Hot swapping an HBA 78
- ◆ R1/R2 boot failover support 80
- ◆ Setting the naming convention to enclosure-based names 83
- ◆ SCSI-2 reservation support in PowerPath for Linux 83
- ◆ Audit and error messages 83

Management daemon using SNMP

The PowerPath management daemon provides an event and alerting mechanism for multipathing events listed in the `emcp_mond.conf` file. These SNMP traps are viewable through an SNMP manager and display alerts based on filtering criteria, such as a path going dead.

The management daemon requires the `net-snmp` package to be installed and configured according to the `net-snmp` documentation. The management daemon sends SNMPv2c traps.

Monitored events

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

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

- ◆ A path is dead for one minute.
- ◆ All paths to a volume are dead for five minutes.
- ◆ An I/O has taken more time than the set threshold on a path.

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

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

Table 1 Management daemon event causes (1 of 2)

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

Table 1 Management daemon event causes (2 of 2)

Monitored event	Event cause	Event ID	SNMP Trap
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
Path Latency Max Exceeded	An I/O has taken more time than the set threshold on a path.	0x80000,0x80000,0	Yes

The time threshold for each event can be modified as described in ["Editing the event"](#) on page 69.

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.

In the configuration file, some un-supported events such as **bus is dead** and **vol is dead** are also included. But these will not work even if they are enabled.

Configuring the management daemon

After PowerPath Version 5.5 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.

PowerPath management daemon on SLES 11 SP1 PPC64 hosts

The PowerPath management daemon requires the 32-bit version of the SNMP libraries, but if you are running SLES 11 SP1 PPC64 architecture, SLES 11 SP1 does not install the 32-bit SNMP libraries. To use the PowerPath management daemon on SLES11 SP1 PPC64 hosts, use YaST to install the 32-bit version of the SNMP libraries (**libsnmpp15-32bit**) or install the 32-bit version of the SNMP libraries rpm package, and then configure the management daemon.

Procedure To configure the management daemon:

1. Rename or copy the `/etc/opt/emcpower/emcp_mond-sample.conf` file to `/etc/opt/emcpower/emcp_mond.conf`.
2. Edit `/etc/opt/emcpower/snmpd.conf` to identify where the trap should be sent.

```
< trap2sink <Trap_Destination_IP> <Trap_destination_port#>
```

where:

- `<Trap_Destination_IP>` is the IP address of the host where the SNMP manager is installed.
- `<Trap_destination_port#>` is 162. This is the default port number of the host where the SNMP manager is listening.

For example:

```
< trap2sink 192.xxx.xxx.xxx 162
```

3. To configure snmpd persistent across reboot.
 - a. Copy `/etc/opt/emcpower/snmpd.conf` to default snmp configuration file (`/etc/snmp/snmp.conf`).



CAUTION

Caution: If there is an `snmpd.conf` file already located in `/etc/snmp/snmp.conf`, the previous file will be overwritten.

- b. Check that the snmpd service is enabled by using `chkconfig` command; if not then enable the service.

c. Start the snmpd daemon: **/etc/init.d/snmpd start**

If you are not planning to configure snmpd persistent across reboot, type: **snmpd -c /etc/opt/emcpower/snmpd.conf** to start the SNMP daemon.

4. Type **ps -ef | grep snmpd** to check the status.

Output similar to the following appears:

```
'root      17448  1  0  Oct08?   00:00:08 /usr/sbin/snmpd -c
            /etc/opt/emcpower/snmpd.conf
```

5. Run **/etc/init.d/emcp_mond.sh start** to start the emcp_mond daemon.

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

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

```
to:
```

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

3. Close the emcp_mond.conf file.
4. Restart the daemon:
 - a. Run **/etc/init.d/emcp_mond.sh stop**.
 - b. Run **/etc/init.d/emcp_mond.sh start**.

Installing QLogic qla2300 HBAs after PowerPath installation

If you install QLogic qla2300 HBAs and drivers after PowerPath has been installed on a host connected to an HP StorageWorks EVA array, add the following lines to the `/etc/modprobe.conf.pp` file:

```
install qla2xxx /sbin/modprobe emcpsf; /sbin/modprobe qla2xxx --ignore-install
install emcp /sbin/modprobe qla2xxx; /sbin/modprobe emcp --ignore-install
```

Saving new configurations

When new devices are added to or old devices are removed from the host, the configuration changes.

Any change in the device configuration in the host requires you to carry out an explicit **powermt save** to save the changed configuration.

Configuring LVM2 support

You must modify the `/etc/lvm/lvm.conf` file to filter out `sd` device nodes from its internal cache so that LVM2 recognizes a PowerPath pseudo device as the single path to LUN. The following sections describe how to modify `/etc/lvm/lvm.conf` on a host where the root file system is:

- ◆ not mounted on a logical volume.
- ◆ is mounted on a logical volume.
- ◆ is mounted on mirrored logical volumes.

Note: If duplicate PV errors occur in a boot from SAN configuration, modify the recommended filters shown in this section to prevent LVM2 from scanning the pseudo device partition node (which causes this issue).

For example, for RHEL 5, replace the recommended filter for the root file system mounted on a logical volume with the following:

```
filter = [ "a/sda[1-9]$/", "r/emcpowera2/", "r/sd.*/",
"r/disk.*/", "a/.*/" ]
```

Root file system not mounted on a logical volume

Modify the **filters** field of the `/etc/lvm/lvm.conf` file to prevent LVM2 from scanning `sd` device nodes.

1. Modify the **filter** field in the `/etc/lvm/lvm.conf` file. Replace:

```
filter = [ "a/.*/" ]
```

with the appropriate line for your Linux kernel, as follows:

Kernel	Filter
RHEL 5 (5.0-5.5)	<code>filter = ["r/sd.*/", "r/disk.*/", "a/.*/"]</code>
SLES 10 SP4 SLES 11 SP1	<code>filter = ["r /dev/.*/by-path/.*/", "r /dev/.*/by-id/.*/", "r /dev/.*/by-name/.*/", "r/sd.*/", "a/.*/"]</code>

2. Run `vgscan -v` to rebuild the LVM2 cache.
3. Verify that the filter field is working correctly.

4. Run **lvmdiskscan** and verify that the *filtered* device nodes are not listed in the command output.

Root file system mounted on a logical volume

Identify the underlying devices for the root/swap logical volumes and the /boot devices (if any). Modify the **filters** field of the /etc/lvm/lvm.conf file to prevent LVM2 from scanning all sd device nodes except for root/swap and /boot devices:

1. Run **df -k** or **mount** to identify the root/swap logical volumes and the /boot devices (if any).
2. Identify the underlying devices for the root/swap logical volumes. For example, if the root file system is mounted on logical volume /dev/vg01/lv01, type:

```
vgdisplay -v /dev/vg01/lv01
```

3. Modify the **filter** field in the /etc/lvm/lvm.conf file to prevent LVM2 from scanning all sd device nodes except for root/swap and /boot devices identified in steps 1 and 2.

Note: Modify the filter as needed using standard shell-scripting regular expressions. For example, to include partitions sda1 to sda9 for LVM2 while filtering out the remaining sd device nodes, set the filter field to:
`filter=["a/sda[1-9]$/", "r/sd.*/", "a.*"/]`

For example, if the underlying device for the root/swap file system is /dev/sda2 and /boot is mounted on /dev/sda3, set the filter as follows:

Kernel	Filter
RHEL 5 (5.0-5.5)	<code>filter = ["a/sda[1-9]\$/", "r/sd.*/", "r/disk.*/", "a.*/"]</code>
SLES 10 SP4 SLES 11 SP1	<code>filter = ["a/sda[1-9]\$/", "r /dev/.*/by-path/.*/", "r /dev/.*/by-id/.*/", "r /dev/.*/by-name/.*/", "r/sd.*/", "a.*/"]</code>

Note: In the above filter value, sda is an internal hard disk that is not managed by PowerPath. Therefore, the lvm filter value should include the boot device and all its partitions as sda [1-9].

4. Run **vgscan -v** to rebuild the LVM2 cache.

5. Verify that the **filter** field is working correctly.
6. Run **lvmdiskscan** and verify that the **sd** device nodes containing the root/swap/boot devices identified in steps 1 and 2 are listed in the command output, and that the *filtered* device nodes are not listed in the command output.

Root file system mounted on mirrored logical volumes

Identify the underlying devices for the root/swap logical volumes and the /boot devices (if any). Modify the filters field of the **/etc/lvm/lvm.conf** file to prevent LVM2 from scanning all **sd** device nodes except for root/swap and /boot devices:

1. Run **df -k** or **mount** to identify the root/swap logical volumes and the /boot devices (if any).
2. Identify the underlying devices for the root/swap and mirror logical volumes. For example, if the root file system is mounted on logical volume **/dev/vg01/lv01** and its mirror is mounted on **/dev/vg01/lv02**, type:

```
vgdisplay -v /dev/vg01/lv01
vgdisplay -v /dev/vg01/lv02
```

3. Modify the **filter** field in the **/etc/lvm/lvm.conf** file to prevent LVM2 from scanning all **sd** device nodes except for root/swap and /boot devices identified in steps 1 and 2. For example, if the underlying logical device for the root/swap file system is **/dev/sda2** and its mirror is **/dev/sdb3**, set the filter field to:

Kernel	Filter
RHEL 5 (5.0-5.5)	<code>filter = ["a/sda[1-9]\$/", "a/sdb[1-9]\$/", "r/sd.*/", "r/disk.*/", "a./.*/"]</code>
SLES 10 SP4 SLES 11 SP1	<code>filter = ["a/sda[1-9]\$/", "a/sdb[1-9]\$/", "r /dev/.*/by-path/.*/", "r /dev/.*/by-id/.*/", "r /dev/.*/by-name/.*/", "r/sd.*/", "a./.*/"]</code>

Note: Modify the filter as needed using standard shell-scripting regular expressions. For example, to include partitions **sda1** to **sda9** for LVM2 while filtering out the remaining **sd** device nodes, set the filter field to:
`filter=["a/sda[1-9]$/", "r/sd.*/", "a./.*/"]`

4. Run **vgscan -v** to rebuild the LVM2 cache.

5. Verify that the **filter** field is working correctly.
6. Run **lvmdiskscan** and verify that the sd device nodes containing the root/swap/boot devices identified in steps 1 and 2 are listed in the command output, and that the *filtered* device nodes are not listed in the command output.

Dynamic LUN addition and removal

PowerPath supports dynamically adding and removing LUNs when PowerPath is installed on RHEL 5, SLES 10 SP4, and SLES 11 SP1 systems.



IMPORTANT

EMC recommends quiescing I/O to the maximum extent possible during any dynamic addition or removal operations. Cluster or database integrity may be adversely affected if moderate or heavy I/O is running while LUN configuration is changed.

Adding a LUN

To add a new LUN to an existing PowerPath configuration:

1. Scan the bus for the LUNs.
2. Ensure that the new LUN is created on the storage side and recognized by Linux.

RHEL 5.x, SLES 10 SP4, and SLES 11SP1

For example, if `powermt display` lists HBAs 8 and 9, initiate the scan by typing:

```
echo "--" > /sys/class/scsi_host/host8/scan
```

```
echo "--" > /sys/class/scsi_host/host9/scan
```

For iSCSI software initiator, type:

```
echo "--" > /sys/class/scsi_host/host<controller>/scan
```

3. Run `powermt config` to configure the new LUN.
4. Run `powermt display` to confirm the addition of the new LUN.

Removing a LUN

To remove a LUN and all associated PowerPath and Linux devices from the host environment on RHEL 5, SLES 10 SP4, and SLES 11 SP1:



IMPORTANT

It is critical to follow the procedures in exact order because deviating from the procedures can cause the host to panic.

1. Stop any I/O to the device being removed.



CAUTION

It is critical to stop all I/O on the device that is being removed.

2. Run **powermt display dev=emcpower<id>** to determine which native SCSI devices are mapped to the pseudo device.

where <id> corresponds to the LUN to be removed.

3. Run **powermt remove dev=emcpower<id>**.

where <id> corresponds to the LUN to be removed.

4. Run **powermt release**.

Failing to run this command results in the pseudo device still being visible in /dev and /sys/block and may lead to some complications when the new devices are dynamically added.

5. Run **echo 1 > /sys/block/sd<id>/device/delete** for each native sd device path that was associated with the LUN.

This command causes PowerPath to stop processing I/O for the native device and removes it from /dev and /sys/block. Failing to run this command may lead to some complications when the new devices are dynamically added.

6. Remove the device from the storage array.
7. Run **powermt display** to verify that the device has been removed.

Note that on Linux systems it is possible that mapping of the native device name to a LUN can change across reboots if any configuration undergoes a change that can lead to a change in device scan order. This can affect the way PowerPath pseudo devices are mapped to corresponding native device names.

For example, when a LUN, for instance LUN 2, is attached to a host, PowerPath might configure it as pseudo device *emcpowera* with the following native devices: *sdc, sdd, sde, sdf*.

When another LUN, for instance LUN 3, is added to the host, PowerPath might configure it as pseudo device *emcpowerb* with the following native devices: *sdg, sdh, sdi, sdj*.

If the system is rebooted at that time, *emcpowera* remains mapped to LUN 2 and *emcpowerb* remains mapped to LUN 3, but the native devices may no longer be mapped as before. The native devices might be mapped as follows after the reboot:

LUN 2 — *emcpowera* — *sdc, sde, sdg, sdi*

LUN 3 — *emcpowerb* — *sdd, sdf, sdh, sdj*

To avoid this issue in mapping changes, using the pseudo device names (*emcpower*) or UUID instead of native device names (*sd*) in */etc/fstab* is recommended.

The *EMC PowerPath Family Product Guide* has more information on adding new logical devices to the PowerPath configuration.

Hot swapping an HBA

PowerPath 5.5 supports the Linux PCI hot plug feature on RHEL 5, and SLES 10 SP4 and SLES 11 SP1, which allows you to hot swap an HBA card using Fujitsu hardware and drivers.

Removing an HBA

To remove an HBA:

1. Verify that the Fujitsu drivers **shfchp** and **pciehpc** are loaded.
2. Identify the HBA to be removed:
 - a. Run **powermt** to show the HBA number that maps to those entries found under */sys/class/scsi_host*.
 - b. Type **cat /sys/class/scsi_host/<host>/info** to extract the PCI bus number information.

For example, If **powermt** display lists HBA 8, the output from **cat /sys/class/scsi_host/host8/info** will be:

```
[root@l82bi194 host8]# cat info
```

```
Emulex LP10000 2Gb PCI-X Fibre Channel Adapter on
PCI bus 10 device 08 irq 70
```

The relevant information is PCI bus 10.

- c. Translate the hexadecimal bus number from 0x10 to decimal 16.

This will give the corresponding PCI slot number.

For example, the available slots given from

```
[root@l82bi194 host8]# ls -F /sys/bus/pci/slots
will be:
```

```
0009_0000/ 0019_0019/ 0041_0196/ 0046_0198/
0191_0019/ 0196_0003/ 0043_0019/ 0048_0199/
```

The correct slot is 0196_0003.

3. Type **powermt disable hba=#** to stop all I/O associated with the HBA that you want to remove.
4. Shut off power to the associated HBA. Type:

```
echo 0 > /sys/bus/pci/slots/<0196_0003>/power
```

where <0196_0003> is the slot number identified in step 2c.

Running **powermt display** confirms that the HBA has been removed. In addition, SCSI sd devices are no longer present in `/sys/block` and `/dev`.

Adding an HBA

To add an HBA:

1. Verify that the Fujitsu drivers **scpchp** and **pciehpc** are loaded.
2. Power on the new HBA. Type:

```
echo 1 > /sys/bus/pci/slots/[pci bus #_pci slot #]/power
```

3. Run **powermt display** to verify the device configuration.
4. Run **powermt config**.

Note that this process does not include re-enabling the HBA through PowerPath by running:

```
powermt enable hba=#
```

According to PowerPath, the original HBA number has been removed from the system. When the new HBA is added, it is not assigned original host# but rather the next available host identifier.

R1/R2 boot failover support

If a storage system device corresponding to a bootable emcpower device is mirrored through SRDF[®], it is possible in the event of a server failure at the local storage system to fail over the boot disk to the remote mirror disk and then boot the server on an identical remote host.

Contact EMC Customer Support for assistance when configuring R1/R2 boot disk failover.

R1/R2 supported configurations

EMC supports the following specific R1/R2 configurations:

- ◆ Each boot host must be connected to only one Symmetrix.
- ◆ The two hosts must have identical hardware.
- ◆ All R1 devices reside on one Symmetrix, *Symmetrix A*, and are visible only to a single host, *Host A*.
- ◆ All R2 devices reside on a separate Symmetrix, *Symmetrix B*, and are visible only to the identical host in reserve, *Host B*.
- ◆ Each R1 device has only one mirror. (Concurrent SRDF is not supported.)
- ◆ Neither host has non-mirrored devices, BCVs, or gatekeepers.
- ◆ SRDF is managed from either of the following two facilities:
 - EMC ControlCenter[®] Management Server
 - Symmetrix Service Processor

Note: When the system is booted on the R2 device, PowerPath does not support reconfiguring or administering devices (for example, adding or removing devices) on the R1 side. Device configuration and administration must be done on the R2 side.

R1/R2 boot procedure

Note: In an R1/R2 environment:

- The R1/R2 hosts must have the same host architecture. For example, if the R1 host is an x86_64 architecture, then the R2 host must be the same.
- HBA adapters must also be identical on the R1/R2 hosts.

R1/R2 boot support assumes that the systems are configured to boot from a Linux LVM device (that is, a Linux Logical Volume Manager device).

When the host on the R2 side boots, it is connected to a different Symmetrix system and set of volume IDs. Therefore, the `powermt.custom` file (which is identical to the R1 files since the boot disk is identical) is modified to create a valid mapping between the `emcpower` device and native path device for both R1 and R2 locations. PowerPath will determine which Symmetrix volume IDs are valid (that is, the visible ones) and will act accordingly when either the R1 or the R2 host is booted.

Configure the network interface on the R1 host

To ensure R1/R2 boot failover support, you must configure the network interface on the R1 host. Refer to the appropriate section for your Linux kernel.

RHEL 5

To configure the R1 host on RHEL 5:

1. Type `chkconfig kudzu off` to disable `kudzu`.
2. Type `chkconfig --list kudzu` to verify that `kudzu` is disabled.

Output similar to the following appears:

```
kudzu 0:off 1:off 2:off 3:off 4:off 5:off 6:off
```

3. Change directories to `/etc/sysconfig/network-scripts`. Type:


```
cd /etc/sysconfig/network-scripts
```
4. Type `# cat ifcfg-eth0`.
5. Comment out the `HWADDR` line by adding a `#` symbol at the beginning of the line. For example:

```
#HWADDR=00:0D:60:14:C2:98
```

SLES 10 SP4 and SLES 11 SP1

To configure the R1 host on SLES 10 SP4, and SLES 11 SP1:

1. Using YaST, go to **Network Devices > Network Card**, and then select **Change**.
2. Delete any previously configured network interface. For example, if the **eth0** interface is configured, delete it from YaST.
3. Click **Finish** to save the configuration.
4. On SLES 10 hosts only, change the persistent binding setting to **no**, as follows:
 - a. Open the config file for editing:

```
vi /etc/sysconfig/network/config
```

- b. In the config file, change the persistent binding value to **no**:

```
FORCE_PERSISTENT_NAMES=no
```

5. Open **Network Devices > Network Card** and select the network interface you want to configure, and then select **Edit** (SLES 10 SP4).
6. In the **Configuration Name** field, enter the number of the interface you deleted earlier.

For example, to configure the eth0 interface:

- On SLES 9, enter **eth0**.
 - On SLES 10, enter **0**.
7. Enter the IP address in the **IP Address** field, and update other network data as necessary.
 8. Select **Next** and then **Finish** to save the configuration.
 9. Verify the network service has started correctly. Type:

```
/etc/sysconfig/network # ifconfig
```

The network data appears for the eth0 interface in the ifconfig output.

Setting the naming convention to enclosure-based names

Due to a known Symantec issue with VxVM 5.0 MP3 and above, the entry in the `vxdisk list` command output does not change to enclosure-based naming when the naming system is changed through `vxddladm set namingscheme=ebn` or `vxdiskadm`.

To change the enclosure-based naming system for emcpower devices with VxVM 5.0 MP3 and above:

1. Type `vxddladm set namingscheme=ebn` or `vxdiskadm` with the proper option, to set the naming convention to enclosure-based names.
2. Type `vxdlmpadm setattr enclosure <enclosure_name> tpdmode=native` to set the device name to native names on which enclosure names work.

For example:

```
vxddladm set namingscheme=ebn persistence=yes
vxdlmpadm setattr enclosure pp_emc_CLARiiON0 tpdmode=native
```

SCSI-2 reservation support in PowerPath for Linux

PowerPath for Linux does not support SCSI-2 reservation. Therefore, in a VCS (Veritas Cluster Server) configuration, DiskReservation Agent is not supported as this uses SCSI-2 reservation.

Audit and error messages

The *EMC PowerPath Family CLI and System Messages Reference Guide* on Powerlink describes audit logging, including instructions for setting up logging in your environment. It also lists error messages, descriptions, and suggested corrective action. By default, messages are logged to the `/var/log/messages` file.

Removing PowerPath from a Linux host

This chapter includes the following sections:

- ◆ Before you remove PowerPath 86
- ◆ Uninstalling PowerPath 87

Before you remove PowerPath

Before you remove PowerPath from the host:

- ❑ If the **powermt display every=<time_in_seconds>** command is running in a loop, stop the command. The *EMC PowerPath Family CLI and System Messages Reference Guide* and the `powermt(1)` man page contains information about `powermt display`.)
- ❑ Make sure no PowerPath devices are in use. For example, unmount a mounted file system.
- ❑ Stop the Navisphere agent (CLARiiON only).
- ❑ Manually remove references to PowerPath pseudo devices from system configuration files. For example, `/etc/fstab`.
- ❑ If you removed a PowerPath-managed native device file (that is, `/dev/sd[a-z][a-z]`) while PowerPath was loaded, you must use the **mknod** command to recreate that native device file with its default major/minor numbers before trying to unload or uninstall PowerPath.
- ❑ If the host is configured to boot from a Symmetrix or CLARiiON storage system refer to "[Configuring a PowerPath root device on Linux](#)" on page 41 for removal in a boot-from-SAN setup.
- ❑ Undo any LVM filtering, as the pseudos will no longer be available after deinstallation.

Uninstalling PowerPath

To uninstall PowerPath:

1. Log in as root.
2. Display the package name. Type:

```
rpm -qa | grep EMCpower.LINUX
```

3. Remove the software. Type:

```
rpm -ev EMCpower.LINUX.5.5.x.xx.xx_<build>
```

Note: If the PowerPath uninstall program displays a message saying that it could not unload the emcp module, run the **lsmod | grep emcp** command to determine if any PowerPath modules are loaded on the host. If so, restart the host after the uninstallation completes to unload the modules.

Saved configuration files

The PowerPath uninstall process saves the PowerPath configuration files in the `/etc/emc/archive` directory. If necessary, you can use these configuration files to restore the PowerPath configuration on the host, or upgrade to a later version of PowerPath without PowerPath installed on the host while re-using the configuration.

After you remove PowerPath

After you remove PowerPath, you may need to:

- ◆ Remove PowerPath configuration and registration files.

Removing PowerPath files

The PowerPath un-installation process saves the following files, adding the extension `.saved`:

- ◆ `/etc/emc/archive/emcp_devicesDB.dat.saved`
- ◆ `/etc/emc/archive/emcp_devicesDB.idx.saved`
- ◆ `/etc/emc/archive/emcp_registration.saved`
- ◆ `/etc/emc/archive/mpaa.excluded.saved`
- ◆ `/etc/emc/archive/mpaa.lams.saved`
- ◆ `/etc/emc/archive/powermt.custom.saved`

- ◆ /etc/emc/powerkmd.custom

If you are removing PowerPath from the host entirely (that is, you are not planning to re-install PowerPath), use `emcp_cleanup` to remove all these saved files from the host.

1. Type `/sbin/emcp_cleanup` to remove the PowerPath license and all supporting files.

`emcp_cleanup` issues the following warning:



WARNING

Warning: Attempt to remove saved configuration files for previously installed PowerPath version. Removing these files will result in loss of saved configuration. Proceed with caution!!

Do you want to continue [y,n,?] (default: n): Rvig PowerPath

2. Type `y` and press **Enter**.
3. After running `emcp_cleanup` and verifying no files remain, manually delete `/sbin/emcp_cleanup`.

This chapter includes the following sections:

- ◆ Recovering from a failed PowerPath installation..... 90
- ◆ Troubleshooting a PowerPath upgrade..... 91
- ◆ Missing or deleted files 93
- ◆ Incorrect PowerPath configuration 94
- ◆ The `_netdev` option in `/etc/fstab` 95
- ◆ Problems booting the host 95
- ◆ System hangs 95

Recovering from a failed PowerPath installation

If PowerPath installation fails, correct any error conditions reported by the install program. Then complete the following steps. After completing the steps, install PowerPath following the instructions in [Chapter 1, "Installing PowerPath on a Linux host."](#)

1. Edit the appropriate file for your Linux version (shown in the following table):

Linux version	File
RHEL 5 Asianux 3 OEL 5	<code>/etc/rc.sysinit</code> (linked to <code>/etc/rc.d/rc.sysinit</code>)
SLES 10 SP4 SLES 11 SP1	<code>/etc/init.d/boot.d/S04boot.powerpath</code> (linked to <code>/etc/init.d/boot.powerpath</code>)

and remove the following lines:

```
# Configure and initialize PowerPath.
if [ -f /etc/init.d/PowerPath ];
    then /etc/init.d/PowerPath start
fi
```

If the RPM installation succeeds yet the error message indicating that the patch failed is displayed, you should verify that the file you modified contains the lines above.

2. Edit the `/etc/modprobe.conf` file and remove the following lines:

```
###BEGINPP
include /etc/modprobe.conf.pp
###ENDPP
```

3. Type `lsmod | grep emc` to determine if the PowerPath modules are loaded.

If the modules are loaded, type `/etc/init.d/PowerPath stop` to unload them.

4. Rename the `/etc/init.d/PowerPath` file. Type:

```
mv /etc/init.d/PowerPath /etc/init.d/PowerPath.saved
```

5. Check the `/etc/opt/emcpower/log` file for additional information about the failure.

6. Remove the PowerPath entry in the RPM database (if it exists). Type the following commands to remove the entry:

```
rpm -qa | grep EMCpower.LINUX  
rpm -ev EMCpower.LINUX-5.5.x.xx.xx-<build>
```

Note: If this command fails, use the `rpm -ev --noscripts` command.

```
rm -rf /etc/opt/emcpower
```

Note: Ensure that you have unloaded the PowerPath modules and renamed the `/etc/init.d/PowerPath` file before running these commands.

Troubleshooting a PowerPath upgrade

This section describes how to resolve problems that can occur when upgrading PowerPath.

PowerPath custom settings not preserved

If your custom PowerPath settings are lost during the upgrade, check the following for the possible source of the error:

- ◆ The configuration files are corrupt.
- ◆ The PowerPath configuration recorded in the `powermt.custom` file is outdated and does not match the current PowerPath configuration.
- ◆ The file system where the `powermt.custom` file is stored or the file system where PowerPath is being upgraded is full.

- ◆ Connectivity problems on the SAN that cause some devices not to be seen by the host.
- ◆ A fibre driver was not properly loaded when the configuration file was saved or when the upgrade was performed.

Corrupt configuration files

If the PowerPath configuration files are corrupt, your custom configuration is no longer available after you upgrade PowerPath. If an application was configured to use PowerPath pseudo devices, you must reconfigure that application to use the appropriate PowerPath pseudo devices after upgrading PowerPath (if the native-to-pseudo device mapping has changed). Also, run **powermt set policy/priority=<policy>/<0-9>** to reset the desired policy/priority settings for the storage devices in your PowerPath configuration.

Outdated configuration files

If you change your PowerPath configuration but do not run **powermt save**, and then upgrade PowerPath, the custom configuration recorded in the saved `powermt.custom` file is not applied during the PowerPath upgrade.

Run **powermt set policy/priority=<policy>/<0-9>** to reset the desired policy/priority settings for the storage devices in your PowerPath configuration.

Full file system

Expand the space in the file system. Then run **powermt set** to reset the desired policy/priority settings for the storage devices in your PowerPath configuration.

SAN connectivity problems/unloaded Fibre Channel driver

If there is a connectivity problem in the SAN, you must first repair that problem. If a fibre driver was not properly loaded, load the fibre driver. Once all devices are visible to the host, complete the following steps to reinstall PowerPath on the host:

1. Copy the `/etc/emc/archive/powermt.custom.saved` file to a different directory. Then copy the `/etc/emcp_registration` file to this same directory, renaming it `emcp_registration.saved`.
2. Uninstall PowerPath following the instructions in [Chapter 5, "Removing PowerPath from a Linux host."](#)
3. Remove any file saved to the `/etc/emc/archive` directory by the uninstall process.

4. Copy the **powermt.custom.saved** and **emcp_registration.saved** files that you move to a different directory in step 1 back to the `/etc/emc/archive` directory.
5. Reinstall PowerPath on the host as described in "Installing PowerPath" on page 19.

Multiple PowerPath entries in the RPM database

Using the **-i** option to upgrade PowerPath creates multiple PowerPath entries in the RPM database, which can cause ongoing maintenance issues.

To prevent creating multiple PowerPath entries while performing an upgrade, type:

```
rpm -U EMCpower.LINUX-5.5.x.xx.xx-<build>
```

Removing the multiple PowerPath entries

To remove PowerPath:

1. Type `rpm -qa | grep EMCpower.LINUX`.
2. Type `rpm -ev EMCpower.LINUX-5.5.x.xx.xx-<build>`.

Missing or deleted files

If PowerPath files are missing or deleted after installation, PowerPath may not run correctly (or at all). If this happens, remove and then reinstall PowerPath:

1. Type `/etc/init.d/PowerPath stop` to stop PowerPath.

Note: Use the script only for EMC PowerPath module loading and unloading.

2. Remove PowerPath. Type:

```
rpm -qa | grep EMCpower.LINUX
rpm -ev EMCpower.LINUX-5.5.x.xx.xx-<build>
```

Note: If this command fails, use the `rpm -ev --noscripts` command.

3. Remove `/etc/opt/emcpower`. Type:

```
rm -rf /etc/opt/emcpower
```



CAUTION

Remove the `emcpower` directory *only* if you cannot restore the deleted or corrupted files from backups.

4. Reinstall PowerPath. Refer to "[Installing PowerPath](#)" on page 19 for more information.

Incorrect PowerPath configuration

If the Navisphere agent (or any application that holds devices open) is running when you try to configure PowerPath devices, the configuration will be incorrect.

To resolve this problem:

1. Stop the Navisphere agent and any other application that may be holding devices open.

Note: The Navisphere agent is not applicable for IBM PowerPC hosts. Therefore, stopping the Navisphere agent is not necessary for a PowerPC CPU architecture host.

2. Type `/etc/init.d/PowerPath stop` to stop PowerPath.

Note: Use the script only for EMC PowerPath module loading and unloading.

3. Type `/etc/init.d/PowerPath start` to start PowerPath.
4. Restart the applications you stopped in step 1.

The `_netdev` option in `/etc/fstab`

For RHEL 5, PowerPath devices should be mounted with the `_netdev` option instead of the `defaults` option in the `/etc/fstab` file. This will ensure that `fsck` is run later in the boot sequence.

Problems booting the host

If you have previously used `emcpower` devices and subsequently uninstalled PowerPath without removing these `emcpower` devices, the host system fails to boot properly.

Manually remove references to PowerPath pseudo devices (`emcpower` devices) from system configuration files such as `/etc/fstab`, `/etc/auto.master`, and `/etc/*auto*`.

System hangs

System hangs may be caused by a number of reasons. This section offers solutions to some scenarios you may encounter.

Unsupported HBA drivers

If your computer hangs after you install PowerPath or after the first time you reboot after installation, check the E-Lab Interoperability Navigator and verify that your fibre driver is currently supported.

Note: The E-Lab Interoperability Navigator, which is updated every month, is available on the Powerlink website.

The driver version that may have come with your Linux distribution might not be supported by EMC. In this case, use the following recovery procedure:

1. Detach all connections to the storage system before booting.
2. Reboot to a safe kernel and remove the `qla` entry from the `/etc/modprobe.conf` or `/etc/sysconfig/kernel` file.
3. If a RAM disk is used, rebuild it.
4. Restart the host with the new RAM disk.

5. Build the supported driver and incorporate it into a new RAM disk as desired.
6. Reattach the connections to the storage system.
7. Restart the host with the storage attached.

Other causes

If your computer hangs and you have the correct HBA drivers, check the Issue Tracker application on Powerlink for other possible causes.

Files changed by PowerPath

This chapter includes the following sections:

- ◆ Linux files modified by PowerPath installation..... 98
- ◆ Files created by PowerPath installation..... 99

Linux files modified by PowerPath installation

The following files are modified when PowerPath is installed on a Linux host:

- ◆ `/etc/modprobe.conf`
- ◆ `/etc/init.d/boot.localfs`
- ◆ `/etc/profile`
- ◆ `/etc/rc.d/rc.sysinit`
- ◆ `/lib/modules/*/modules.dep` — For each `/lib/modules` subdirectory that gets a set of PowerPath drivers installed.

Files created by PowerPath installation

The files in this section are created when PowerPath is installed on a Linux host:

- /etc**
- ◆ **emc/powerkmd.custom**
 - ◆ emc/mpaa.excluded
 - ◆ emc/mpaa.lams
 - ◆ emc/ppme
 - ◆ emcp_devicesDB.dat
 - ◆ emcp_devicesDB.idx
 - ◆ emcp_registration
 - ◆ modprobe.conf.pp
 - ◆ powermt.custom
 - ◆ **rc.d/init.d/PowerMig**
 - ◆ **rc.d/init.d/PowerIscsi**
 - ◆ **rc.d/init.d/PowerPath**
 - ◆ init.d/emcp_mond.sh
 - ◆ **init.d/PowerMig**
 - ◆ **init.d/PowerIscsi**
 - ◆ **init.d/PowerPath**

RHEL only

- ◆ rc.d/rc3.d/S24PowerIscsi
- ◆ rc.d/rc3.d/S99PowerMig
- ◆ rc.d/rc3.d/K99PowerMig
- ◆ rc.d/rc3.d/S51emcp_mond
- ◆ rc.d/rc3.d/K51emcp_mond
- ◆ rc.d/rc5.d/S24PowerIscsi
- ◆ rc.d/rc5.d/S99PowerMig
- ◆ rc.d/rc5.d/K99PowerMig
- ◆ rc.d/rc5.d/S51emcp_mond
- ◆ rc.d/rc5.d/K51emcp_mond

SLES only

- ◆ init.d/emcp_mond.sles
- ◆ init.d/boot.powerpath
- ◆ init.d/poweriscsi.sles
- ◆ init.d/powermig.sles
- ◆ init.d/boot.d/S06boot.powerpath
- ◆ init.d/boot.d/K16boot.powerpath
- ◆ rc.d/rc3.d/Sxxpoweriscsi.sles
- ◆ rc.d/rc3.d/Kxxpoweriscsi.sles
- ◆ rc.d/rc3.d/Sxxpowermig.sles
- ◆ rc.d/rc3.d/Kxxpowermig.sles
- ◆ rc.d/rc3.d/Sxxemcp_mond
- ◆ rc.d/rc3.d/Kxxemcp_mond

- ◆ rc.d/rc5.d/Sxxpoweriscsi.sles
- ◆ rc.d/rc5.d/Kxxpoweriscsi.sles
- ◆ rc.d/rc5.d/Sxxpowermig.sles
- ◆ rc.d/rc5.d/Kxxpowermig.sles
- ◆ rc.d/rc5.d/Sxxemcp_mond.sles
- ◆ rc.d/rc5.d/Kxxemcp_mond.sles

RHEL 5 and SLES 10 only

- ◆ udev/rules.d/05-udev-pp.rules

/etc/opt/emcpower

- ◆ .os_version
- ◆ .pp_version
- ◆ .tmp/rollbackfiles.tar
- ◆ .tmp/rollbackfiles12549.tar
- ◆ .tmp/undolist
- ◆ .tmp/undolist12549
- ◆ emcpmgr
- ◆ **emcp_db_global_lock**
- ◆ **emcp_db_lock**
- ◆ **EULA.pdf**
- ◆ **server.pem**
- ◆ emcp_mond-sample.conf
- ◆ log
- ◆ patch.log
- ◆ powercf
- ◆ snmpd.conf

/etc/opt/emcpower/EMCpower.LINUX-5.5.0.00.00

- ◆ **boot-powerpath.sh**
- ◆ **EULA.pd**
- ◆ **setup-powerpath.sh**
- ◆ .pp_build_version
- ◆ .pp_vendor
- ◆ .pp_version
- ◆ boot.powerpath
- ◆ emcp_mond.sh
- ◆ emcp_mond.sles
- ◆ enable
- ◆ man/.man
- ◆ man/emcpadm.1
- ◆ man/emcpreg.1
- ◆ man/emcpupgrade.1
- ◆ man/powerformat.1
- ◆ man/powermig.1
- ◆ man/powermt.1
- ◆ man/powerprotect.1

- ◆ modprobe.conf.pp
- ◆ PowerIscsi
- ◆ poweriscsi.sles
- ◆ PowerMig
- ◆ powermig.sles
- ◆ PowerPath.rhel
- ◆ PowerPath.sles
- ◆ pp_udev.sh
- ◆ udev-pp.rules

/etc/opt/emcpower/EMCpower.LINUX-5.5.0.00.00/bin

- ◆ .cmds
- ◆ .drivers_base
- ◆ .drivers_ext
- ◆ .drivers_mgr
- ◆ .sharedlibs
- ◆ **.sharedlibs32**
- ◆ .staticlibs
- ◆ **cmds/emcp_cleanup**
- ◆ **cmds/server.pm**
- ◆ cmds/emcpadm
- ◆ cmds/emcpdiscover
- ◆ cmds/emcpmgr
- ◆ cmds/emcpmigd
- ◆ cmds/emcp_mond
- ◆ cmds/emcp_mond-sample.conf
- ◆ cmds/emcppurge
- ◆ cmds/emcpreg
- ◆ cmds/powercf
- ◆ cmds/powermig
- ◆ cmds/powermt
- ◆ cmds/powerprotect
- ◆ cmds/pp_inq
- ◆ cmds/snmpd.conf
- ◆ **/lib/emcp_mond_netsnmp.so**
- ◆ **/lib/libemcp_mpapi_rtl.so**
- ◆ /lib/libemcp_core.so
- ◆ /lib/libemcp_lam.so
- ◆ /lib/libemcp_lic_rtl.so
- ◆ /lib/libemcp_mp_rtl.so
- ◆ /lib/libemcpmp.so
- ◆ /lib/libemcp.so
- ◆ /lib/libpn.so

/etc/opt/emcpower/EMCpower.LINUX-5.5.0.00.00/i18n/catalog

- ◆ de/LC_MESSAGES/EMCpower.mo
- ◆ de/LC_MESSAGES/PowerPath.mo
- ◆ es/LC_MESSAGES/EMCpower.mo
- ◆ es/LC_MESSAGES/PowerPath.mo
- ◆ fr/LC_MESSAGES/EMCpower.mo
- ◆ fr/LC_MESSAGES/PowerPath.mo
- ◆ it/LC_MESSAGES/EMCpower.mo
- ◆ it/LC_MESSAGES/PowerPath.mo
- ◆ ja/LC_MESSAGES/EMCpower.mo
- ◆ ja/LC_MESSAGES/PowerPath.mo
- ◆ ko/LC_MESSAGES/EMCpower.mo
- ◆ ko/LC_MESSAGES/PowerPath.mo
- ◆ pt/LC_MESSAGES/EMCpower.mo
- ◆ pt/LC_MESSAGES/PowerPath.mo
- ◆ zh/LC_MESSAGES/EMCpower.mo
- ◆ zh/LC_MESSAGES/PowerPath.mo

/lib

- ◆ **libemcp_mpapi_rtl.so**
- ◆ libemcp_core.so
- ◆ libemcp_lam.so
- ◆ libemcp_lic_rtl.so
- ◆ libemcp_mp_rtl.so
- ◆ libemcpcm.so
- ◆ libemcp.so
- ◆ libpn.so

/sbin

- ◆ **emcp_cleanup**
- ◆ emcp_mond
- ◆ emcpmigd
- ◆ emcpadm
- ◆ emcpdiscover
- ◆ emcppurge
- ◆ emcpreg
- ◆ powermig
- ◆ powermt
- ◆ powerprotect

/usr/share/man/man1

- ◆ emcpadm.1
- ◆ emcpreg.1
- ◆ emcpupgrade.1
- ◆ powerformat.1
- ◆ powermig.1
- ◆ powermt.1
- ◆ powerprotect.1

/var/lock/subsys

- ◆ PowerIscsi
- ◆ PowerMig

x86_64 hosts only

/lib64

- ◆ **emcp_mond_netsnmp.so**
- ◆ **libemcp_mpapi_rtl.so**
- ◆ libemcp_core.so
- ◆ libemcp_lam.so
- ◆ libemcp_lic_rtl.so
- ◆ libemcp_mp_rtl.so
- ◆ libemcphp.so
- ◆ libemcp.so
- ◆ libpn.so

ia64 hosts only

/emul/ia32-linux/lib

- ◆ libemcp_core.so
- ◆ libemcp_lam.so
- ◆ libemcp_lic_rtl.so
- ◆ libemcp_mp_rtl.so
- ◆ libemcphp.so
- ◆ libemcp.so
- ◆ libpn.so

x86_64 and ia64 hosts

/etc/opt/emcpower/EMCpower.LINUX-5.5.0.00.00/bin

- ◆ .sharedlibs32
- ◆ lib32/libemcp_core.so
- ◆ lib32/libemcp_lam.so
- ◆ lib32/libemcp_lic_rtl.so
- ◆ lib32/libemcp_mp_rtl.so
- ◆ lib32/libemcphp.so
- ◆ lib32/libemcp.so
- ◆ lib32/libpn.so

Driver files

`/etc/opt/emcpower/EMCpower.LINUX-5.5.0.00.00/bin/driver/<kernel_type>`

- ◆ `emcpdm.ko`
- ◆ `emcpgpx.ko`
- ◆ `emcpio.ko`
- ◆ `emcp.ko`
- ◆ `emcpx.ko`

Where `<kernel_type>` is one of the following depending on the architecture:

RHEL 5 i386

`rhel5-PAE-i686, rhel5-i686, rhel5-xen-i686`

RHEL 5 x86_64

`rhel5-x86_64, rhel5-xen-x86_64`

RHEL 5 ia64

`rhel5-ia64, rhel5-xen-ia64`

RHEL 5 ppc64

`rhel5-ppc64`

SLES 10 SP3 i386

`sles10sp3bigsm, sles10sp3default, sles10sp3smp, sles10sp3xen, sles10sp3xenpae`

SLES 10 SP3 x86_64

`sles10sp3smp_x8664, sles10sp3xensmp_x8664`

SLES 10 SP3 ia64

`sles10sp3default_ia64`

SLES 10 SP3 ppc64

`sles10sp3ppc64`

SLES 11 SP1 i386

`sles11default, sles11smp, sles11xen, sles11xenpae`

SLES 11 SP1 x86_64

`sles11smp_x8664, sles11xensmp_x8664`

SLES 11 SP1 ia64

`sles11default_ia64`

SLES 11 SP1 ppc64
sles11ppc64***/lib/modules/<kernel_version>/powerpath***

- ◆ emcpdm.ko
- ◆ emcpgpx.ko
- ◆ emcpio.c.ko
- ◆ emcp.ko
- ◆ emcpx.ko

Where *<kernel_version>* is the version of the kernel currently running on a host.

Encryption files

Encryption is supported on RHEL 5 x86_64 and SLES 10 x86_64 only.

/etc	<ul style="list-style-type: none"> ◆ /init.d/PowerXcrypt ◆ rc.d/init.d/PowerXcrypt
/etc/emc/rsa	<ul style="list-style-type: none"> ◆ cst/config ◆ cst/lib/libCSP-NSS.so ◆ cst/lib/libexpat.so ◆ cst/lib/libCSP-lb.so ◆ cst/lib/libCSP-CatalogService.so ◆ cst/lib/libCSP-SM.so ◆ cst/lib/libCSP-LogService.so ◆ cst/lib/libccme_eccaccel.so ◆ cst/lib/libCSP-DestinationFile.so ◆ cst/lib/libCSP-RT.so ◆ cst/lib/libccme_eccnistaccel.so ◆ cst/lib/CSP-MessageCatalogen-US.cat ◆ cst/lib/libCSP-Configuration.so ◆ cst/lib/libCSP-Security.so ◆ cst/lib/libCSP-PAM.so ◆ cst/lib/libCSP-PAMAuth.so ◆ cst/lib/libCSP-XML.so ◆ cst/lib/libccme_ecc.so ◆ cst/lib/libCSP-DestinationSyslog.so ◆ cst/lib/libCSP-TrustedUser.so ◆ cst/lib/libCST-Admin.so ◆ cst/lib/libCSPI.so ◆ cst/lib/libCSP-LocalizationService.so ◆ cst/lib/libosysrxml.so ◆ cst/lib/libCSP-PositionalFormatter.so ◆ cst/lib/libcryptocme2.sig ◆ cst/lib/pam_cst.so ◆ cst/lib/libCSP-SprintfFormatter.so ◆ cst/lib/libCSP-LayoutRaw.so ◆ cst/lib/libCSP-AuthenticationService.so ◆ cst/lib/libCSP-DestinationEventLog.so ◆ cst/lib/libCSP-UserManagement.so ◆ cst/lib/libCSP-LDAP.so ◆ cst/lib/libCSP-OSLogin.so ◆ cst/lib/libCSP-LocalDirectory.so ◆ cst/lib/libccme_base.so ◆ cst/lib/libCSP-CatalogPosix.so ◆ cst/lib/libCSP-DestinationAlert.so ◆ cst/lib/libCSP-C.so ◆ cst/lib/libcryptocme2.so

- ◆ cst/lib/libCSP-CatalogWindows.so
- ◆ cst/lib/libCSP-RoleManagement.so
- ◆ cst/lib/libosysrt.so
- ◆ cst/lib/libCSP-LayoutFlat.so
- ◆ cst/lib/libCSP.so
- ◆ cst/lib/libCSP-CertService.so
- ◆ cst/lib/libCSP-DestinationODBC.so
- ◆ cst/lib/libCSP-LDAPClient.so
- ◆ cst/lib/libCSP-LayoutXML.so
- ◆ cst/xml/CSP-LocalDirectory.xml
- ◆ cst/xml/CSP-Logging.xml
- ◆ cst/xml/Config-Template.xml
- ◆ cst/xml/CSP-LayoutFlat.xml
- ◆ cst/xml/CSP-saml-schema-assertion-2.0.xsd
- ◆ cst/xml/CSP-DestinationEventLog.xml
- ◆ cst/xml/CSP-UserManagement.xml
- ◆ cst/xml/CSP-saml-schema-assertion-1.1.xsd
- ◆ cst/xml/xmldsig-core-schema.xsd
- ◆ cst/xml/Config.xml
- ◆ cst/xml/CSP-CertificateManagement.xml
- ◆ cst/xml/CSP-LocalizationService.xml
- ◆ cst/xml/CSP-LayoutRaw.xml
- ◆ cst/xml/CSP.xml
- ◆ cst/xml/csp-app.xml
- ◆ cst/xml/CSP-Config.xml
- ◆ cst/xml/CSP-SecurID.xml
- ◆ cst/xml/CSP-Admin.xml
- ◆ cst/xml/CSP-RoleManagement.xml
- ◆ cst/xml/CSP-LayoutXML.xml
- ◆ cst/xml/CSP-DestinationSyslog.xml
- ◆ cst/xml/CSP-CatalogWindows.xml
- ◆ cst/xml/CSP-SprintfFormatter.xml
- ◆ cst/xml/CSP-PositionalFormatter.xml
- ◆ cst/xml/CSP-CatalogPosix.xml
- ◆ cst/xml/CSP-DestinationAlert.xml
- ◆ cst/xml/CSP-xenc-schema.xsd
- ◆ cst/xml/CSP-SM.xml
- ◆ cst/xml/CSP-LDAP.xml
- ◆ cst/xml/CSP-CatalogService.xml
- ◆ cst/xml/CSP-OSLogin.xml
- ◆ cst/xml/CSP-Authn.xml
- ◆ cst/xml/CSP-TrustedUser.xml
- ◆ cst/xml/CSP-PAM.xml
- ◆ cst/xml/CSP-DestinationODBC.xml
- ◆ cst/xml/CSP-DestinationFile.xml
- ◆ rkm_client/lib/libkmcryptolib.so
- ◆ rkm_client/lib/libCSP-lb.so

- ◆ rkm_client/lib/libkmsvcshlib.so
- ◆ rkm_client/lib/libcryptocme2.sig
- ◆ rkm_client/lib/libkmclient_shared.so
- ◆ rkm_client/lib/libccme_base.so
- ◆ rkm_client/lib/libcryptocme2.so
- ◆ rkm_client/config
- ◆ rkm_client/config/rkm_svc.conf.tpl
- ◆ rkm_client/config/rkm_registration.conf.tpl
- ◆ rkm_client/config/rkm_init.conf.tpl
- ◆ rkm_client/config/rkm_keyclass.conf.tpl

/etc/emc/bin

- ◆ **emcp_setup.sh**
- ◆ **emcp_xcryptd**
- ◆ **xcryptd_adm**

/etc/opt/emcpower/bin

- ◆ **envsubst**
- ◆ **gettext**
- ◆ **gettext.sh**
- ◆ xcrypt_upgrade
- ◆ cstadmin
- ◆ xcrypt_config
- ◆ rkmc_cache_conv
- ◆ ckmadm

/etc/opt/emcpower/EMCpower.LINUX-5.5.0.00.00

- ◆ **PowerXcrypt**
- ◆ man/powervt.1
- ◆ man5/.man5
- ◆ man5/rkm_init.conf.5
- ◆ man5/rkm_keyclass.conf.5
- ◆ man5/rkm_svc.conf.5
- ◆ man8/.man8
- ◆ man8/ckmadm.8
- ◆ man8/cstadmin.8
- ◆ man8/xcrypt_config.8
- ◆ man5/rkm_registration.conf.5

/etc/opt/emcpower/EMCpower.LINUX-5.5.0.00.00/bin

- ◆ ckmadm
- ◆ cstadmin
- ◆ emcp_setup.sh
- ◆ emcp_xcryptd
- ◆ rkmc_cache_conv
- ◆ xcrypt_config
- ◆ xcryptd_adm

- ◆ xcrypt_upgrade
- ◆ cmds/powervt
- ◆ rsa/.cst_libs
- ◆ rsa/.rkmc_config
- ◆ rsa/.cst_xml
- ◆ rsa/.rkmc_libs
- ◆ rsa/rkm_client/lib/libkmcryptolib.so
- ◆ rsa/rkm_client/lib/libCSP-lb.so
- ◆ rsa/rkm_client/lib/libkmsvcshlib.so
- ◆ rsa/rkm_client/lib/libcryptocme2.sig
- ◆ rsa/rkm_client/lib/libkmcclient_shared.so
- ◆ rsa/rkm_client/lib/libccme_base.so
- ◆ rsa/rkm_client/lib/libcryptocme2.so
- ◆ rsa/rkm_client/config/rkm_svc.conf.tpl
- ◆ rsa/rkm_client/config/rkm_registration.conf.tpl
- ◆ rsa/rkm_client/config/rkm_init.conf.tpl
- ◆ rsa/rkm_client/config/rkm_keyclass.conf.tpl
- ◆ rsa/cst/lib/libCSP-NSS.so
- ◆ rsa/cst/lib/libexpat.so
- ◆ rsa/cst/lib/libCSP-lb.so
- ◆ rsa/cst/lib/libCSP-CatalogService.so
- ◆ rsa/cst/lib/libCSP-SM.so
- ◆ rsa/cst/lib/libCSP-LogService.so
- ◆ rsa/cst/lib/libccme_eccaccel.so
- ◆ rsa/cst/lib/libCSP-DestinationFile.so
- ◆ rsa/cst/lib/libCSP-RT.so
- ◆ rsa/cst/lib/libccme_eccnistaccel.so
- ◆ rsa/cst/lib/CSP-MessageCatalogen-US.cat
- ◆ rsa/cst/lib/libCSP-Configuration.so
- ◆ rsa/cst/lib/libCSP-Security.so
- ◆ rsa/cst/lib/libCSP-PAM.so
- ◆ rsa/cst/lib/libCSP-PAMAuth.so
- ◆ rsa/cst/lib/libCSP-XML.so
- ◆ rsa/cst/lib/libccme_ecc.so
- ◆ rsa/cst/lib/libCSP-DestinationSyslog.so
- ◆ rsa/cst/lib/libCSP-TrustedUser.so
- ◆ rsa/cst/lib/libCST-Admin.so
- ◆ rsa/cst/lib/libCSPI.so
- ◆ rsa/cst/lib/libCSP-LocalizationService.so
- ◆ rsa/cst/lib/libosysrtxml.so
- ◆ rsa/cst/lib/libCSP-PositionalFormatter.so
- ◆ rsa/cst/lib/libcryptocme2.sig
- ◆ rsa/cst/lib/pam_cst.so
- ◆ rsa/cst/lib/libCSP-SprintfFormatter.so
- ◆ rsa/cst/lib/libCSP-LayoutRaw.so
- ◆ rsa/cst/lib/libCSP-AuthenticationService.so
- ◆ rsa/cst/lib/libCSP-DestinationEventLog.so

- ◆ rsa/cst/lib/libCSP-UserManagement.so
- ◆ rsa/cst/lib/libCSP-LDAP.so
- ◆ rsa/cst/lib/libCSP-OSLogin.so
- ◆ rsa/cst/lib/libCSP-LocalDirectory.so
- ◆ rsa/cst/lib/libccme_base.so
- ◆ rsa/cst/lib/libCSP-CatalogPosix.so
- ◆ rsa/cst/lib/libCSP-DestinationAlert.so
- ◆ rsa/cst/lib/libCSP-C.so
- ◆ rsa/cst/lib/libcryptocme2.so
- ◆ rsa/cst/lib/libCSP-CatalogWindows.so
- ◆ rsa/cst/lib/libCSP-RoleManagement.so
- ◆ rsa/cst/lib/libosysrt.so
- ◆ rsa/cst/lib/libCSP-LayoutFlat.so
- ◆ rsa/cst/lib/libCSP.so
- ◆ rsa/cst/lib/libCSP-CertService.so
- ◆ rsa/cst/lib/libCSP-DestinationODBC.so
- ◆ rsa/cst/lib/libCSP-LDAPClient.so
- ◆ rsa/cst/lib/libCSP-LayoutXML.so
- ◆ rsa/cst/xml/CSP-LocalDirectory.xml
- ◆ rsa/cst/xml/CSP-Logging.xml
- ◆ rsa/cst/xml/Config-Template.xml
- ◆ rsa/cst/xml/CSP-LayoutFlat.xml
- ◆ rsa/cst/xml/CSP-saml-schema-assertion-2.0.xsd
- ◆ rsa/cst/xml/CSP-DestinationEventLog.xml
- ◆ rsa/cst/xml/CSP-UserManagement.xml
- ◆ rsa/cst/xml/CSP-saml-schema-assertion-1.1.xsd
- ◆ rsa/cst/xml/xmlldsig-core-schema.xsd
- ◆ rsa/cst/xml/Config.xml
- ◆ rsa/cst/xml/CSP-CertificateManagement.xml
- ◆ rsa/cst/xml/CSP-LocalizationService.xml
- ◆ rsa/cst/xml/CSP-LayoutRaw.xml
- ◆ rsa/cst/xml/CSP.xml
- ◆ rsa/cst/xml/csp-app.xml
- ◆ rsa/cst/xml/CSP-Config.xml
- ◆ rsa/cst/xml/CSP-SecurID.xml
- ◆ rsa/cst/xml/CSP-Admin.xml
- ◆ rsa/cst/xml/CSP-RoleManagement.xml
- ◆ rsa/cst/xml/CSP-LayoutXML.xml
- ◆ rsa/cst/xml/CSP-DestinationSyslog.xml
- ◆ rsa/cst/xml/CSP-CatalogWindows.xml
- ◆ rsa/cst/xml/CSP-SprintfFormatter.xml
- ◆ rsa/cst/xml/CSP-PositionalFormatter.xml
- ◆ rsa/cst/xml/CSP-CatalogPosix.xml
- ◆ rsa/cst/xml/CSP-DestinationAlert.xml
- ◆ rsa/cst/xml/CSP-xenc-schema.xsd
- ◆ rsa/cst/xml/CSP-SM.xml
- ◆ rsa/cst/xml/CSP-LDAP.xml

- ◆ rsa/cst/xml/CSP-CatalogService.xml
- ◆ rsa/cst/xml/CSP-OSLogin.xml
- ◆ rsa/cst/xml/CSP-Authn.xml
- ◆ rsa/cst/xml/CSP-TrustedUser.xml
- ◆ rsa/cst/xml/CSP-PAM.xml
- ◆ rsa/cst/xml/CSP-DestinationODBC.xml
- ◆ rsa/cst/xml/CSP-DestinationFile.xml

RHEL 5 x86_64 only

- ◆ rhel5_x8664/emcpvlumd.ko
- ◆ rhel5_x8664/emcpxcrypt.ko
- ◆ rhel5xen_x8664/emcpvlumd.ko
- ◆ rhel5xen_x8664/emcpxcrypt.ko

SLES 10 SP3 x86_64 only

- ◆ sles10sp3smp_x8664/emcpvlumd.ko
- ◆ sles10sp3smp_x8664/emcpxcrypt.ko
- ◆ sles10sp3xensmp_x8664/emcpvlumd.ko
- ◆ sles10sp3xensmp_x8664/emcpxcrypt.ko

/lib/modules/<kernel_version>/powerpath

- ◆ emcpvlumd.ko
- ◆ emcpxcrypt.ko

Where *<kernel_version>* is the version of the kernel currently running on a host.

/sbin

- ◆ powervt

/usr/share/man/man1

- ◆ powervt.1

/usr/share/man/man5

- ◆ rkm_init.conf.5
- ◆ rkm_keyclass.conf.5
- ◆ rkm_registration.conf.5
- ◆ rkm_svc.conf.5

/usr/share/man/man8

- ◆ ckmadm.8
- ◆ cstadmin.8
- ◆ xcrypt_config.8

A

- adding an HBA 79
- adding paths 59

B

- boot-from-SAN
 - removing PowerPath 54
 - removing PowerPath (SLES11) 55
 - upgrading Linux kernel 50
 - upgrading Linux kernel (SLES11) 51
 - upgrading PowerPath 53

C

- CLARiiON storage systems, setting up 13
- configuration, troubleshooting 94

D

- deleted files, troubleshooting 93
- device, installation troubleshooting 61
- DiskReservation Agent support 83
- documentation, related 7
- dynamically adding LUNs 76
- dynamically removing LUNs 77

E

- editing the boot.vxvm file 31
- emcp_mond-sample.conf, renaming 68
- emcpreg 26,27
- enclosure based names, setting 83

F

- files, deleted or missing 93

G

- gettext utilities 17

H

- hot swapping an HBA 78

I

- installing PowerPath
 - installation procedure 19
 - post-installation instructions 28
 - pre-installation instructions 12

L

- LANG environment variable 17
- license registration key 12
- Linux operating system, upgrading 35
- Linux PCI hot plug 78
- Linux, updates 16
- localized versions, PowerPath installer 16
- logical devices
 - adding paths 59
 - removing paths 61
- lsmod command 29
- LVM2 (Logical Volume Manager) 72

M

- maintaining PowerPath 65
- major number 15

management daemon
 configuring 68
 restarting 69
missing files, troubleshooting 93
monitored events (SNMP) 66
multiple PowerPath entries 93

N

netdev option, /etc/fstab 95

P

patches. *See also* service packs 12
PCI hot plug 78
post-installation instructions 28
powermt display dev=all 28
PowerPath boot-from-SAN
 configuring 42
 configuring (RHEL) 47
 configuring (SLES) 43
 configuring (SLES11) 45
 removing 54
 removing (SLES11) 55
pre-installation instructions 12

Q

QLogic HBA 65,70

R

R1/R2 boot failover support 80
reconfiguring device mappings 28
registration key 26
removing
 HBA 78
 paths 61
 PowerPath 87
root device considerations 42

S

SCSI-2 reservation support 83
service packs 12
Sistina LVM 72
starting PowerPath 27
SteelEye LifeKeeper cluster 36
Symmetrix storage systems, setting up 13

T

third-party storage systems, setting up 13
time threshold, editing 69
trap destination IP 68
trap destination port number 68
troubleshooting
 deleted or missing files 93
 Linux host not rebooting 14,32
 multiple PowerPath entries 93
 PowerPath configuration 94

U

uninstalling PowerPath 87
updates, Linux 16
upgrading Linux operating system 35
upgrading PowerPath in a SteelEye Life Keeper
 cluster environment 36

V

VCS (Veritas Cluster Server) 83
verifying extensions are loaded 29
verifying PowerPath capabilities 28
Veritas Storage Foundation 25
VxVM (Veritas Volume Manager) 14