



EMC[®] PowerPath[®] for Linux

Versions 6.2 and Minor Releases

Installation and Administration Guide

P/N 302-003-528

REV 01

EMC²

Copyright © 2006 - 2017 EMC Corporation. All rights reserved. Published in the USA.

Published April 2017

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

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

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

For the most up-to-date regulatory document for your product line, go to the technical documentation and advisories section on the EMC online support website.

CONTENTS

Preface	
Chapter 1	Getting Started
Chapter 2	Installing PowerPath
	Before you install 14
	Installing the RPM File 16
	Installing from a CD 16
	Installing in a Xen environment 17
	Installing PowerPath when Veritas Storage Foundation is already installed .. 18
	Licensing after the installation 19
	Post-Installation Steps 20
Chapter 3	Administering in a Boot from SAN environment
	Before configuring for a root device 22
	Configuring PowerPath in a RHEL 5.x Boot from SAN setup 22
	Disabling dm-multipath on a RHEL 5.x host 25
	Configuring PowerPath in a RHEL 6.x or OL 6.x Boot from SAN setup 25
	Disabling dm-multipath on a RHEL 6.x or OL 6.x host 27
	Configuring PowerPath in a RHEL 7.x or OL 7.x Boot from SAN setup 28
	Disabling dm-multipath on a RHEL 7.x or OL 7.x host 30
	Configuring PowerPath in a SLES 11 SP4 Boot from SAN setup 30
	Disabling dm-multipath to prepare for PowerPath installation on SLES 11 SP4 33
	Configuring PowerPath in a SLES 12 SP2 Boot from SAN setup 33
	Upgrading PowerPath in a RHEL 5.x Boot from SAN setup 34
	Upgrading PowerPath in a RHEL 6.x or OL 6.x Boot from SAN setup 35
	Upgrade using emcpbfs_linux script 35
	Upgrade using manual steps 36
	Upgrading PowerPath in a RHEL 7.x, OL 7.x, or SLES 12 SP2 Boot from SAN setup 37
	Upgrade using emcpbfs_linux script 37
	Upgrade using manual steps 38
	Upgrading PowerPath in a SLES 11 SP4 Boot from SAN setup 39
	Upgrading the Linux kernel in a SLES 11 SP4 Boot from SAN setup 40
	Upgrading the Linux kernel in a Boot from SAN setup 42
	Upgrade using emcpbfs_linux script 42
	Upgrade using manual steps 43
	Uninstalling PowerPath in a RHEL 5.x Boot from SAN setup 44
	Uninstalling PowerPath in a RHEL 6.x or OL 6.x Boot from SAN setup 44
	Uninstall using emcpbfs_linux_script 44
	Uninstall using manual steps 45
	Uninstalling PowerPath in a RHEL 7.x or OL 7.x Boot from SAN setup 45
	Uninstall using emcpbfs_linux_script 46
	Uninstalling using manual steps 46
	Uninstalling PowerPath in a SLES 11 SP4 Boot from SAN setup 47
	Uninstalling PowerPath in a SLES 12 SP2 Boot from SAN setup 48

Chapter 4	Managing paths to a logical device	
	Dynamic LUN addition and removal.....	52
	Adding a LUN	52
	Removing a LUN	52
	Resizing a LUN	53
	Before you add or remove paths.....	53
	Adding new paths	53
	Removing some of the paths	54
Chapter 5	Coexistence with other path management software	
	PowerPath and third-party path management software.....	60
	PowerPath behavior	60
Chapter 6	Maintaining PowerPath	
	Configuring for event monitoring	64
	Monitored events	64
	Configuring the Management Daemon	65
	Editing the event	66
	Changing the language of installation messages	66
	Saving new configurations	67
	Configuring LVM2 support	67
	Root file system not mounted on a logical volume	68
	Root file system mounted on a logical volume	68
	Root file system mounted on a logical volume residing on a local disk not managed by PowerPath.....	68
	Hot swapping an HBA.....	70
	Removing an HBA.....	70
	Adding an HBA.....	71
	Configuring for R1/R2 boot failover support	71
	R1/R2 supported configurations	71
	R1/R2 boot procedure.....	72
	Configure the network interface on the R1 host	72
	SCSI-2 reservation support in PowerPath for Linux	72
	Managing and unmanaging EMC Celerra iSCSI devices	72
	Manage EMC Celerra iSCSI devices under PowerPath	72
	Unmanage EMC Celerra iSCSI devices under PowerPath	73
	Enable scsi name on NetApp arrays.....	73
	Removing devices from PowerPath control	74
	Expand the file system on a pseudo-device	74
	Supporting more than 256 storage devices	74
	/etc/fstab options.....	75
	RHEL5 and RHEL6.....	75
	SLES 12 SP2 and RHEL7	75
	SLES11 SP4	76
Chapter 7	Upgrading PowerPath	
	Preparing for upgrade	78
	Upgrading PowerPath.....	78
	Upgrading PowerPath from a CD.....	79
	Upgrading PowerPath from a compressed archive	80
	Upgrading PowerPath across Linux platforms	80
	Upgrading the Linux kernel in a PowerPath installed host.....	81

	Upgrading PowerPath after upgrading from one major Linux release to the next	81
	Upgrading PowerPath before upgrading to the latest SLES Service Pack	82
	Upgrading PowerPath before upgrading RHEL update on Linux	82
	Upgrading PowerPath in a SteelEye LifeKeeper cluster environment	82
Chapter 8	Removing PowerPath	
	Before you remove PowerPath	86
	Uninstalling PowerPath	86
	To uninstall PowerPath	86
	Saved configuration files	87
	Removing PowerPath files	87
Chapter 9	Troubleshooting PowerPath issues	
	Recovering from a failed PowerPath installation	88
	Troubleshooting a PowerPath upgrade	89
	PowerPath custom settings not preserved	89
	Missing or deleted files	89
	The /etc/fstab	90
	Problems booting the host	90
	Boot issue on RHEL6.x, SLES11 SP4, RHEL7.x, SLES 12 SP2, OL6.x, and OL7.x boot-from-SAN when emcpbfs_linux is used	90
	System hangs	91
	Unsupported HBA drivers	92
	Other causes	92
Appendix A	Files changed by PowerPath	
	Linux files modified by PowerPath installation	96
	Files created by PowerPath installation	96

PREFACE

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

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

Note: This document was accurate at publication time. New versions of this document might be released on the EMC Online Support website. Check the EMC Online Support website to ensure that you are using the latest version of this document.

Audience

This manual is part of the PowerPath documentation set. It is intended for use by storage administrators and other information system professionals responsible for installing, using, and maintaining PowerPath software.

Readers of this manual are expected to be familiar with the host operating system, storage system management, and the applications used with PowerPath software.

Related documentation

The complete set of EMC enterprise storage documentation for PowerPath is available on EMC Online Support. These documents are updated periodically.

Revision History

The following table presents the revision history of this document.

Revision	Date	Description
01	April 2017	First release of the EMC PowerPath for Linux Version 6.2

Conventions used in this document

EMC uses the following conventions for special notices:



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



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

CAUTION

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

NOTICE

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

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

IMPORTANT

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

Typographical conventions

EMC uses the following type style conventions in this document:

Normal	Used in running (nonprocedural) text for: <ul style="list-style-type: none"> • Names of interface elements, such as names of windows, dialog boxes, buttons, fields, and menus • Names of resources, attributes, pools, Boolean expressions, buttons, DQL statements, keywords, clauses, environment variables, functions, and utilities • URLs, pathnames, filenames, directory names, computer names, links, groups, service keys, file systems, and notifications
Bold	Used in running (nonprocedural) text for names of commands, daemons, options, programs, processes, services, applications, utilities, kernels, notifications, system calls, and man pages Used in procedures for: <ul style="list-style-type: none"> • Names of interface elements, such as names of windows, dialog boxes, buttons, fields, and menus • What the user specifically selects, clicks, presses, or types
<i>Italic</i>	Used in all text (including procedures) for: <ul style="list-style-type: none"> • Full titles of publications referenced in text • Emphasis, for example, a new term • Variables
Courier	Used for: <ul style="list-style-type: none"> • System output, such as an error message or script • URLs, complete paths, filenames, prompts, and syntax when shown outside of running text
Courier bold	Used for specific user input, such as commands
<i>Courier italic</i>	Used in procedures for: <ul style="list-style-type: none"> • Variables on the command line • User input variables
< >	Angle brackets enclose parameter or variable values supplied by the user
[]	Square brackets enclose optional values
	Vertical bar indicates alternate selections — the bar means “or”
{ }	Braces enclose content that the user must specify, such as x or y or z
...	Ellipses indicate nonessential information omitted from the example

Where to get help

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

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

<http://support.emc.com>.

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

Your comments

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

techpubcomments@emc.com

CHAPTER 1

Getting Started

Refer to related sections for procedures to install and administer PowerPath on a Linux server for EMC and third-party arrays.

Do you want to
install PowerPath?

Look at
[Installing PowerPath](#) for quick installation with a compressed archive from EMC Online Support.

Also includes steps for installing PowerPath from CD, for a Xen environment, or for an environment where Veritas Storage Foundation is already installed.

administer PowerPath?

[Administering in a Boot from SAN environment](#) for administering PowerPath in a BFS environment, for configuring PowerPath logical device, managing PowerPath, or removing PowerPath from the host.

upgrade PowerPath?

[Upgrading PowerPath](#) for pre-requisites and procedures to upgrade PowerPath.

troubleshoot PowerPath?

[Troubleshooting PowerPath issues](#) for procedures to resolve problems while installing and administering PowerPath.

CHAPTER 2

Installing PowerPath

Ensure that all the requirements are met, and then install PowerPath with the compressed PowerPath archive downloaded from EMC Online Support. After installation, license the PowerPath and configure it for operation.

IMPORTANT

The installation process is the same for major releases, service packs, patches, and hotfixes. You need not reboot the host after installing PowerPath.

- ◆ [Installing PowerPath](#) 13
- ◆ [Before you install](#) 14
- ◆ [Installing the RPM File](#) 16
- ◆ [Installing from a CD](#) 16
- ◆ [Post-Installation Steps](#) 20

Before you install

Before installing, understand the expected environment requirements as per the E-Lab™ Interoperability Navigator and prepare the host system for a successful PowerPath installation.

Storage System related steps

Plan your installation based on your storage system:

- EMC Symmetrix® storage systems: Ensure Symmetrix is working properly before PowerPath installation.
- EMC VNX™ OE storage systems: Contact VNX support to plan your installation. PowerPath installation is an integral part of a VNX OE setup and configuration procedure.
- CLARiiON® storage systems: Contact CLARiiON support to plan your installation. PowerPath installation is an integral part of a CLARiiON setup and configuration procedure.
- EMC VPLEX Storage Systems: Ensure VPLEX storage array is working properly before PowerPath installation.
- EMC XtremIO Storage Systems: Ensure XtremIO storage array is working properly before PowerPath installation.
- EMC Unity storage systems: Contact VNX support to plan your installation. PowerPath installation is an integral part of a VNX OE setup and configuration procedure.
- DELL SC: Ensure Dell SC storage array is working properly before PowerPath installation.
- Third-party storage systems: Ensure the third-party array is working properly before installing PowerPath.
- Ensure that the storage system logical devices support PowerPath. The *EMC Host Connectivity Guide* and the CLARiiON storage-system support website provides additional information.

Support matrix related steps

1. Review the Linux kernel updates as per the *EMC Host Connectivity with QLogic Fibre Channel and iSCSI Host Bus Adapters (HBAs) and Converged Network Adapters (CNAs) for the Linux environment guide*.
2. Ensure PowerPath requirements are met as per the *EMC PowerPath for Linux Release Notes*.
3. Ensure PowerPath requirements are met as per the E-Lab™ Interoperability Navigator.
4. Use EMC Grab utilities, PowerPath Configuration Checker, and EMC Reporting tool and verify your environment meets the requirements.
5. Ensure that HBA BIOS and fibre drivers are configured as per E-Lab Interoperability Navigator. Improper settings could cause unwanted I/O delays.

Host system related steps

1. Ensure that the host has a minimum of 2 GB memory free.
2. Ensure that the root partition has at least 120 MB of free space.

Note: Installation fails in case of inadequate space. [“Recovering from a failed PowerPath installation” on page 88](#) contains details on a related, but misleading message from RPM.

3. Ensure that the host is installed with the default file system layout of the Operating System vendor. If not, PowerPath initialization may show an error message and fail.
4. Ensure any version of EMC Unisphere™ and Navisphere® Application Transparent Failover (ATF) is removed from the host by EMC Professional Services.
5. Native Linux DM-multipathing (DM-MPIO) is not compatible with PowerPath, and installing both on the same host causes system instability. See [“Configuring PowerPath in a RHEL 6.x or OL 6.x Boot from SAN setup” on page 25](#), [“Disabling dm-multipath on a RHEL 6.x or OL 6.x host” on page 27](#), or [“Disabling dm-multipath on a RHEL 7.x or OL 7.x host” on page 30](#).

Note: When installing on SLES 12 SP2, ensure native multipathing is not automatically activated during the OS installation, as there is no way to disable native multipathing after the OS installation is complete.

6. To configure LVM, see the section [“Configuring LVM2 support” on page 67](#).

Licensing related steps

1. Ensure that you have the required license key and adhere to the following:
 - The PowerPath for Linux license registration key is on the License Key Card that you received from EMC.
 - License is not required for a host exclusively connected to a CLARiiON AX-series array in a non-virtual environment. PowerPath provides full support.
 - A separate product license is not required for PowerPath Migration Enabler. The PowerPath multipathing license includes Migration Enabler functionality and allows you to migrate data with either Open Replicator or Invista.
2. Verify that `/usr/include/linux/major.h` and `/proc/devices` do not include or define a major number such as 120.

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

Third-party path management software related steps

1. Check the presence of third-party path management software. PowerPath is not supported on third-party path management environments when the corresponding storage system class type is in the managed state.
2. If third-party path management software exists, ensure that the HBAs are dedicated either to PowerPath or the third-party software; HBAs cannot be used for both packages.

Installation language related steps

1. Select your installation language. For example: En_US.utf8

Note: The [“Changing the language of installation messages”](#) section contains procedure for changing the installation language.

Installing the RPM File

From the EMC Online Support site, download the archive file for your operating system and platform, and perform the installation steps as a root user for installing PowerPath.

Note: All the procedures given for RHEL 5 and RHEL 6 are also applicable to Oracle Linux (OL) 5.7 and Oracle VM (OVM) 3.1. PowerPath supports only the UEK kernel of Oracle Linux.

1. Log in as root.
2. Download the appropriate PowerPath archive for your environment from EMC Online Support site.
3. Untar the PowerPath archive.

```
tar -xzf EMCPower.LINUX-<version>.<build>.tar.gz
```

Where:

 <version> is the product version; for example, PowerPath 6.2.

 <build> is the software build number.

4. Copy the RPM package into a temporary folder.
5. Install PowerPath.

```
rpm -ivh EMCPower.LINUX-<version>.<build>.<platform>.x86_64.rpm
```

 Where <platform> is the Linux distribution, for example OL6 or RHEL6.

6. Enter the 24-character alphanumeric sequence from the License Key Card, and follow the on-screen instructions.

Installing from a CD

Mount the PowerPath installation CD for your platform and copy the archive file to a temporary folder from where you can install PowerPath.

1. Log in as root.
2. Insert the PowerPath installation CD into the CD-ROM drive.
3. Create the directory **/cdrom** to be the mount point for the CD-ROM. Type:

```
mkdir /cdrom
```

4. Mount the PowerPath CD on **/cdrom**. Type:

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

- Change directories to your operating system directory.

- For SLES, type:

```
cd /cdrom/LINUX/pp<version>/sles
```

- For RHEL and Asianux, type:

```
cd /cdrom/LINUX/pp<version>/rhel
```

- Copy the RPM package into a temporary folder.

- Type the following installation command:

```
rpm -ivh EMCPower.LINUX-<version>-<build>.<platform>.x86_64.rpm
```

Where:

<version> is the product version; for example, PowerPath 6.2.

<build> is the software build number.

<platform> is the Linux distribution, for example OL6 or RHEL6.

Installing in a Xen environment

After installing PowerPath in a Xen kernel, create a Xen virtual machine to add the newly added device in the Xen environment.

- Log in as root.
- Verify the following Xen tools and kernel RPMs are installed on the host. If not present, you must install them before installing PowerPath. For more information on Xen kernels, refer to the vendor documentation for your operating system.

- Type **uname -a** to verify your system is running the Xen kernel.

- Copy the RPM package into a temporary folder.

- Type the following installation command:

```
rpm -ivh EMCPower.LINUX-<version>-<build>.<platform>.x86_64.rpm
```

- Create a Xen virtual machine. For more information, refer to the vendor documentation for your operating system.

- Select a pseudo device (for example, /dev/emcpowerd) and open the virtual machine configuration file.

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

where <vm_name> is the name of the virtual machine you created in step 6.

- Locate the following line in the configuration file:

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

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

For example, add the following line:

```
disk = [
'file:/images/lcle055/had.img,hda,w', 'phy:</dev/emcpower>, <hdb>,w'
]
```

where:

- `</dev/emcpower>` 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 used 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 6.

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 when Veritas Storage Foundation is already installed

Installing PowerPath after Veritas Storage Foundation requires reconfiguring Veritas Volume manager to recognize the pseudo devices.

- ◆ For Veritas Storage Foundation 6.1 and later, exclude EMC arrays from DMP path management before and after the reconfiguration:

1. Log in as root.

2. Type the following to exclude EMC arrays from DMP path management:

- For Symmetrix arrays:

```
vxddladm excludearray libname=libvxemc.so
```

- For Unity devices:

```
vxddladm excludearray libname=libvxUnity
```

- For VNX OE and CLARiiON arrays:

```
vxddladm excludearray libname=libvxCLARiiON.so
```

- For VPLEX devices:

```
vxddladm excludearray libname=libvxInvista.so
```

- For XtremIO devices:

```
vxddladm excludearray libname=libvxXtremio.so
```

3. Run **vxctl enable**.

Verify enclosure EMC0 is **disconnected** and all EMC LUNs are on disk enclosure EMC0.

4. Copy the RPM package into a temporary folder.

5. Enter the following installation command:

```
rpm -ivh EMCPower.LINUX-<version>-<build>.<platform>.x86_64.rpm
```

6. Run **vxctl enable** so PowerPath claims the devices.

Verify enclosure EMC0 has moved to pp_emc.

Licensing after the installation

If PowerPath is not registered while installing, you can register it after installation and change the default load balancing and failover policy.

1. Register PowerPath license with the **emcprep -install** command.
 - a. Type **y**
 - b. Enter the 24-character alphanumeric sequence as shown in the License Key Card.

Note: To reduce common typographical errors, the License Key field accepts either uppercase or lowercase letters. Certain numbers and letters are interchangeable. For example, an entry of the alphabetic letters O, I, S, and B is equivalent to an entry of the numbers, 0, 1, 5, and 8.

- c. Follow the on-screen instructions.
2. If PowerPath service is not started, type **/etc/init.d/PowerPath start**, or for RHEL7, OL7, or SLES 12 SP2 type **systemctl start PowerPath.service**. Use the script only to load and unload PowerPath modules.

Note: If the **hald** daemon is running in the system, this command takes longer to complete.

3. Type **powermt set policy** to change the policy if needed. By default, PowerPath will set the appropriate optimal policy based on the license installed.
4. Type **powermt save** to ensure PowerPath configuration settings are persist across host restarts.

The PowerPath management daemon starts along with the PowerPath service and the default settings are considered for the initial startup.

Post-Installation Steps

After installing and licensing PowerPath, configure PowerPath for real-time operation. The configuration enables PowerPath to use the set load balancing and failover policy with the required HBA modules and also to send SNMP traps.

1. If necessary, configure the management daemon to send SNMP traps when path events occur. For steps on configuring, refer to [“Configuring the Management Daemon” on page 65](#).
2. Use the **powermt display dev=all** command and verify that all PowerPath capabilities, including the load balancing and failover policies, are given as expected for your license.

CHAPTER 3

Administering in a Boot from SAN environment

You can configure a PowerPath pseudo device (emcpower) in a Boot from SAN setup to be a root device and provide load balancing and failover of paths for EMC storage systems. You can also upgrade or uninstall PowerPath, or upgrade the Linux kernel in the Boot from SAN environment.

- ◆ Before configuring for a root device 22
- ◆ Configuring PowerPath in a RHEL 5.x Boot from SAN setup 22
- ◆ Configuring PowerPath in a RHEL 6.x or OL 6.x Boot from SAN setup 25
- ◆ Configuring PowerPath in a RHEL 7.x or OL 7.x Boot from SAN setup 28
- ◆ Configuring PowerPath in a SLES 11 SP4 Boot from SAN setup 30
- ◆ Configuring PowerPath in a SLES 12 SP2 Boot from SAN setup 33
- ◆ Upgrading PowerPath in a RHEL 5.x Boot from SAN setup 34
- ◆ Upgrading PowerPath in a RHEL 6.x or OL 6.x Boot from SAN setup..... 35
- ◆ Upgrading PowerPath in a RHEL 7.x, OL 7.x, or SLES 12 SP2 Boot from SAN setup ... 37
- ◆ Upgrading PowerPath in a SLES 11 SP4 Boot from SAN setup 39
- ◆ Upgrading the Linux kernel in a SLES 11 SP4 Boot from SAN setup 40
- ◆ Upgrading the Linux kernel in a Boot from SAN setup 42
- ◆ Uninstalling PowerPath in a RHEL 5.x Boot from SAN setup..... 44
- ◆ Uninstalling PowerPath in a RHEL 6.x or OL 6.x Boot from SAN setup 44
- ◆ Uninstalling PowerPath in a SLES 11 SP4 Boot from SAN setup 47
- ◆ Uninstalling PowerPath in a SLES 12 SP2 Boot from SAN setup 48

Before configuring for a root device

- ◆ For SLES 11, use a PowerPath pseudo (emcpower) device instead of the default `/dev/disk/by-id` device in the `/etc/fstab` file to mount the `/boot` file system.
- ◆ For RHEL 5.x, use the `uuid` in the `/etc/fstab` file to mount the `/boot` file system.

CAUTION

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.

[“Configuring LVM2 support” on page 67](#) provides information on modifying the `/etc/lvm/lvm.conf` file.

Note: For additional information on configuring OVM, see OVM documentation.

Configuring PowerPath in a RHEL 5.x Boot from SAN setup

This section includes procedures for configuring a PowerPath root device using the LVM on a RHEL 5.x host.

Note: Boot from SAN is not supported on OVM 3.x.x. For additional information on configuring OVM, see OVM documentation.

To configure a PowerPath root device using LVM on a RHEL 5.x only host:

1. Ensure that fibre channel cable is plugged only into the first HBA of the host. If the host uses iSCSI HBAs or FCoE adapters for storage connectivity, ensure that only the first port is connected to the network.
2. Enable the BIOS on the first HBA and disable it on all others.
3. Ensure only one LUN is presented to the host.

Note: During configuration, it is mandatory that only one LUN be presented to the host using only a single path. Additional LUNs and paths can be added after the OS installation.

4. Select LUN 0 in the BIOS on the first HBA for boot.
5. Boot the host and start RHEL5.x installation.
6. Create a custom partition setup on the target storage device for two partitions:
 - The `/boot` partition formatted as an ext3 file system with at least 100 MB in size.
 - A second partition of type 0x8E Linux LVM using the remaining space on the storage device.
7. Create a volume group on the LVM partition.

8. From the volume group, create two logical volumes - one for the root file system and the other for swap. Format the root volume as ext3 and the other as swap and designate the mount points for each volume.
9. Ensure that the default boot loader is GRUB and that the boot loader location specifies the MBR of the desired device.
10. Complete the operating system installation.
11. Ensure dm-multipath is not configured. If “multipath -ll” returns configured devices, refer to [“Disabling dm-multipath on a RHEL 5.x host” on page 25](#).
12. Install and configure PowerPath following the instructions in [“Installing the RPM File” on page 16](#) 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.
13. If necessary, save a backup of the /etc/fstab file. For RHEL 5.x, the /etc/fstab should be modified to mount /boot partition by UUID. By default, the /boot partition is configured to mount by LABEL. The following steps show an example of 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/sdp1 99M   22M   73M   23% /boot
# ls -l /dev/disk/by-uuid/ | grep sdp1
# powermt display dev=sdp
Pseudo name=emcpowera
Symmetrix ID=000194901016
Logical device ID=0036
state=alive; policy=SymmOpt; queued-IOs=0

=====
----- Host ----- - Stor - -- I/O Path -- -- Stats ---
### HW Path          I/O Paths   Interf.  Mode   State  Q-IOs Errors
=====
   4 qla2xxx          sdar       FA  7eA  active  alive   0    0
   5 qla2xxx          sdae       FA  7eA  active  alive   0    0
   5 qla2xxx          sdp        FA  8eA  active  alive   0    0

# ls -al /dev/disk/by-uuid|grep emcpowera
lrwxrwxrwx 1 root root 16 Jun 16 07:06 a4ef34a8-4634-4a4a-bd8f-c79026b63224 -> ../../emcpowera
```

The UUID for the partition is:

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

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

```
/dev/system/LogVol00 / 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/LogVol01 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.

14. If necessary, save a backup of the /etc/lvm/lvm.conf file. For RHEL 5, edit the /etc/lvm/lvm.conf file and replace the default lvm filter with the following:

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

In this example, the volume group containing the root logical volume and the swap logical volume is assumed to be on the 2nd partition of the LUN whose native device is /dev/sda and pseudo device is /dev/emcpowera.

15. Add additional paths to the boot LUN, perform a bus scan, and configure the newly added paths in PowerPath using **powermt config**.
16. Add additional LUNs to the host.
17. Perform a bus scan.
18. Configure the newly added LUNs with **powermt config**.
19. Run **powermt save**.
20. Reboot the host, enable the BIOS for the other HBAs and add LUN 0 on those HBAs.

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

- 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.
- 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.

21. As a necessary precaution, perform the following steps and if this problem is encountered, the host can be booted and repaired.
 - a. Make a copy of the original ramdisk image found in /boot folder. Ensure that the LVM filter in the original image accepts all devices:

```
# cp /boot/initrd-2.6.18-194.el5.img
/boot/initrd-2.6.18-194.el5.img.original
```

- b. Create a new entry in /boot/grub/menu.lst file with the original ramdisk image created in the previous step.

```
Red Hat Enterprise Linux Server (2.6.18-194.el5)
root (hd0,0)
kernel /vmlinuz-2.6.18-194.el5 ro
root=/dev/VolGroup00/LogVol100 rhgb quiet
initrd /initrd-2.6.18-194.el5.img
```

```
Red Hat Enterprise Linux Server (2.6.18-194.el5)
root (hd0,0)
kernel /vmlinuz-2.6.18-194.el5 ro root=/dev/VolGroup00/LogVol100
rhgb quiet
initrd /initrd-2.6.18-194.el5.img.original
```

Disabling dm-multipath on a RHEL 5.x host

1. Edit `/etc/fstab` and replace `/dev/mapper/mpath<#>p<#>` by UUID
2. Blacklist all devices in `/etc/multipath.conf`
 - a. Save a copy of `/etc/multipath.conf` as `/etc/multipath.conf.bak` for backup if it already exists.
 - b. Edit `/etc/multipath.conf` so that it contains only the following lines:

```

blacklist {
    devnode "*"
}

```
3. Ensure dm-multipath does not start automatically at boot.

```

# chkconfig multipathd off
# chkconfig --list multipathd    -- - -Service should be off on all
run levels

```
4. Recreate `initrd` so that dm-multipath is excluded from the `initrd`.

```

# mkinitrd /boot/initrd-wo-DM-$(uname -r).img $(uname -r)

```
5. Add a new boot entry with the `initrd` file created in previous step in `/boot/grub/menu.lst` file.
6. Reboot the host with the new `initrd` image and ensure dm-multipath does not have any devices configured.

```

# multipath -ll    --- should not return any dm devices

```

Configuring PowerPath in a RHEL 6.x or OL 6.x Boot from SAN setup

To configure a PowerPath root device using the LVM on a RHEL 6.x host:

1. Ensure that fibre channel cable is plugged only into the first HBA of the host. If the host uses iSCSI HBAs or FCoE adapters for storage connectivity, ensure that only the first port is connected to the network.
2. Enable the BIOS on the first HBA and disable it on all others.
3. Ensure only one LUN is presented to the host, and only one path is available from the host to the array.

Note: During configuration, it is advisable to have only the boot LUN connected to the host. This makes it easier to identify the boot disk during installation.

It is also advisable to have only one path from the array to the host while installing the operating system. This prevents the OS installer from automatically activating native multipathing on the host. If multiple paths are available to the LUNs during installation, native multipathing is activated and extra steps are required to disable it after the OS installation is complete.

4. Select LUN 0 in the BIOS on the first HBA for boot.

5. Boot the host and start RHEL 6.x or OL 6.x installation.
6. Create a custom partition setup on the target storage device for two partitions:
 - The /boot partition formatted as an ext3 or ext4 file system with at least 500 MB in size.
 - A second partition of type 0x8E Linux LVM using the remaining space on the storage device.
7. Create a volume group on the LVM partition.
8. From the volume group, create two logical volumes - one for the root file system and the other for swap. Format the root volume as ext 3 or ext4 and the other as swap and designate the mount points for each volume.
9. 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.
10. Complete the operating system installation.
11. Ensure dm-multipath is not configured. If "multipath -ll" returns configured devices, refer to [“Disabling dm-multipath on a RHEL 6.x or OL 6.x host” on page 27](#).
12. If the OS installation was completed with only a single path boot LUN, add additional paths to the boot LUN and any additional LUNs.
13. Install and configure PowerPath following the instructions in [“Installing PowerPath” on page 13](#) 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.
14. To complete the configuration of the boot LUN under PowerPath, use one of the following methods:

- ◆ Using the emcpbfs_linux script

Note: In a multi-boot environment ensure that the default value in the /boot/grub/menu.lst is set properly with the current boot kernel before running emcpbfs_linux config/remove.

- Run `/etc/opt/emcpower/emcpbfs_linux config` to configure PowerPath.

emcpbfs_linux config adds a new boot entry for the new initrd image in /boot/grub/menu.lst and sets the proper 'default' value. If older entries are not required, manually clean up and specify the proper 'default' value .

- ◆ Using manual steps

- a. If necessary, save a backup of the /etc/fstab file. Then, edit /etc/fstab to mount /boot by pseudo-device name (/dev/emcpowera):

```
#UUID=110223d8-f5f3-4058-812f-16d9e9ee7eab /boot ext4 defaults 1 2
/dev/emcpowera1 /boot ext4 defaults 1 2
```

The following step prevents the ramdisk image from becoming NULL.

- b. Remount /boot so that it is mounted on a pseudo device.

```
# umount /boot/
# mount /boot/
```

- c. If necessary, save a backup of the `/etc/lvm/lvm.conf` file. Then, change the LVM filter in `/etc/lvm/lvm.conf` to accept only pseudo-devices:

```
#filter = [ "a./.*/" ]
filter = [ "a/emcpower.*/", "r/sd.*/", "r/disk.*/" ]
```

- d. Run `vgscan -v` to rebuild the LVM2 cache.
- e. Run `lvmdiskscan` to ensure that the filter field is working correctly. Verify that the filtered device nodes are not listed in the output.
- f. Build a new initramfs image to include PowerPath modules:

```
# dracut /boot/initramfs-PP-$(uname -r).img $(uname -r)
```

This generates a new initramfs image named:

```
initramfs-PP-<kernel version>.<arch>.img
```

- g. Add a new boot entry with initramfs image name in `/boot/grub/menu.lst` created in the previous step and make it the default entry.

```
Red Hat Enterprise Linux (2.6.32-220.el6.x86_64)
root (hd0,0)
kernel /vmlinuz-2.6.32-220.el6.x86_64 ro
  root=/dev/mapper/vg_libl044-lv_root rd_NO_LUKS  KEYBOARDTYPE=pc
  KEYTABLE=us LANG=en_US.UTF-8 rd_NO_MD quiet
  SYSFONT=latarcyrheb-sun16 rhgb crashkernel=auto
  rd_LVM_LV=vg_libl044/lv_swap rd_LVM_LV=vg_libl044/lv_root rd_NO_DM
initrd /initramfs-PP-2.6.32-220.el6.x86_64.img
```

```
Red Hat Enterprise Linux (2.6.32-220.el6.x86_64)
root (hd0,0)
kernel /vmlinuz-2.6.32-220.el6.x86_64 ro
  root=/dev/mapper/vg_libl044-lv_root rd_NO_LUKS  KEYBOARDTYPE=pc
  KEYTABLE=us LANG=en_US.UTF-8 rd_NO_MD quiet
  SYSFONT=latarcyrheb-sun16 rhgb crashkernel=auto
  rd_LVM_LV=vg_libl044/lv_swap rd_LVM_LV=vg_libl044/lv_root rd_NO_DM
initrd /initramfs-2.6.32-220.el6.x86_64.img
```

15. Reboot the host, enable the BIOS for the other HBAs and add LUN 0 on those HBAs. After boot, the new initramfs image that contains PowerPath modules will be loaded.

Note: If the root volume group is expanded later by adding more LUNs to it, then the `initrd` image needs to be updated. Either run `emcpbfs_linux` script with `config` option after deleting the `/etc/emc/.emcpbfs_linux_done` file or execute the following steps manually.

16. Take a backup of the existing `initrd` image.
17. Update the initramfs image by following [step f](#) in the manual steps detailed in the previous procedure.

Disabling dm-multipath on a RHEL 6.x or OL 6.x host

1. Blacklist all devices in `/etc/multipath.conf`
 - a. Save a copy of `/etc/multipath.conf` as `/etc/multipath.conf.bak` for backup if it already exists.

- b. Edit /etc/multipath.conf so that it contains only the following lines:

```
blacklist {
    devnode "*"
}
```

2. Ensure dm-multipath does not start automatically at boot.

```
# chkconfig multipathd off
# chkconfig --list multipathd    -- - -Service should be off on
all run levels
```

3. Recreate initramfs so that dm-multipath is excluded from initramfs.

```
# dracut /boot/initramfs-wo-DM-$(uname -r).img $(uname -r)
```

4. Add a new boot entry with the initramfs file created in previous step in the /boot/grub/menu.lst file.

5. Reboot the host with the new initramfs image and ensure dm-multipath does not have any devices configured.

```
# multipath -ll    --- should not return any dm devices
```

Configuring PowerPath in a RHEL 7.x or OL 7.x Boot from SAN setup

To configure a PowerPath root device using the LVM:

1. Ensure that the fibre channel cable is plugged only into the first HBA of the host. If the host uses iSCSI HBAs or FCoE adapters for storage connectivity, ensure that only the first port is connected to the network.
2. Enable the BIOS on the first HBA and disable it on all others.
3. Ensure only one LUN is presented to the host and only one path is available from the host to the array.

Note: When installing on SLES 12 SP2, ensure native multipathing is not automatically activated during the OS installation, as there is no way to disable native multipathing after the OS installation is complete.

Note: During configuration, it is advisable to have the boot LUN connected to the host to make it easier to identify the boot disk during installation.

It is advisable to have only one path from the array to the host when installing the operating system. This prevents the OS installer from automatically activating native multipathing on the host. If multiple paths are available to the LUNs during installation, native multipathing is automatically activated and extra steps are required to disable it after the OS installation is complete.

4. At installation, select **Basic Storage Devices** as the type of device for installation.
5. Proceed with installation, creating a /boot as basic partition of 500MB (ext4), and Logical Volumes for / (root) and swap on Volume Group created from remaining space of the disk.

6. After OS Installation is complete, add additional paths to the boot LUN and attach additional LUNs to the host.
7. Install and configure PowerPath.
8. Complete the configuration of the boot LUN under PowerPath:
 - To automatically complete the configuration:
 - a. Run `/etc/opt/emcpower/emcpbfs_linux config` to configure PowerPath.
 - To manually complete the configuration:
 - a. Edit `/etc/fstab` to mount `/boot` by pseudo-device name (for example, `/dev/emcpowerX`).


```
#UUID=110223d8-f5f3-4058-812f-16d9e9ee7eab /boot ext4 defaults
1 2
/dev/emcpowera1 /boot ext4 defaults 1 2
```
 - b. Unmount and mount `/boot` so that it is mounted on pseudo device.


```
# umount /boot/
# mount /boot/
```
 - c. Setup the LVM filter in `/etc/lvm/lvm.conf`.


```
global_filter = [ "a/emcpower.*/", "r/sd.*/", "r/disk.*/" ]
```
 - d. Run `vgscan -v` to rebuild the LVM2 cache.
 - e. Run `lvmdiskscan` to ensure that the filter field is working correctly. Verify that the filtered device nodes are not listed in the output.
 - f. Build the initramfs to include PowerPath modules.

For example on RHEL7.x/OL7.x:

```
# dracut /boot/initramfs-PP-$(uname -r).img $(uname -r)
```

For example on SLES 12 SP2:

```
# dracut /boot/initramfs-$(uname -r) $(uname -r)
```

This generates a new initramfs file named `initramfs-PP-<kernelversion>.<arch>.img`
 - g. Add a new boot entry with initramfs file name in `/etc/grub.d/40_custom`.


```
#!/bin/sh
exec tail -n +3 $0
# This file provides an easy way to add custom menu entries.
Simply type the
# menu entries you want to add after this comment. Be careful
not to change
# the 'exec tail' line above.

menuentry 'PowerPath - Red Hat Enterprise Linux Server' {
    load_video
    set gfxpayload=keep
    insmod gzio
    insmod part_msdos
    insmod xfs
    set root='hd0,msdos1'
    if [ x$feature_platform_search_hint = xy ]; then
```

```

        search --no-floppy --fs-uuid --set=root
--hint-bios=hd8,msdos1 --hint-efi=hd8,msdos1
--hint-baremetal=ahci8,msdos1 --hint='hd0,msdos1'
c556a69a-6af5-4ab7-b973-ed3bde2208ae
    else
        search --no-floppy --fs-uuid --set=root
c556a69a-6af5-4ab7-b973-ed3bde2208ae
    fi
    linux16 /vmlinuz-3.10.0-121.el7.x86_64
root=UUID=f78a46ff-10d0-4bfb-a2a5-9e451e03a6d9 ro
rd.lvm.lv=rhel/root crashkernel=auto rd.lvm.lv=rhel/swap
vconsole.font=latarcyrheb-sun16 vconsole.keymap=us rhgb quiet
LANG=en_US.UTF-8
    initrd16 /initramfs-PP-3.10.0-121.el7.x86_64.img

```

9. Update the new entry in the grub.cfg file and make this entry as default so that system boots up with the new initrd image automatically.

```

# grub2-mkconfig -o /boot/grub2/grub.cfg
# grub2-set-default 'PowerPath - Red Hat Enterprise Linux Server'

```

10. Reboot the host to have the new initramfs image that has PowerPath modules takes effect.

Disabling dm-multipath on a RHEL 7.x or OL 7.x host

1. Blacklist all devices in /etc/multipath.conf
 - a. Save a copy of /etc/multipath.conf as /etc/multipath.conf.bak for backup if it already exists.
 - b. Edit /etc/multipath.conf so that it contains only the following lines:

```

blacklist {
    devnode "*"
}

```

2. Ensure dm-multipath does not start automatically at boot.

```
# systemctl disable multipathd
```

3. Recreate initramfs so that dm-multipath is excluded from initramfs.

```
# dracut /boot/initramfs-wo-DM-$(uname -r).img $(uname -r)
```

4. Add a new boot entry with the initramfs file created in previous step in the /etc/grub.d/40_custom file, and then update the grub.cfg file.
5. Reboot the host with the new initramfs image and ensure dm-multipath does not have any devices configured.

```
# multipath -ll --- should not return any dm devices
```

Configuring PowerPath in a SLES 11 SP4 Boot from SAN setup

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

1. Ensure that the fibre channel cable is plugged only into the first HBA of the host. If the host uses iSCSI HBAs or FCoE adapters for storage connectivity, ensure that only the first port is connected to the network.

2. Enable the BIOS on the first HBA and disable it on all others.
3. Ensure only one LUN is presented to the host and only one path is available from the host to the array.

Note: During configuration, it is advisable to have the boot LUN connected to the host to make it easier to identify the boot disk during installation.

It is advisable to have only one path from the array to the host when installing the operating system. This prevents the OS installer from automatically activating native multipathing on the host. If multiple paths are available to the LUNs during installation, native multipathing is activated and extra steps are required to disable it after the OS installation is complete.

4. Select LUN 0 in the BIOS on the first HBA for boot.
5. Boot the host and start SLES installation.
6. 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.
7. Create a volume group on the LVM partition.
8. From the volume group, create two logical volumes - one for the root file system and the other for swap. Format the root volume as ext3 and the other as swap and designate the mount points for each volume.
9. 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.
10. Complete the operating system installation.
11. Ensure dm-multipath is not configured. If “multipath -ll” returns configured devices, refer to [“Disabling dm-multipath to prepare for PowerPath installation on SLES 11 SP4” on page 33](#).
12. If the OS installation was completed with only the boot LUN with a single path, add any additional paths to the boot LUN and any additional LUNs.
13. Install and configure PowerPath following the instructions in [“Installing PowerPath” on page 13](#) 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*.
14. To complete the configuration of the boot LUN under PowerPath, use one of the following methods:
 - ◆ Using the emcpbfs_linux config script

Note: In a multi-boot environment ensure that the default value in the /boot/grub/menu.lst is set properly with the current boot kernel before running emcpbfs_linux config/remove.

- Run `/etc/opt/emcpower/emcpbfs_linux config` to configure PowerPath.

Note: `emcpbfs_linux` config adds a new boot entry for the new `initrd` image in `/boot/grub/menu.lst` and sets the proper 'default' value. If older entries are not required, manually clean up and specify the proper 'default' value.

- ◆ Using manual steps
 - a. If necessary, save a backup of the `/etc/fstab` file. Then, 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-3600601605c7019002820bd32b767dell-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
```

Change 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-3600601605c7019002820bd32b767dell-part1
/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
```

- b. If necessary, save a backup of the `/etc/lvm/lvm.conf` file. Then, change the LVM filter in `/etc/lvm/lvm.conf` to accept only pseudo-devices:

```
# filter = [ "a./.*/" ]
filter = [ "a/emcpower.*/", "r/sd.*/", "r/disk.*/" ]
```

- c. Run `vgscan -v` to rebuild the LVM2 cache.
- d. Run `lvmdiskscan` to ensure that the filter field is working correctly. Verify that the filtered device nodes are not listed in the output.
- e. Create a new `initrd` image with the PowerPath feature to include PowerPath modules.

```
mkinitrd -i /boot/initrd-pp-3.0.76-0.9-default -k
/boot/vmlinuz-3.0.76-0.9-default -f powerpath
```

- f. Add a new boot entry for the new `initrd` image created in the previous step in `/boot/grub/menu.lst` and make it the default entry.

```
SUSE Linux Enterprise Server 11 SP3 - 3.0.76-0.9
root (hd0,0)
kernel /vmlinuz-3.0.76-0.9-default root=/dev/system/root
resume=/dev/system/swap splash=silent crashkernel=256M-:128M
showopts
```

```

initrd /initrd-pp-3.0.76-0.9-default

SUSE Linux Enterprise Server 11 SP3 - 3.0.76-0.9
root (hd0,0)
kernel /vmlinuz-3.0.76-0.9-default root=/dev/system/root
resume=/dev/system/swap splash=silent crashkernel=256M-:128M
showopts
initrd /initrd-3.0.76-0.9-default

```

15. Reboot the host, enable the BIOS on the other HBAs and select LUN 0 on those HBAs. After boot, the new initrd image that contains PowerPath modules will be loaded.

Note: If the root volume group is expanded later by adding more LUNs to it, then the initrd image needs to be updated. Either run `emcpbfs_linux` script with config option after deleting the `/etc/emc/.emcpbfs_linux_done` file or execute the following steps manually.

1. Take a backup of the existing initrd image.
2. Update the initrd image by repeating [step e](#) in the manual steps detailed in the previous procedure.

Disabling dm-multipath to prepare for PowerPath installation on SLES 11 SP4

1. Blacklist all devices in `/etc/multipath.conf`
 - a. Save a copy of `/etc/multipath.conf` as `/etc/multipath.conf.bak` for backup if it already exists.
 - b. Edit `/etc/multipath.conf` so that it contains only the following lines:

```

blacklist {
    devnode "*"
}

```

2. Ensure dm-multipath does not start automatically at boot.

```

# chkconfig multipathd off
# chkconfig --list multipathd      --- Service should be off on all
run levels

```

3. Recreate initrd so that dm-multipath is excluded from initrd.

```

mkinitrd -i /boot/initrd-wo-DM-3.0.76-0.9-default -k
/boot/vmlinuz-3.0.76-0.9-default

```

4. Add a new boot entry with the initrd file created in the previous step in `/boot/grub/menu.lst` file.
5. Reboot the host with the new initrd image and ensure dm-multipath has not configured any device.

```

# multipath -ll      --- should not return any dm devices

```

Configuring PowerPath in a SLES 12 SP2 Boot from SAN setup

1. Configure PowerPath on a Boot from SAN host.
- ```

/etc/opt/emcpower/emcpbfs_linux config

```

For example,

```
/etc/opt/emcpower/emcpbfs_linux config

PHASE 1: VERIFICATION PHASE

Verifying PowerPath Installed
Verifying Boot LUN emcpowera:libg210_BFS_0
Verifying Root LVM LUN(s) emcpowera:libg210_BFS_0
Verifying PPME Migrations No migration
Verifying /etc/fstab /boot found
Verifying /boot space Available
Verifying /etc/lvm/lvm.conf Filter found
Verifying /boot/grub2/grub.cfg Validated

PHASE 2: CONFIGURATION PHASE

Updating /etc/fstab Done
Updating /etc/lvm/lvm.conf Done
Creating /boot/initrd-PP-3.12.49-11-default Done
Updating /boot/grub2/grub.cfg Done

INFO: Successfully Completed.
Restart the server for the configuration to take effect.

#
```

If the script fails to complete successfully while creating the initrd image or updating the grub.cfg file, run **mount -a** to mount all of the file systems and then run **/etc/opt/emcpower/emcpbfs\_linux config**. For all other causes of failure, fix those issues as indicated by the failure message.

2. Verify that an initrd image with PP in its name is created in /boot.

```
ls -l /boot/*PP*
```

For example,

```
ls -l /boot/*PP*
-rw-r--r-- 1 root root 20320656 Jul 18 03:07 /boot/initrd-
PP-3.12.59-60.41-default
#
```

3. Verify that a grub entry for starting the host with PowerPath is added to grub.cfg.

```
grep PowerPath /boot/grub2/grub.cfg
```

For example,

```
grep PowerPath /boot/grub2/grub.cfg
menuentry 'PowerPath, Linux 3.12.59-60.41-default' --class sles
--class gnu-linux --class gnu --class os $menuentry_id_option
'gnulinux-simple-205d1894-e29f-4bd8-afc3-f575b195e222' {
```

If either step 2 or step 3 fails, then remove the marker file with **rm /etc/emc/.emcpbfs\_linux\_done** and then restart this procedure.

## Upgrading PowerPath in a RHEL 5.x Boot from SAN setup

1. Perform one of the following steps.
  - For RHEL 6.x, rename **/etc/init.d/PowerPath** to **/etc/init.d/PowerPath.old**
  - For RHEL7, OL7, or SLES 12 SP2 run the following commands:

```
systemctl disable PowerPath.service
```

```
systemctl disable PowerPathPost-start.service
systemctl disable EmcpMond.service
```

2. Edit the `/etc/fstab` file to comment out any references to PowerPath pseudo (emcpower) devices and enable (or add) the native device entries.
3. Comment out the currently used filter, restore default filter in `/etc/lvm/lvm.conf` such that LVM filter accepts all devices.

```
filter = ["a./.*/"]
#filter = ["a/sda2$/", "r/emcpowera2/", "r/sd.*/", "r/disk.*/",
"a./.*/"]
```

4. Reboot the host.
5. After the host boots, run `rpm -Uvh <new_powerpath_package>` to upgrade PowerPath.
6. Revert the LVM filter change done in step 3:

```
#filter = ["a./.*/"]
filter = ["a/sda2$/", "r/emcpowera2/", "r/sd.*/", "r/disk.*/",
"a./.*/"]
```

7. Revert the changes to the `/etc/fstab` file made in step 2.
8. Run `/etc/init.d/PowerPath start` (RHEL 6.x) or `systemctl start PowerPath.service` (RHEL7.x, OL7, or SLES 12 SP2) to start PowerPath.

Note: If issues are observed during or after starting PowerPath and if the old PowerPath version needs to be restored, the following steps should be done to preserve the saved configuration settings:

- a. Copy contents of `/etc/emc/archive` to another folder such as `/tmp/pp_config_backup`.
- b. Uninstall the new version of PowerPath.
- c. Delete contents of `/etc/emc/archive` folder since that now contains config files updated by the new PowerPath version.
- d. Copy the backed up contents from step a to `/etc/emc/archive` folder.
- e. Start PowerPath. While starting, PowerPath will automatically restore the config files from `/etc/emc/archive` and the old settings will be restored.

While upgrading, PowerPath uses the license key entered previously. Therefore, re-registration is not required.

## Upgrading PowerPath in a RHEL 6.x or OL 6.x Boot from SAN setup

To upgrade PowerPath in a boot-from-SAN setup on a RHEL6.x or OL 6.x host, use one of the following methods:

- ◆ [“Upgrade using emcpbfs\\_linux script”](#)
- ◆ [“Upgrade using manual steps”](#)

### Upgrade using emcpbfs\_linux script

1. Run `/etc/opt/emcpower/emcpbfs_linux remove` to prepare the host for PowerPath upgrade.

If the **ERROR: PowerPath is not configured on Boot-From-SAN.** message is displayed, then it indicates that `emcpbfs_linux` config has not been used during PowerPath configuration. Follow the steps to manually upgrade or create the file below and re-run the script.

```
touch /etc/emc/.emcpbfs_linux_done
```

2. Comment all other entries in the `/etc/fstab` file that refer to PowerPath pseudo (emcpower) devices.

---

**Note:** `emcpbfs_linux` script modifies `'/boot'` entry only in the `/etc/fstab` file

---

3. Reboot the host.
4. Upgrade PowerPath.
5. Run `/etc/opt/emcpower/emcpbfs_linux config` to configure PowerPath.
6. Reboot the host.

---

**Note:** `emcpbfs_linux config/remove` adds a new boot entry for the new `initramfs` image in `/boot/grub/menu.lst` and sets the proper 'default' value.

---

## Upgrade using manual steps

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

```
UUID=110223d8-f5f3-4058-812f-16d9e9ee7eab /boot ext4 defaults 1 2
#/dev/emcpowera1 /boot ext4 defaults 1 2
```

3. Comment out the currently used filter, restore default filter in `/etc/lvm/lvm.conf` such that LVM filter accepts all devices.

```
filter = ["a./."]
#filter = ["a/emcpower./.", "r/sd./.", "r/disk./."]
```

4. Edit `/boot/grub/menu.lst` and make the entry that contains the `initramfs` image without the PowerPath modules as the default entry.
5. Reboot the host.
6. After the host boots, run `rpm -Uvh <new_powerpath_package>` to upgrade PowerPath.
7. Run `/etc/init.d/PowerPath start` to start PowerPath.
8. Edit `/etc/fstab` to mount `/boot` by pseudo (emcpower) devices.

```
#UUID=110223d8-f5f3-4058-812f-16d9e9ee7eab /boot ext4 defaults 1 2
/dev/emcpowera1 /boot ext4 defaults 1 2
```

9. Remount `/boot` so that it is mounted on a pseudo device.

```
umount /boot/
mount /boot/
```

10. Setup the LVM filter in `/etc/lvm/lvm.conf` to accept pseudo (emcpower) devices:

```
filter = ["a./.*/"]
filter = ["a/emcpower.*/", "r/sd.*/", "r/disk.*/"]
```

11. Delete the old initramfs containing PowerPath modules that was created in step 12 in "[Configuring PowerPath in a RHEL 6.x or OL 6.x Boot from SAN setup](#)" on page 25

```
rm -f /boot/initramfs-PP-$(uname -r).img
```

12. Build the initramfs to include upgraded PowerPath modules:

```
dracut /boot/initramfs-PP-$(uname -r).img $(uname -r)
```

This generates a new initramfs file named:

```
initramfs-PP-<kernel version>.<arch>.img
```

13. The /boot/grub/menu.lst already contains an entry with initramfs created during initial setup, see "[Configuring PowerPath in a RHEL 6.x or OL 6.x Boot from SAN setup](#)" on page 25". Make it the default entry. Add a new boot entry with initramfs image containing PowerPath modules, if it is missing.
14. Reboot the host so that the new initramfs image containing updated PowerPath modules takes effect.

## Upgrading PowerPath in a RHEL 7.x, OL 7.x, or SLES 12 SP2 Boot from SAN setup

To upgrade PowerPath in a boot-from-SAN setup on a RHEL7.x, OL 7.x and SLES 12 SP2 host, use one of the following methods:

- ◆ [“Upgrade using emcpbfs\\_linux script”](#)
- ◆ [“Upgrade using manual steps”](#)

### Upgrade using emcpbfs\_linux script

1. Run `/etc/opt/emcpower/emcpbfs_linux remove` to prepare the host for PowerPath upgrade.

If the **ERROR: PowerPath is not configured on Boot-From-SAN.** message is displayed, then it indicates that emcpbfs\_linux config has not been used during PowerPath configuration. Follow the steps to manually upgrade or create the file below and re-run the script.

```
touch /etc/emc/.emcpbfs_linux_done
```

2. Comment all other entries in the /etc/fstab file that refer to PowerPath pseudo (emcpower) devices.

---

**Note:** `emcpbfs_linux` script modifies '/boot' entry only in the /etc/fstab file

---

3. Reboot the host.
4. Upgrade PowerPath.
5. Run `/etc/opt/emcpower/emcpbfs_linux config` to configure PowerPath.
6. Reboot the host.

**Note:** `emcpbfs_linux config/remove` adds a new boot entry for the new initramfs image in `/boot/grub/menu.lst` and sets the proper 'default' value.

## Upgrade using manual steps

1. Disable PowerPath service across host reboot.

```
systemctl disable PowerPath.service
systemctl disable PowerPathPost-start.service
systemctl disable EmcpMond.service
```

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

```
UUID=110223d8-f5f3-4058-812f-16d9e9ee7eab /boot ext4 defaults 1 2
#/dev/emcpowera1 /boot ext4 defaults 1 2
```

3. Comment out the currently used filter, restore default filter in `/etc/lvm/lvm.conf` such that LVM filter accepts all devices.

```
global_filter = ["a./."]
#global_filter = ["a/emcpower./", "r/sd./", "r/disk./"]
```

4. Set the default GRUB entry that contains the initramfs image without the PowerPath modules.

```
#grub2-set-default "Red Hat Enterprise Linux Server"
```

5. Reboot the host.

6. After the host boots, run `rpm -Uvh <new_powerpath_package>` to upgrade PowerPath.

7. Run `systemctl start PowerPath.service` to start PowerPath.

8. Edit `/etc/fstab` to mount `/boot` by pseudo (emcpower) devices.

```
#UUID=110223d8-f5f3-4058-812f-16d9e9ee7eab /boot ext4 defaults 1 2
#/dev/emcpowera1 /boot ext4 defaults 1 2
```

9. Remount `/boot` so that it is mounted on a pseudo device.

```
umount /boot/
mount /boot/
```

10. Setup the LVM filter in `/etc/lvm/lvm.conf` to accept pseudo (emcpower) devices:

```
#global_filter = ["a./."]
global_filter = ["a/emcpower./", "r/sd./", "r/disk./"]
```

11. Delete the old initramfs containing PowerPath modules that was created in step 12 in [“Configuring PowerPath in a RHEL 7.x or OL 7.x Boot from SAN setup”](#).

```
rm -f /boot/initramfs-PP-$(uname -r).img
```

12. Build the initramfs to include upgraded PowerPath modules.

For example on RHEL 7.x/OL 7.x:

```
dracut /boot/initramfs-PP-$(uname -r).img $(uname -r)
```

For example on SLES 12 SP2:

```
dracut /boot/initramfs-$(uname -r) $(uname -r)
```

13. The `/etc/grub.d/40_custom` already contains an entry with `initramfs` created during initial setup, see [“Configuring PowerPath in a RHEL 7.x or OL 7.x Boot from SAN setup”](#). Make it the default entry. Add a new boot entry with `initramfs` image containing PowerPath modules, if it is missing.
14. Reboot the host so that the new `initramfs` image containing updated PowerPath modules takes effect.

## Upgrading PowerPath in a SLES 11 SP4 Boot from SAN setup

To upgrade PowerPath in a SLES 11 SP4 Boot from SAN setup:

1. Run `/etc/opt/emcpower/emcpbfs_linux remove` to prepare the host for PowerPath upgrade.

If the **ERROR: PowerPath is not configured on Boot-From-SAN.** message is displayed, then it indicates that `emcpbfs_linux` config has not been used during PowerPath configuration. Follow the steps to manually upgrade or create the file below and re-run the script.

```
touch /etc/emc/.emcpbfs_linux_done
```

2. Comment all other entries in the `/etc/fstab` file that refer to PowerPath pseudo (emcpower) devices.

---

**Note:** `emcpbfs_linux` script modifies `'/boot'` entry only in the `/etc/fstab` file

---

3. Reboot the host.
4. Upgrade PowerPath.
5. Run `/etc/opt/emcpower/emcpbfs_linux config` to configure PowerPath.
6. Reboot the host.

`emcpbfs_linux` config adds a new boot entry for the new `initrd` image in `/boot/grub/menu.lst` and sets the proper ‘default’ value.

### Upgrade using manual steps (RHEL 6.x)

1. Rename `/etc/init.d/PowerPath` to `/etc/init.d/PowerPath.old`
2. 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 ac1,user_xattr 1 2
#/dev/emcpowera1 /boot ext3 ac1,user_xattr 1 2
```

3. Comment out the currently used filter, restore default filter in `/etc/lvm/lvm.conf` such that LVM filter accepts all devices.

```
filter = ["a./."]
#filter = ["a/emcpower./.", "r/sd./.", "r/disk./."]
```

4. Edit `/boot/grub/menu.lst` and make the entry that contains the `initrd` image without the PowerPath modules as the default entry.

5. Reboot the host.
6. After the host boots, run `rpm -Uvh <new_powerpath_package>` to upgrade PowerPath.
7. Run `/etc/init.d/PowerPath start` to start PowerPath.
8. Edit `/etc/fstab` to mount `/boot` by pseudo (emcpower) devices.

```
#/dev/disk/by-id/scsi-3600601605c7019002820bd32b767de11-part1
/boot ext3 ac1,user_xattr 1 2
/dev/emcpowera1 /boot ext3 ac1,user_xattr 1 2
```

9. Remount `/boot` so that it is mounted on a pseudo device.

```
umount /boot/
mount /boot/
```

10. Setup the LVM filter in `/etc/lvm/lvm.conf` to accept pseudo (emcpower) devices:

```
filter = ["a./.*/"]
filter = ["a/emcpower.*/", "r/sd.*/", "r/disk.*/"]
```

11. Delete the old initrd image containing PowerPath modules that was created in step 12 in [“Configuring PowerPath in a SLES 11 SP4 Boot from SAN setup” on page 30](#).

```
rm -f /boot/initrd-pp-3.0.13-0.27-default
```

12. Create a new initrd image to include upgraded PowerPath modules:

```
mkinitrd -i /boot/initrd-pp-3.0.76-0.9-default -k
/boot/vmlinuz-3.0.76-0.9-default -f powerpath
```

13. The `/boot/grub/menu.lst` already contains an entry with initrd image created during initial setup, see [“Configuring PowerPath in a SLES 11 SP4 Boot from SAN setup” on page 30](#). Make it the default entry. Add a new boot entry with the initrd image containing PowerPath modules, if it is missing.

14. Reboot the host so that the new initrd image containing updated PowerPath modules takes effect.

## Upgrading the Linux kernel in a SLES 11 SP4 Boot from SAN setup

To upgrade the Linux kernel in a SLES 11 SP4 boot from SAN setup, use one of the following methods:

- ◆ [“Upgrade using emcpbfs\\_linux script”](#)
- ◆ [“Upgrade using manual steps”](#)

### Upgrade using emcpbfs\_linux script

1. Run `/etc/opt/emcpower/emcpbfs_linux upgrade` to prepare the host for kernel upgrade.

If the **ERROR: PowerPath is not configured on Boot-From-SAN.** message is displayed, then it indicates that `emcpbfs_linux` config has not been used during PowerPath configuration. Follow the steps to manually upgrade or create the file below and re-run the script.

```
touch /etc/emc/.emcpbfs_linux_done
```

2. Comment all other entries in the `/etc/fstab` file that refer to PowerPath pseudo (emcpower) devices.

**Note:** `emcpbfs_linux` script modifies the `/boot` entry only in the `/etc/fstab` file

3. Upgrade the kernel as described in Novell documentation and reboot the host.
4. Run `/etc/opt/emcpower/emcpbfs_linux config` to configure PowerPath.
5. Reboot the host.

`emcpbfs_linux config` adds a new boot entry for the new `initrd` image in `/boot/grub/menu.lst` and sets the proper `'default'` value. If older entries are not required, manually cleanup and specify the proper `'default'` value.

## Upgrade using manual steps

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
/dev/emcpowera1 /boot ext3 acl,user_xattr 1 2
/dev/emcpowera1 /boot ext3 acl,user_xattr 1 2
```

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

```
filter = ["a./."]
filter = ["a/emcpower.*/", "r/sd.*/", "r/disk.*/"]
```

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
/dev/emcpowera1 /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 = ["a/emcpower.*/", "r/sd.*/", "r/disk.*/"]
```

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 features.

```
mkinitrd -i /boot/initrd-pp-3.0.76-0.9-default -k
/boot/vmlinuz-3.0.76-0.9-default -f powerpath
```

9. Add a new boot entry which has PowerPath `initrd` image (`initrd-pp-3.0.13-0.27-default`) in `/boot/grub/menu.lst` and make it the default.

```
SUSE Linux Enterprise Server 11 SP3 - 3.0.76-0.9
root (hd0,0)
```

```
kernel /vmlinuz-3.0.76-0.9-default root=/dev/system/root
resume=/dev/system/swap splash=silent crashkernel=256M-:128M
showopts
initrd /initrd-pp-3.0.76-0.9-default

SUSE Linux Enterprise Server 11 SP3 - 3.0.76-0.9
root (hd0,0)
kernel /vmlinuz-3.0.76-0.9-default root=/dev/system/root
resume=/dev/system/swap splash=silent crashkernel=256M-:128M
showopts
initrd /initrd-3.0.76-0.9-default
```

10. Reboot the host.

## Upgrading the Linux kernel in a Boot from SAN setup

---

**Note:** Before upgrading the kernel on RHEL5, if the root filesystem is on a Logical Volume Manager (LVM), change the LVM filter to accept all SD devices.

---

When upgrading a kernel on a host using an online option such as Yum, Satellite, or Yast on a RHEL, OL, or SLES 12 SP2 host, PowerPath will automatically reboot and start after the upgrade is complete without requiring any extra steps.

When upgrading a kernel on a host using an offline option such as booting the host from a CD, or if the host is running SLES11 SP4, then the following steps need to be completed to ensure that PowerPath will function normally after the upgrade:

- ◆ [“Upgrade using emcpbfs\\_linux script”](#)
- ◆ [“Upgrade using manual steps”](#)

### Upgrade using emcpbfs\_linux script

1. Run `/etc/opt/emcpower/emcpbfs_linux upgrade` to prepare the host for kernel upgrade.

If the **ERROR: PowerPath is not configured on Boot-From-SAN.** message is displayed, then it indicates that emcpbfs\_linux config has not been used during PowerPath configuration. Follow the steps to manually upgrade or create the file below and re-run the script.

```
touch /etc/emc/.emcpbfs_linux_done
```

2. Comment all other entries in the `/etc/fstab` file that refer to PowerPath pseudo (emcpower) devices.

---

**Note:** `emcpbfs_linux` script modifies `/boot` entry only in the `/etc/fstab` file

---

3. Upgrade the kernel as described in RedHat documentation and reboot the host.
4. Run `/etc/opt/emcpower/emcpbfs_linux config` to configure PowerPath.
5. Reboot the host.

`emcpbfs_linux config/remove` adds a new boot entry for the new initrd image in `/boot/grub/menu.lst` and sets the proper `'default'` value. If older entries are not required, manually cleanup and specify the proper `'default'` value.

## Upgrade using manual steps

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/mapper/vg_lib1048-Lib1048_Root / ext4 defaults 1 1
UUID=8321d7bc-6f48-4359-a0a3-26887d39e0ad /boot ext4 defaults 1 2
#/dev/emcpowera1 /boot ext4 defaults 1 2
/dev/mapper/vg_lib1048-Swap swap swap defaults 0 0
tmpfs /dev/shm tmpfs defaults 0 0
devpts /dev/pts devpts gid=5,mode=620 0 0
sysfs /sys sysfs defaults 0 0
proc /proc proc defaults 0 0
```

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

```
filter = ["a./.*/"]
filter = ["a/emcpower.*/", "r/sd.*/", "r/disk.*/"]
```

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 Oracle or Red Hat documentation and then reboot the host.
5. Run `/etc/init.d/PowerPath start`, or for RHEL7, OL7, or SLES 12 SP2 type `systemctl start PowerPath.service` to start PowerPath.

Note: If issues are observed during or after starting PowerPath and if the old PowerPath version needs to be restored, the following steps should be done to preserve the saved configuration settings:

- a. Copy contents of `/etc/emc/archive` to another folder such as `/tmp/pp_config_backup`.
  - b. Uninstall the new version of PowerPath.
  - c. Delete contents of `/etc/emc/archive` folder since that now contains config files updated by the new PowerPath version.
  - d. Copy the backed up contents from step a to `/etc/emc/archive` folder.
  - e. Start PowerPath. While starting, PowerPath will automatically restore the config files from `/etc/emc/archive` and the old settings will be restored.
6. Edit the `/etc/fstab` file to mount the `/boot` partition on a PowerPath pseudo device.

```
/dev/mapper/vg_lib1048-Lib1048_Root / ext4 defaults 1 1
#UUID=8321d7bc-6f48-4359-a0a3-26887d39e0ad /boot ext4 defaults 1 2
/dev/emcpowera1 /boot ext4 defaults 1 2
/dev/mapper/vg_lib1048-Swap swap swap defaults 0 0
tmpfs /dev/shm tmpfs defaults 0 0
devpts /dev/pts devpts gid=5,mode=620 0 0
sysfs /sys sysfs defaults 0 0
proc /proc proc defaults 0 0
```

7. Remount /boot so that it is mounted on a pseudo device.

```
umount /boot/
mount /boot/
```

8. Change the LVM filter in /etc/lvm/lvm.conf to accept only pseudo-devices:

```
filter = ["a./."]
filter = ["a/emcpower.*/", "r/sd.*/", "r/disk.*/"]
```

9. Build a new initramfs image to include PowerPath modules:

```
dracut /boot/initramfs-PP-$(uname -r).img $(uname -r)
```

10. Add a new boot entry which has the new initramfs image built in step 9 in /boot/grub/menu.lst and make it the default entry.

11. Reboot the host.

## Uninstalling PowerPath in a RHEL 5.x Boot from SAN setup

Also see the section [“Uninstalling PowerPath” on page 86](#).

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 /etc/lvm/lvm.conf file and modify the default lvm filter to accept all devices.
4. Reboot the host.
5. Uninstall PowerPath.

## Uninstalling PowerPath in a RHEL 6.x or OL 6.x Boot from SAN setup

To uninstall PowerPath from a Boot from SAN setup on RHEL 6.x or OL 6.x hosts, use one of the following methods:

- ◆ [“Uninstall using emcpbfs\\_linux\\_script”](#)
- ◆ [“Uninstall using manual steps”](#)

### Uninstall using emcpbfs\_linux\_script

---

**Note:** In a multi-boot environment ensure that the default value in the /boot/grub/menu.lst is set properly with the current boot kernel before running emcpbfs\_linux config/remove.

---

1. Run `/etc/opt/emcpower/emcpbfs_linux remove` to uninstall PowerPath.

If the **ERROR: PowerPath is not configured on Boot-From-SAN.** message is displayed, then it indicates that emcpbfs\_linux config has not been used during PowerPath configuration. Follow the steps to manually upgrade or create the file below and re-run the script.

```
touch /etc/emc/.emcpbfs_linux_done
```

2. Comment all other entries in the `/etc/fstab` file that refer to PowerPath pseudo (emcpower) devices.

---

**Note:** `emcpbfs_linux` script modifies `/boot` entry only in the `/etc/fstab` file

---

3. Reboot the host.

---

**Note:** `emcpbfs_linux config/remove` adds a new boot entry for the new initrd image in `/boot/grub/menu.lst` and sets the proper `'default'` value. If older entries are not required, manually clean up and specify the proper `'default'` value .

---

4. Uninstall PowerPath.

## Uninstall using manual steps

1. Move the `/etc/init.d/PowerPath` script to a different location.
2. Edit the `/etc/lvm/lvm.conf` file to restore the filter and accept SCSI device nodes.
3. Edit `/etc/fstab` to restore the original option of mounting `/boot` by UUID.

```
UUID=110223d8-f5f3-4058-812f-16d9e9ee7eab /boot ext4 defaults 1 2
#/dev/emcpowera1 /boot ext4 defaults 1 2
```

4. Run `vgscan -v` and then verify the filter by running the `lvmdiskscan` command.
5. Reboot the host using the original initramfs.

If the original initramfs is missing, re-create it using the following steps:

- a. Move dracut configuration files for PowerPath out of `/usr/share/dracut/modules.d/` by entering the following:

```
mv /usr/share/dracut/modules.d/90powerpath ~/dracut/
```

- b. Generate the initramfs file:

```
dracut /boot/initramfs-PP-$(uname -r).img $(uname -r)
```

- c. Add a new boot entry for the new initramfs image created `/boot/grub/menu.lst`.
- d. Reboot the host and select the boot loader entry with initramfs image in the previous step.

6. Uninstall PowerPath.

## Uninstalling PowerPath in a RHEL 7.x or OL 7.x Boot from SAN setup

Use one of the following to uninstall PowerPath from a Boot from SAN setup on RHEL 7.x or OL 6.x host:

- ◆ [“Uninstall using emcpbfs\\_linux\\_script”](#)
- ◆ [“Uninstalling using manual steps”](#)

## Uninstall using emcpbfs\_linux\_script

**Note:** In a multi-boot environment ensure that the default value in the `/boot/grub/menu.lst` is set properly with the current boot kernel before running `emcpbfs_linux config/remove`.

1. Run `/etc/opt/emcpower/emcpbfs_linux remove` to uninstall PowerPath.

If the **ERROR: PowerPath is not configured on Boot-From-SAN.** message is displayed, then it indicates that `emcpbfs_linux config` has not been used during PowerPath configuration. Follow the steps to manually upgrade or create the file below and re-run the script.

```
touch /etc/emc/.emcpbfs_linux_done
```

2. Comment all other entries in the `/etc/fstab` file that refer to PowerPath pseudo (emcpower) devices.

**Note:** `emcpbfs_linux` script modifies `/boot` entry only in the `/etc/fstab` file

3. Reboot the host.

**Note:** `emcpbfs_linux config/remove` adds a new boot entry for the new initrd image in `/boot/grub/menu.lst` and sets the proper `'default'` value. If older entries are not required, manually clean up and specify the proper `'default'` value .

Uninstall PowerPath.

## Uninstalling using manual steps

1. Disable the following PowerPath services:

```
systemctl disable PowerPath.service
systemctl disable PowerPathPost-start.service
systemctl disable EmcpMond.service
```

2. Edit `/etc/fstab` to restore the original option for mounting `/boot` by UUID.

```
UUID=b4ca11e-9aa1-4777-bd5a-012e827514b1 /boot xfs defaults 1 2
#/dev/emcpowera1 /boot xfs defaults 1 2
```

3. Edit the `/etc/lvm/lvm.conf` file to disable the `global_filter`.

4. If the original `initramfs` image is missing, re-create it using the following steps:

- a. Move `dracut` configuration files for PowerPath out of `/usr/lib/dracut/modules.d/`

```
mv /usr/lib/dracut/modules.d/90powerpath ~/dracut/
```

- b. Generate the new `initramfs` file.

For example on RHEL7.x/OL7.x:

```
dracut /boot/initramfs-PP-$(uname -r).img $(uname -r)
```

For example on SLES 12 SP2:

```
dracut /boot/initramfs-$(uname -r) $(uname -r)
```

- Remove the entry from the 40\_custom file in the /etc/grub.d folder.

```
[root@lib1038 ~]# cat /etc/grub.d/40_custom
#!/bin/sh
exec tail -n +3 $0
This file provides an easy way to add custom menu entries. Simply
type the
menu entries you want to add after this comment. Be careful not
to change
the 'exec tail' line above.
```

- Recreate the /boot/grub2/grub.cfg file to remove the PowerPath boot entry.

```
grub2-mkconfig -o /boot/grub2/grub.cfg
```

- Set the original boot entry as the default entry.

```
grub2-set-default 'Red Hat Enterprise Linux Server, with Linux
3.10.0-121.el7.x86_64'
```

- Reboot the host and then uninstall PowerPath.

## Uninstalling PowerPath in a SLES 11 SP4 Boot from SAN setup

Uninstall PowerPath using one of the following methods:

- ◆ [“Uninstall using emcpbfs\\_linux\\_script”](#)
- ◆ [“Uninstall using manual steps”](#)

### Uninstall using emcpbfs\_linux\_script

---

**Note:** In a multi-boot environment, ensure that the default value in the /boot/grub/menu.lst is set properly with the current boot kernel before running emcpbfs\_linux config/remove.

---

- Run `/etc/opt/emcpower/emcpbfs_linux remove` to uninstall PowerPath.

If the **ERROR: PowerPath is not configured on Boot from SAN.** message is displayed, then it indicates that emcpbfs\_linux config has not been used during PowerPath configuration. Follow the steps to manually upgrade or create the file below and re-run the script.

```
touch /etc/emc/.emcpbfs_linux_done
```

- Comment all other entries in the /etc/fstab file that refer to PowerPath pseudo (emcpower) devices.

---

**Note:** emcpbfs\_linux script modifies ‘/boot’ entry only in the /etc/fstab file

---

- Reboot the host.

---

**Note:** emcpbfs\_linux remove adds a new boot entry for the new initrd image in /boot/grub/menu.lst and sets the proper ‘default’ value. If older entries are not required, manually clean up and specify the proper ‘default’ value .

---

- Uninstall PowerPath.

## Uninstall using manual steps

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-3.0.13-0.27-default -k
/boot/vmlinuz-3.0.13-0.27-default
```
7. Add a new boot entry for the new `initrd` image created in the previous step in `/boot/grub/menu.lst` and make it the default entry.
8. Reboot the host.
9. Uninstall PowerPath.

## Uninstalling PowerPath in a SLES 12 SP2 Boot from SAN setup

Due to issues with `/boot` getting automatically unmounted while running `dracut` when PowerPath is active, the perform the following to uninstall PowerPath from a Boot From SAN host.

1. Unconfigure PowerPath.

```
/etc/opt/emcpower/emcpbfs_linux remove
```

For example,

```
~ # /etc/opt/emcpower/emcpbfs_linux remove

PHASE 1: VERIFICATION PHASE

Verifying PowerPath Installed
Verifying Boot LUN
emcpoweraq:libg210_BFS_0
Verifying Root LVM LUN(s) emcpoweraq:libg210_BFS_0
Verifying PPME Migrations No migration
Verifying /etc/fstab /boot found
Verifying /boot space Available
Verifying /etc/lvm/lvm.conf Filter found
Verifying /boot/grub2/grub.cfg Validated

PHASE 2: REMOVAL PHASE

Updating /etc/fstab Done
Updating /etc/lvm/lvm.conf Done
Creating /boot/initrd-3.12.49-11-default Done
Updating /boot/grub2/grub.cfg Done
Disabling PowerPath systemctl services ... Done

INFO: Successfully Completed.
 Restart the server and then un-install PowerPath.

```

If there are any issues reported, fix the issues, run **mount -a** and then rerun **/etc/opt/emcpower/emcpbfs\_linux remove**.

2. Run **mount -a**.

3. Remove the PowerPath dracut configuration file.

```
rm -rf /usr/lib/dracut/modules.d/90powerpath
```

4. Generate an initrd image without the PowerPath modules.

```
dracut --force /boot/initrd-$(uname -r) $(uname -r)
```

5. Reinitialize the 40\_custom file.

```
echo '#!/bin/sh' > /etc/grub.d/40_custom
echo 'exec tail -n +3 $0' >> /etc/grub.d/40_custom
```

6. Recreate the grub.cfg file.

```
grub2-mkconfig > /boot/grub2/grub.cfg
```

7. Set the default grub entry.

```
grep menuentry /boot/grub2/grub.cfg
```

For example,

```
grep menuentry /boot/grub2/grub.cfg
if [x"${feature_menuentry_id}" = xy]; then
 menuentry_id_option="--id"
 menuentry_id_option=""
export menuentry_id_option
 menuentry "If OK, run 'snapper rollback' and reboot." { true; }
menuentry 'SLES 12-SP2' --class sles --class gnu-linux --class gnu
--class os $menuentry_id_option
'gnulinux-simple-205d1894-e29f-4bd8-afc3-f575b195e222' {
submenu 'Advanced options for SLES 12-SP2' --hotkey=1
$menuentry_id_option
'gnulinux-advanced-205d1894-e29f-4bd8-afc3-f575b195e222' {
 menuentry 'SLES 12-SP2, with Linux 3.12.49-11-default'
--hotkey=2 --class sles --class gnu-linux --class gnu --class os
$menuentry_id_option
'gnulinux-3.12.49-11-default-advanced-205d1894-e29f-4bd8-afc3-f575b
195e222' {
libg210:~ #
libg210:~ # grub2-set-default 'SLES 12-SP2, with Linux
3.12.49-11-default' 'SLES 12-SP2, with Linux 3.12.49-11-default'
libg210:~ #
```

8. Reboot the host.

9. When the host comes up, uninstall PowerPath.

```
rpm -ev EMCpower.LINUX
```



# CHAPTER 4

## Managing paths to a logical device

A logical device of an EMC array can have multiple paths from the host. You can add new paths to the logical device or remove some of the existing paths from it. Changes affect the physical paths, logical device, HBAs, attached cables, storage array fabric zoning, and associated applications.

- ◆ [Dynamic LUN addition and removal](#)..... 52
- ◆ [Before you add or remove paths](#)..... 53
- ◆ [Adding new paths](#) ..... 53
- ◆ [Removing some of the paths](#) ..... 54

## Dynamic LUN addition and removal

PowerPath supports dynamically adding and removing LUNs when PowerPath is installed on RHEL 5.x, RHEL 6.x, RHEL 7.x, OL 6.x, OL 7.x, SLES 12 SP2, and SLES 11 SP4 systems.

**Note:** It is recommended to slow down 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

#### WARNING

EMC recommends that all I/O on the SCSI devices should be quiesced prior to attempting to rescan the SCSI bus.

1. Add the new LUNs into the configuration.
2. Scan for the LUNs by running the following script:

```
/etc/opt/emcpower/emcplun_linux <operation> <option> [-noprompt]
```

Where <operation> is **scan**

Where <option> can be **hba, iscsi, open-fcoe, all, aid=<array id>, <#>, hba=<#> tgt=<#>**

*For example:*

For scanning hba, type:

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

For scanning iSCSI software initiator, type:

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

If no operation is given, the emcplun\_linux script runs in interactive mode.

Use **/etc/opt/emcpower/emcplun\_linux help** for usage of the script.

Use **/etc/opt/emcpower/emcplun\_linux help scan** for information on the scan operation.

### Removing a LUN

The emcplun\_linux script can be found in the /etc/opt/emcpower/ directory.

1. Identify the device to be removed. Run **/etc/opt/emcplun\_linux list emcpower <id>**. Where <id> denotes the LUN to be removed.
2. Stop any I/O to the device being removed and ensure emcpower<id> is not in use.
3. Remove the LUN by running the following script:

```
/etc/opt/emcpower/emcplun_linux remove <pseudo-list> [-noprompt]
```

Where <pseudo-list> can be emcpowerX emcpowerY.

4. Remove LUNs from the storage side that correspond to the emcpower<id>.

## Resizing a LUN

To perform an online expansion of a LUN on RHEL 6, SLES 11 SP4, RHEL 7, SLES 12 SP2, OL 6.x, and OL 7.x:

1. Expand the LUN from the array.
2. For all native devices associated with the LUN, run the `#echo 1 > /sys/block/sd*/device/rescan` command.
3. Check the size of the pseudo and native devices using `parted` or `fdisk -l` on the devices.

## Before you add or remove paths

Before and after every change in the paths to a pseudo device, monitor and ensure that settings are as expected. Incorrect changes may result in problems such as data being written to the wrong logical device on the assumption that the native device is associated with that logical device.

- ◆ 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.
- ◆ This path reconfiguration is applicable only to logical devices connected to EMC arrays that are already under the managed system classes.
- ◆ After making configuration changes that affect host-to-storage-system connectivity or logical device identification, PowerPath must be reconfigured.

If PowerPath is not reconfigured after making configuration changes, many of the changes are treated as unintentional, and PowerPath tries to prevent them from affecting applications.

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

## Adding new paths

You can add new paths to a logical device that is already configured with at least one path.

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



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

**Procedure** To add 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.

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 and configure the additional paths.
  - ◆ To automatically scan and configure the additional paths, run the **/etc/opt/emcpower/emcplun\_linux scan <option> [-noprompt]** command, where <option> can be hba, iscsi, open-fcoe, all, aid=<array id>, <#>, hba=<#> tgt=<#>.
  - ◆ To manually scan and configure the additional paths:
    - a. Scan the HBA Bus on which new paths are configured.  
`echo "- - -" > /sys/class/scsi_host/host<#>/scan`
    - b. Run **powermt config** to configure the new paths under PowerPath control.
    - c. Inspect the new PowerPath configuration:
      - a. Run **powermt display dev=all**.
      - b. The new paths should be displayed with the alive state.
      - c. Run **powermt restore** to test all paths.
      - d. Scan operating system error logs to ensure no errors are logged against the new paths.
    - d. Correct any issues detected.
    - e. Run **powermt save** to save the new configuration.

## Removing some of the paths

You can remove some of the existing paths of a logical device 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.

**Procedure** To remove some paths to PowerPath logical devices:

**CAUTION**

**Failure to follow all the steps in this procedure can result in errors that affect the operating system, the PowerPath Management daemon, and other system applications.**

1. Stop all applications that use the native device associated with paths to be removed.

**CAUTION**

**If the path being removed is associated with a native device being used by any application, this procedure may fail. Select different paths to remove or stop using that path.**

2. Ensure that at least one path is active on the logical device.

**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).**

3. Run **`powermt display dev=all`** to confirm the configuration of the logical devices from which paths will be removed.
4. Check the number of existing paths.  
The path state should be alive for known good paths and dead for known bad paths.
5. If there is a problem, correct it before proceeding.
6. 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).
7. 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.

8. Run **`pp_inq -btl`** to obtain bus, TID and LUN. Save the output of the command.
9. If using character I/O (sg) devices on Linux, run **`sg_map -i -x`**, and save the output.
10. 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.
11. 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`**.

12. 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. The PowerPath Management Daemon will also issue errors when paths are physically disconnected. These error messages can be disregarded, and the errors do not affect applications as long as the steps in this procedure are followed correctly.

13. 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.
14. 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.

---

15. Remove the stale or unused native devices associated with the paths that were removed (the native devices names that are noted in Step 11).

**Note:** This step is not applicable for RHEL 6.x, OL 6.x, RHEL7, SLES 12 SP2, SLES11 SP4, and OL7 as stale native devices are removed automatically.

---

**IMPORTANT**

---

It is critical to remove the dead paths from PowerPath as mentioned in [step 14](#) before removing the stale native devices from the kernel. If not, the host can panic.

---

16. To remove the stale native devices from the kernel:

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

17. Inspect the new PowerPath configuration:

- a. Run **powermt display**

The output should show fewer total paths than before. All paths should reflect the 'optimal' state.

- b. Run **powermt display dev=all**

All remaining paths associated with the affected logical devices should reflect the 'alive' state.

18. Correct any issues detected above before saving the PowerPath configuration or using the new logical devices.

19. Run **powermt save** to save the new configuration.



# CHAPTER 5

## Coexistence with other path management software

If third-party path management software is present, plan about managing the classes, HBAs, and attached arrays to avoid undefined multipathing behavior.

- ◆ [PowerPath and third-party path management software.....](#) 60

## PowerPath and third-party path management software

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

- ◆ IBM Subsystem Device Driver (SDD)
- ◆ HP StorageWorks Secure Path
- ◆ HP StorageWorks AutoPath XP Virtual Array

---

**Note:** PowerPath is not compatible with the native Linux device mapper (DM-MPIO). EMC recommends that you blacklist all non-local devices on a host on which PowerPath will be installed.

---

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 class is disabled in PowerPath                      |
|-------------------------------------------|----------------------------------------------------------------------|
| IBM SDD                                   | ess                                                                  |
| HP StorageWorks Secure Path               | hpxp                                                                 |
| HP StorageWorks AutoPath XP Virtual Array | Class(es) that would ordinarily handle the MPxIO-configured array(s) |

### 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.

Because 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                                    |
|----------------------------------------|---------------------------------------------------------------------------------------------------|
| IBM SDD                                | powermt manage class=ess                                                                          |
| HP StorageWorks XP Auto Path           | powermt manage class=hpxp                                                                         |
| MPxIO                                  | powermt manage class=<class><br>Where <class> is any class that handles an MpxIO-configured array |

Because PowerPath and the third-party software cannot co-manage devices, do not initiate co-management by executing the **powermt manage class** 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 with PowerPath 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        |
|--------------------------------------------|-----------------------------|
| IBM SDD                                    | powermt unmanage class=ess  |
| HP StorageWorks XP AutoPath                | powermt unmanage class=hpxp |



# CHAPTER 6

## Maintaining PowerPath

You can configure various components in your PowerPath environment to maintain expected multipathing activities. These configuration activities could include enabling SNMP event messaging, providing logical volume manager support, host swapping of an HBA, R1/R2 boot failover support, or SCSI-2 reservation support.

- ◆ [Configuring for event monitoring .....](#) 64
- ◆ [Changing the language of installation messages .....](#) 66
- ◆ [Saving new configurations .....](#) 67
- ◆ [Configuring LVM2 support .....](#) 67
- ◆ [Hot swapping an HBA.....](#) 70
- ◆ [Configuring for R1/R2 boot failover support .....](#) 71
- ◆ [SCSI-2 reservation support in PowerPath for Linux .....](#) 72
- ◆ [Managing and unmanaging EMC Celerra iSCSI devices .....](#) 72
- ◆ [Enable scsi name on NetApp arrays.....](#) 73
- ◆ [Removing devices from PowerPath control .....](#) 74
- ◆ [Expand the file system on a pseudo-device .....](#) 74
- ◆ [Supporting more than 256 storage devices .....](#) 74
- ◆ [/etc/fstab options.....](#) 75

## Configuring for event monitoring

The PowerPath management daemon provides an event and alerting mechanism for multipathing events listed in the `emcp_mond.conf` file. This includes the ability for the host to be connected to a PowerPath Viewer console as well as send SNMP 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 cancelling an 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 64](#) lists 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

| Monitored event           | Event cause                                                  | Event ID          | SNMP Trap |
|---------------------------|--------------------------------------------------------------|-------------------|-----------|
| Path is Dead              | Path is dead for one minute.                                 | 0x2,0x4,60        | Yes       |
|                           | Path is dead and then recovers within one minute.            | 0x2,0x4,60        | No        |
| All Paths Dead            | All paths are dead for five minutes.                         | 0x20,0x40,300     | Yes       |
|                           | All paths are dead and then 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 66](#).

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, unsupported 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

The PowerPath management daemon provides an event and alerting mechanism for various multipathing events. The management daemon gets started along with the PowerPath service and the default settings are considered for the initial startup. Configure this daemon, if required.

The configuration can be enabled either by manually editing the `emcp_mond-sample.conf` file or by using the `emcp_mond_edit` (`sbin/emcp_mond_edit`).

**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 the `/etc/opt/emcpower/emcp_mond.conf` file or alternatively use the `emcp_mond_edit` utility to uncomment `PPSNMPTrapSvcEnable=TRUE` entry before starting the `emcp_mond.sh` service.

**Note:** Step 2 is necessary for Trap messages to be received for the dead paths.

3. 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
```

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

### **CAUTION**

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

- b. Verify that the `snmpd` service is enabled by using `chkconfig` command; if not, then enable the service.
  - c. Run `/etc/init.d/snmpd start` to start the snmpd daemon. If you are not planning to configure snmpd persistent across reboot, type `snmpd -c /etc/opt/emcpower/snmpd.conf` to start the SNMP daemon.
5. Type `ps -ef | grep snmpd` to check the status.

Output similar to the following is displayed:

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

6. To restart the management daemon:
  - On RHEL7, OL7, or SLES 12 SP2, run **systemctl restart EmcpMond.service**.
  - On all other kernels, run the following commands:

```
/etc/init.d/emcp_mond.sh stop
```

```
/etc/init.d/emcp_mond.sh start
```

## 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**

## Changing the language of installation messages

If the `LANG` environment variable is set to a locale ID listed in [Table 2 on page 67](#), 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.

---

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

**Table 2** 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 |          |

## Saving new configurations

When new devices are added to a host or old devices removed from the host, the configuration is modified.

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

The LVM2 (logical volume manager) can recognize PowerPath as the single path to LUN and use the PowerPath pseudo devices. For LVM2 to recognize only the PowerPath pseudo devices, modify the `/etc/lvm/lvm.conf` file in the host, and filter out the `sd` device nodes from the LVM2 internal cache.

---

**Note:** If you see `dm-multipath` devices after modifying the `lvm.conf` file, stop and remove `dm-multipath` before continuing. To disable native multipathing see [“Configuring PowerPath in a RHEL 6.x or OL 6.x Boot from SAN setup” on page 25](#), or [“Configuring PowerPath in a RHEL 7.x or OL 7.x Boot from SAN setup” on page 28](#).

---

Modification of this file in the host differs based on the configuration of the root file system of the host. The root file system can be configured on the host in any of the following three ways:

- ◆ root file system not mounted on a logical volume
- ◆ root file system mounted on a logical volume managed by PowerPath
- ◆ root file system mounted on a logical volume residing on a local disk not managed by PowerPath
- ◆ root file system mounted on mirrored logical volumes

**Note:** If duplicate PV messages are displayed while running LVM commands, ensure that the LVM filter is set correctly as detailed below:

## 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                                                                   |
|------------------|--------------------------------------------------------------------------|
| SLES 11 SP4      | <code>filter = [ "a/emcpower.*/", "r/sd.*/", "r/disk.*/" ]</code>        |
| SLES 12 SP2      | <code>global_filter = [ "a/emcpower.*/", "r/sd.*/", "r/disk.*/" ]</code> |
| RHEL 5.x         | <code>filter = [ "r/sd.*/", "r/disk.*/", "a./.*/" ]</code>               |
| RHEL 6.x         | <code>filter = [ "a/emcpower.*/", "r/sd.*/", "r/disk.*/" ]</code>        |
| RHEL 7.x         | <code>global_filter = [ "a/emcpower.*/", "r/sd.*/", "r/disk.*/" ]</code> |
| Oracle Linux 6.x | <code>filter = [ "a/emcpower.*/", "r/sd.*/", "r/disk.*/" ]</code>        |
| Oracle Linux 7.x | <code>global_filter = [ "a/emcpower.*/", "r/sd.*/", "r/disk.*/" ]</code> |

2. Run `vgscan -v` to rebuild the LVM2 cache.
3. Run `lvmdiskscan` and verify that the filtered device nodes are not listed in the command output.

## Root file system mounted on a logical volume

For details on the LVM filter that must be set, refer to [“Configuring PowerPath in a RHEL 6.x or OL 6.x Boot from SAN setup” on page 25](#), or [“Configuring PowerPath in a RHEL 7.x or OL 7.x Boot from SAN setup” on page 28](#).

## Root file system mounted on a logical volume residing on a local disk not managed by PowerPath

Identify the underlying devices for the root/swap logical volumes and the `/boot` devices, if any. Modify the LVM filter 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: `lvdisplay -m /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, on RHEL5:

To include partitions `sda1` to `sda9` for LVM2 while filtering the remaining SD device nodes, set the LVM filter field to: `["a/sda[1-9]$/", "a/emcpower.*/", "r/sd.*/", "r/disk.*/" ]`

For example, on SLES 12 SP2, SLES11, RHEL6, RHEL7 and OL7:

To include all `sda` partitions, run the following command and identify the SCSI ID of `sda` `/lib/udev/scsi_id --whitelisted --replace-whitespace --device=/dev/sda`

**SSEAGATE\_ST373207LC\_3KT4PMAN**

set the LVM filter field to:

`["a|/dev/.*/by-id/scsi-SSEAGATE_ST373207LC_3KT4PMAN.*|", "a/emcpower.*/", "r/sd.*/", "r/disk.*/" ]`

Modify the appropriate line as follows in the `/etc/lvm/lvm.conf` file for your Linux kernel:

| Kernel Filter    | Filter                                                                                                                           |
|------------------|----------------------------------------------------------------------------------------------------------------------------------|
| RHEL 5.x         | <code>filter = ["a/sda[1-9]\$/", "r/sd.*/", "r/disk.*/", "a.*/" ]</code>                                                         |
| RHEL 6.x         | <code>filter = [ "a /dev/.*/by-id/scsi-SSEAGATE_ST373207LC_3KT4PMAN.* ", "a/emcpower.*/", "r/sd.*/", "r/disk.*/" ]</code>        |
| SLES SP4         | <code>filter = [ "a /dev/.*/by-id/scsi-SSEAGATE_ST373207LC_3KT4PMAN.* ", "a/emcpower.*/", "r/sd.*/", "r/disk.*/" ]</code>        |
| Oracle Linux 6.x | <code>filter = [ "a /dev/.*/by-id/scsi-SSEAGATE_ST373207LC_3KT4PMAN.* ", "a/emcpower.*/", "r/sd.*/", "r/disk.*/" ]</code>        |
| RHEL 7.x         | <code>global_filter = [ "a /dev/.*/by-id/scsi-SSEAGATE_ST373207LC_3KT4PMAN.* ", "a/emcpower.*/", "r/sd.*/", "r/disk.*/" ]</code> |
| SLES 12 SP2      | <code>global_filter = [ "a /dev/.*/by-id/scsi-SSEAGATE_ST373207LC_3KT4PMAN.* ", "a/emcpower.*/", "r/sd.*/", "r/disk.*/" ]</code> |
| Oracle Linux 7.x | <code>global_filter = [ "a /dev/.*/by-id/scsi-SSEAGATE_ST373207LC_3KT4PMAN.* ", "a/emcpower.*/", "r/sd.*/", "r/disk.*/" ]</code> |

**Note:** In the previous 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. Run **lvmdiskscan** and verify that the SD and emcpower device nodes containing the root/swap/boot devices identified in steps 1 and 2 are listed in the command output.

## Hot swapping an HBA

PowerPath supports the Linux PCI hot plug feature 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 **pciehp** 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**.

This process does not include re-enabling the HBA through PowerPath by running:

```
powermt enable hba=#
```

The original HBA number is considered as removed from the system and when the new HBA is added, it is not assigned original host# but rather the next available host identifier.

## Configuring for R1/R2 boot failover support

A storage device can be associated with a bootable emcpower device. That storage device can also be mirrored through SRDF<sup>®</sup>. If it is mirrored, in case of a server failure, the boot device can fail over to the remote mirror disk, and can boot the server on an identical remote host.

---

**Note:** 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*.
- ◆ All R1 devices on Symmetrix A must have the same Host LUN Addresses as those of the R2 devices in the Symmetrix B. Ensure that Host Lun Addresses are consistent by using the following command on both Symmetrix A and Symmetrix B:
 

```
symaccess -sid <#> show view <view name>
```
- ◆ Each R1 device has only one mirror. (Concurrent SRDF is not supported.)
- ◆ Neither host has non-mirrored devices, BCVs, or gatekeepers.

---

**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.

Under SRDF R2 operation, device status may change from read-only to read-write, but the OS and PowerPath may not track that change.

Running `blockdev --rereadpt /dev/emcpowerX`, where `emcpowerX` is the R2 PowerPath pseudo device name, will update the device status correctly.

## Configure the network interface on the R1 host

You must configure the network interface on the R1 host appropriately or use DHCP so that network interface IP on R2 host should be configuration when R2 host boot up.

## SCSI-2 reservation support in PowerPath for Linux

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

## Managing and unmanaging EMC Celerra iSCSI devices

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

### Manage EMC Celerra iSCSI devices under PowerPath

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

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

2. Run `powermt init`.

3. Run **powermt config**.
4. To restart PowerPath or reboot the host:
  - a. Run the **/etc/init.d/PowerPath stop** or for RHEL7, OL7, or SLES 12 SP2 type **systemctl stop PowerPath**.
  - b. Run **/etc/init.d/PowerPath start** or for RHEL7, OL7, or SLES 12 SP2 type **systemctl start PowerPath.service**.

## Unmanage EMC Celerra iSCSI devices under PowerPath

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

```
managed:generic:EMC:Celerra
```
2. Run **powermt init**.
3. Run **powermt config**.
4. To restart PowerPath or reboot the host:
  - a. Run the **/etc/init.d/PowerPath stop**, or for RHEL7, OL7, or SLES 12 SP2 type **systemctl stop PowerPath**.
  - b. Run **/etc/init.d/PowerPath start**, or for RHEL7, OL7, or SLES 12 SP2 type **systemctl start PowerPath.service**.

## Enable scsi name on NetApp arrays

PowerPath reports NetApp ID in **powermt** output and in user interface only if the Report SCSI Name String Inquiry Descriptor setting is enabled for the initiator group through which host is connected to LUNs.

The following command needs to be run at both the controllers of the NetApp HA pair separately to enable the property. The command needs to be run on the NetApp array before the initiators are added to the initiator group else a reboot is required. There is no way to enable this property from the host.

Enable `report_scsi_name` by entering the following NetApp command:

```
igroup set igroup_name report_scsi_name yes
```

Use the following NetApp command to check the current value of the setting:

```
igroup show -v igroup_name
```

The setting is disabled in the following example.

```
system> igroup show -v
fcplnx (FCP):
OS Type: linux
Member: 21:00:00:24:ff:17:d7:11 (not logged in)
Member: 10:00:00:00:d9:e6:c1:b1 (logged in on: 0a)
UUID: ab7b40ac-917c-17e0-b240-123478563412
ALUA: Yes
Report SCSI Name String Inquiry Descriptor: NO
```

## Removing devices from PowerPath control

On PowerPath for Linux, devices removed using `powermt remove` will be automatically reconfigured in the following scenarios:

- ◆ For a pseudo device, if one or more native device paths are removed from PowerPath control using `powermt remove dev` those path(s) will be reconfigured in PowerPath as soon as IO is attempted to the native device.
- ◆ If a pseudo device is removed from PowerPath control, it will be added again or reconfigured in PowerPath when I/O is sent to any of the associated native devices, and only those native paths with IO sent to it will be reconfigured (unless `powermt config` is run, in which case all paths will be configured).

In order to remove a device from PowerPath control, use `powermt unmanage dev` to prevent it from being automatically reconfigured in PowerPath when IO is sent to one of its native device paths.

Some devices may also be automatically configured if PowerPath is running and new devices are provisioned to the server. In the case of VNX/CLARiiON devices, LUN nice name (s) and storage group (s) will be not be updated until users run `powermt config` or `powermt update lun_names`.

## Expand the file system on a pseudo-device

1. Expand the LUN on the array.
2. Unmount the file system if it is mounted.
3. Run `blockdev --rereadpt /dev/⟨pseudo_device⟩` to re-read the partition table of the pseudo device.
4. Run `echo 1>/sys/block/⟨native_device⟩/device/rescan` to scan the native path.
5. Follow the operating system's recommended steps to resize the file system:
  - `resize2fs -p /dev/⟨pseudo_device⟩` for the ext2, ext3, or ext4 file system
  - `xfs_growfs /dev/⟨pseudo_device⟩` for the XFS file system
  - `btrfs filesystem resize <size> /dev/⟨pseudo_device⟩` for the BTRFS file system
6. Run `mount /dev/⟨pseudo_device⟩ <mount_point>`.

## Supporting more than 256 storage devices

Complete the following steps to support more than 256 storage devices on a host.

### Emulex HBAs

If you are using Emulex HBAs, set the `lpfc_max_luns` parameter to **1024**, **2048**, or **4096**.

### QLogic HBAs

If you are using QLogic HBAs, change the system parameters under `/sys/module/scsi_mod/max_luns` and `max_report_luns` to the recommended value.

For more information, refer to the *EMC Host Connectivity Guide for Linux* or to the vendor website.

To make the parameter settings permanent across reboot on Emulex or QLogic HBAs:

### On RHEL 5.x

1. Edit the `/etc/modprobe.conf` file and set the `lpfc_max_luns` parameter to 1024, 2048, or 4096.
2. Recreate the `initrd`.
3. Point the `initrd` to the new `initrd` when booting.

### On RHEL 6.x

1. Create the `lpfc.conf` file in the `/etc/modprobe.d/` directory.
2. Edit the file and set the `lpfc_max_luns` parameter to 512.
3. Recreate the `initramfs` by running the `dracut -f` command.
4. Reboot the host.

For more information, refer to the *EMC Host Connectivity Guide for Linux* or the vendor website.

## /etc/fstab options

### RHEL5 and RHEL6

#### FC, iSCSI and Open-FCoE (RHEL6) devices

For all file systems residing on FC, iSCSI and Open-FCoE devices use `_netdev` option in the `/etc/fstab`

Example of a PowerPath device in `/etc/fstab`:

```
/dev/emcpowera1 /mnt/emcpowera1 ext3 _netdev 0 0
```

Example of a LVM created on a PowerPath device in `/etc/fstab`:

```
/dev/VG1/LV01 /mnt/VG1_LV01 ext3 _netdev 0 0
```

### SLES 12 SP2 and RHEL7

#### FC, iSCSI and Open-FCoE devices

For all file systems residing on FC, iSCSI and Open-FCoE devices use `nofail` option in the `/etc/fstab`

Example of a PowerPath device in `/etc/fstab`:

```
/dev/emcpowera1 /mnt/emcpowera1 ext3 nofail 0 2
```

Example of an LVM created on a PowerPath device in `/etc/fstab`:

```
/dev/VG1/LV01 /mnt/VG1_LV01 ext3 nofail 0 3
```

## SLES11 SP4

### FC devices

For all file systems residing on FC devices use defaults option in the /etc/fstab

Example of a PowerPath device in /etc/fstab:

```
/dev/emcpowera1 /mnt/emcpowera1 ext3 defaults 0 2
```

Example of an LVM created on a PowerPath device in /etc/fstab:

```
/dev/VG1/LV01 /mnt/VG1_LV01 ext3 defaults 0 2
```

### ISCSI and Open-FCoE Devices

For all file systems residing on ISCSI and Open-FCoE devices use nofail option in the /etc/fstab

Example of a PowerPath device in /etc/fstab:

```
/dev/emcpowera1 /mnt/emcpowera1 ext3 nofail 0 2
```

Example of an LVM created on a PowerPath device in /etc/fstab:

```
/dev/VG1/LV01 /mnt/VG1_LV01 ext3 nofail 0 2
```

# CHAPTER 7

## Upgrading PowerPath

Ensure that all the requirements are met, and then based on your requirement select a suitable upgrade procedure.

- ◆ [Preparing for upgrade .....](#) 78
- ◆ [Upgrading PowerPath.....](#) 78

## Preparing for upgrade

Before upgrading, ensure that the expected environment requirements are met as per the *E-Lab Interoperability Navigator* and prepare the host system for PowerPath upgrade.

- ◆ Download the latest version of EMC Grab utilities, available on the EMC Online Support site, and then run the PowerPath Configuration Checker (PPCC). You can upgrade to the latest version of PowerPath from PowerPath 5.0 or later. This may require upgrading the operating system before upgrading to the latest PowerPath version. 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 EMC Online Support site for the most current information.
- ◆ While upgrading PowerPath for Linux, the upgrade process saves the current configuration settings before beginning the upgrade.
- ◆ Backup `/etc/emcp_registration` file and `/etc/powermt.custom` or `/etc/powermt_custom.xml` files.
- ◆ Document your existing power device 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.
- ◆ Run the **`lsof`** command and ensure that none of the storage devices managed by the PowerPath are listed. If the **`lsof`** command lists the processes along with opened PowerPath managed storage devices, then close such applications/processes and ensure that PowerPath managed storage devices are not in use.
- ◆ Localized versions of the PowerPath installer are available in Brazilian Portuguese, French, German, Italian, Korean, Japanese, Spanish, and simplified Chinese. Refer [“Changing the language of installation messages” on page 66](#) for information on using a localized version of the PowerPath installer when upgrading PowerPath.
- ◆ Refer [“Troubleshooting PowerPath issues” on page 87](#) if you encounter any problems while upgrading PowerPath.

## Upgrading PowerPath

Based on your requirements, select a suitable upgrade method for your environment.

- ◆ [“Upgrading PowerPath from a CD” on page 79](#)
- ◆ [“Upgrading PowerPath from a compressed archive” on page 80](#)
- ◆ [“Upgrading PowerPath across Linux platforms” on page 80](#)
- ◆ [“Upgrading PowerPath before upgrading RHEL update on Linux” on page 82](#)

## Upgrading PowerPath from a CD

By using the installation CD, you can upgrade to the latest version of PowerPath.

1. Insert the PowerPath 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 SLES, type:

```
cd /cdrom/LINUX/pp<version>/sles
```

- For RHEL type:

```
cd /cdrom/LINUX/pp<version>/rhel
```

5. Log in as root.
6. To upgrade to the new PowerPath version, type:

```
rpm -Uhv EMCPower.LINUX-<version>-<build>.<platform>.x86_64.rpm
```

Where:

`<version>` is the product version; for example, PowerPath 6.2.

`<build>` is the software build number.

`<platform>` is the Linux distribution, for example OL6 or RHEL6.

7. Run **`/etc/init.d/PowerPath start`** (RHEL 6.x) or **`systemctl start PowerPath.service`** (RHEL7.x, OL7, or SLES 12 SP2) to start PowerPath.

Note: If issues are observed during or after starting PowerPath and if the old PowerPath version needs to be restored, the following steps should be done to preserve the saved configuration settings:

- a. Copy contents of `/etc/emc/archive` to another folder such as `/tmp/pp_config_backup`.
  - b. Uninstall the new version of PowerPath.
  - c. Delete contents of `/etc/emc/archive` folder since that now contains config files updated by the new PowerPath version.
  - d. Copy the backed up contents from step a to `/etc/emc/archive` folder.
  - e. Start PowerPath. While starting, PowerPath will automatically restore the config files from `/etc/emc/archive` and the old settings will be restored.
8. If the PowerPath configuration includes any custom policies or settings, run **`powermt load`** to load these settings.



- ◆ [“Upgrading PowerPath before upgrading to the latest SLES Service Pack ” on page 82](#)
- ◆ [“Upgrading PowerPath before upgrading RHEL update on Linux” on page 82](#)
- ◆ [“Upgrading PowerPath in a SteelEye LifeKeeper cluster environment” on page 82](#)

## Upgrading the Linux kernel in a PowerPath installed host

**Note:** Before upgrading the kernel on RHEL5, if the root filesystem is on a Logical Volume Manager (LVM), change the LVM filter to accept all SD devices.

To upgrade the Linux kernel in a PowerPath installed setup:

1. Log in as root.
2. Upgrade the kernel to a PowerPath supported version as described in your operating system documentation.
3. Restart the host.

## Upgrading PowerPath after upgrading from one major Linux release to the next

To upgrade PowerPath after upgrading from major Linux release to next major Linux release, for example, from RHEL6.5 to RHEL 7.0, perform the following steps:

1. Perform one of the following steps.
  - For RHEL 6.x, rename `/etc/init.d/PowerPath` to `/etc/init.d/PowerPath.old`
  - For RHEL7, OL7, or SLES 12 SP2 run the following commands:
 

```
systemctl disable PowerPath.service
systemctl disable PowerPathPost-start.service
systemctl disable EmcpMond.service
```
2. Upgrade the operating system to a PowerPath supported version as described in the RedHat and Novell documentation.

**Note:** The `upgradeany` and `ext4migrate` boot options enable upgrading from a previous major version of RHEL to RHEL 6.x. These options also enable migrating the existing ext3 file system to ext4 during the upgrade to RHEL 6.x, if you choose to migrate the file system to ext4. If you opt to use the `ext4migrate` boot option to migrate the ext4 file system, edit the `/etc/fstab` file manually after the upgrade. This is to ensure that the root file system is mounted as ext4 on the next reboot.

3. Upgrade PowerPath to the latest version as described in [“Upgrading PowerPath from a CD” on page 79](#).
4. Run `/etc/init.d/PowerPath start` (RHEL 6.x) or `systemctl start PowerPath.service` (RHEL7.x, OL7, or SLES 12 SP2) to start PowerPath.

Note: If issues are observed during or after starting PowerPath and if the old PowerPath version needs to be restored, the following steps should be done to preserve the saved configuration settings:

- a. Copy contents of `/etc/emc/archive` to another folder such as `/tmp/pp_config_backup`.

- b. Uninstall the new version of PowerPath.
  - c. Delete contents of /etc/emc/archive folder since that now contains config files updated by the new PowerPath version.
  - d. Copy the backed up contents from step a to /etc/emc/archive folder.
  - e. Start PowerPath. While starting, PowerPath will automatically restore the config files from /etc/emc/archive and the old settings will be restored.
5. After you have checked to make sure that the upgrade is complete, delete /etc/init.d/PowerPath.old.

## Upgrading PowerPath before upgrading to the latest SLES Service Pack

To upgrade PowerPath before upgrading Linux SLES on a host:

1. Upgrade PowerPath to the latest supported version for the current operating system version.

---

**Note:** Do not start PowerPath service at this time.

---

2. Upgrade the operating system to a PowerPath supported version.
3. Reboot the host.

## Upgrading PowerPath before upgrading RHEL update on Linux

To upgrade PowerPath before upgrading RHEL on a host:

1. Log in as root.
2. To upgrade to the new PowerPath version, type:

```
rpm -Uhv EMCPower.LINUX-<version>-<build>.<platform>.x86_64.rpm
```

Where:

    <version> is the product version; for example, PowerPath 6.2.

    <build> is the software build number.

    <platform> is the Linux distribution, for example OL6 or RHEL6.

3. Start PowerPath
4. Upgrade the operating system to a PowerPath supported version.
5. Reboot the host.

## Upgrading PowerPath in a SteelEye LifeKeeper cluster environment

In a SteelEye Life Keeper cluster environment, first backup the LifeKeeper configuration, stop the existing PowerPath for each node, and then upgrade the node to the PowerPath of later version.

---

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

---

1. Log in as root.

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

---

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

---

3. Run **lkstop -r** on the node you are upgrading to stop LifeKeeper.
4. Uninstall PowerPath.
  - a. Type **rpm -qa | grep -i emcpower**
  - b. Type **rpm -ev --nodeps <package\_name>**.

5. Copy the RPM package into a temporary folder.

6. Enter the following install command:

```
rpm -ivh EMCPower.LINUX-<version>-<build>.<platform>.x86_64.rpm
```

Where:

<version> is the product version; for example, PowerPath 6.2.

<build> is the software build number.

<platform> is the Linux distribution, for example OL6 or RHEL6.

7. Reboot the host.
8. Verify the PowerPath configuration:
  - a. Run **powermt display**.
  - b. Run **powermt display dev=all**.
9. Run **lkstart** to start SteelEye LifeKeeper.
10. Repeat steps 1-9 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. Manual switchover reduces any impact to running applications as opposed to automatic switchover.

---



# CHAPTER 8

## Removing PowerPath

Plan and prepare the host for removing or uninstalling PowerPath. You can either keep the existing PowerPath configuration and registration files for restoring in the next installation or completely remove them.

- ◆ [Before you remove PowerPath](#) ..... 86
- ◆ [Uninstalling PowerPath](#) ..... 86

## Before you remove PowerPath

- ❑ If the `powermt display every=<time_in_seconds>` command is running in a loop, stop the command.
- ❑ Ensure no PowerPath devices are in use. For example, unmount a mounted file system or deactivate LVM volume groups that use PowerPath pseudo devices (`vgchange -an <group>`).
- ❑ Run the `lsof` command and ensure that none of the storage devices managed by the PowerPath are listed. If the `lsof` command lists the processes along with opened PowerPath managed storage devices, then close such applications/processes and ensure that PowerPath managed storage devices are not in use.
- ❑ Manually remove references to PowerPath pseudo devices from system configuration files. For example, `/etc/fstab`.
- ❑ Undo any LVM filtering, as the pseudos will no longer be available after uninstallation.

## Uninstalling PowerPath

To uninstall a host booting from a local disk:

1. Uninstall PowerPath using the steps in [“To uninstall PowerPath” on page 86](#).

To uninstall a host booting from a SAN setup:

1. If the host is a boot from SAN set up, prepare the host to uninstall PowerPath using the procedures in:
  - [“Uninstalling PowerPath in a RHEL 5.x Boot from SAN setup” on page 44](#)
  - [“Uninstalling PowerPath in a RHEL 6.x or OL 6.x Boot from SAN setup” on page 44](#)
  - [“Uninstalling PowerPath in a RHEL 7.x or OL 7.x Boot from SAN setup” on page 45](#)
  - [“Uninstalling PowerPath in a SLES 11 SP4 Boot from SAN setup” on page 47](#)
  - [“Uninstalling PowerPath in a SLES 12 SP2 Boot from SAN setup” on page 48](#)
2. Uninstall PowerPath using the steps in [“To uninstall PowerPath” on page 86](#).

## To uninstall PowerPath

1. Log in as root.
2. Remove any LVM filtering, as the pseudos will no longer be available after uninstallation.
3. Display the package name. Type:

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

4. Remove the software. Type:

```
rpm -ev EMCPower.LINUX-<version>_<build>
```

Uninstalling PowerPath saves its configuration files for reuse. If required, you can also remove those configuration files after removing PowerPath from the host.

---

**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.

## Removing PowerPath files

The PowerPath uninstallation process adds the extension **.saved** and saves the PowerPath configuration and registration files. If required, you can remove these files after PowerPath uninstallation.

- ◆ `/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.xml.saved`
- ◆ `/etc/emc/powerkmd.custom`

If you are removing PowerPath from the host entirely, that is, if 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: 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):
```

2. Type **y** and press **Enter**.
3. Verify the cleanup and that no files remain, manually delete **/sbin/emcp\_cleanup**.



# CHAPTER 9

## Troubleshooting PowerPath issues

You can resolve possible problems that might occur while installing, upgrading, or when PowerPath is functioning on your host.

- ◆ [Recovering from a failed PowerPath installation](#) ..... 88
- ◆ [Missing or deleted files](#) ..... 89
- ◆ [The /etc/fstab](#) ..... 90
- ◆ [Problems booting the host](#) ..... 90
- ◆ [System hangs](#) ..... 91

## Recovering from a failed PowerPath installation

If PowerPath installation fails, correct error conditions reported by the installation program.

Install PowerPath following the instructions in [“Installing the RPM File” on page 16](#)

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

If the modules are loaded, type `/etc/init.d/PowerPath stop`, or for RHEL7, OL7, or SLES 12 SP2 type `systemctl stop PowerPath` to unload them.

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

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

3. Check the `/etc/opt/emcpower/log` file for additional information about the failure.
4. 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-<version>-<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

Learn about possible problems that might occur while upgrading PowerPath, and also understand how to resolve them.

### 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.xml` file is outdated and does not match the current PowerPath configuration.
- ◆ The file system where the `powermt_custom.xml` 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.

### 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 LUN-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=<policy>** to reset the desired policy 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 settings for the storage devices in your PowerPath configuration.

## 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. To stop PowerPath, type **/etc/init.d/PowerPath stop**, or for RHEL7, OL7, or SLES 12 SP2 type **systemctl start PowerPath.service**.

Use the script only to load and unload PowerPath modules.

---

**Note:** If the 'hald' daemon is running in the system, this command takes longer to complete.

---

- To remove PowerPath, type:

```
rpm -qa | grep EMCpower.LINUX
rpm -ev EMCpower.LINUX-<version>-<build>
```

---

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

---

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

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



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

---

- Reinstall PowerPath. Refer to “[Installing the RPM File](#)” on page 16 for more information.

## The `/etc/fstab`

For information, refer to the *EMC Host Connectivity Guide for Linux*.

## 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*`.

## Boot issue on RHEL6.x, SLES11 SP4, RHEL7.x, SLES 12 SP2, OL6.x, and OL7.x boot-from-SAN when emcpbfs\_linux is used

The emcpbfs\_linux takes the backup of system configuration files and keeps the backup files in the same location and in the timestamp format listed below.

**‘/etc/opt/emcpower/emcpbfs\_linux config’**

```
/etc/fstab_bkp_emcpbfs_linux_config_Jul112013_14h34m09s
/etc/lvm/lvm.conf_bkp_emcpbfs_linux_config_Jul112013_14h34m09s
```

For RHEL 6.x and OL 6.x:

```
/boot/grup/menu.lst_bkp_emcpbfs_linux_config_Jul112013_14h20m36s
```

For RHEL 7.x, SLES 12 SP2, and OL 7.x:

```
/boot/grub2/grub.cfg_bkp_emcpbfs_linux_config_Nov262014_15h29m07s
```

**‘/etc/opt/emcpower/emcpbfs\_linux remove’**

```
/etc/fstab_bkp_emcpbfs_linux_remove_Jul112013_14h20m36s
```

```
/etc/lvm/lvm.conf_bkp_emcpbfs_linux_remove_Jul112013_14h20m36s
```

For RHEL 6.x and OL 6.x:

```
/etc/init.d/PowerPath.bkp_emcpbfs_linux_remove_Jul112013_14h20m36s
```

```
/boot/grup/menu.lst_bkp_emcpbfs_linux_remove_Jul112013_14h20m36s
```

For RHEL 7.x, SLES 12 SP2, and OL 7.x:

```
/boot/grub2/grub.cfg_bkp_emcpbfs_linux_remove_Nov262014_14h39m50s
```

**‘/etc/opt/emcpower/emcpbfs\_linux upgrade’**

```
/etc/fstab_bkp_emcpbfs_linux_upgrade_Jul152013_14h43m46s
```

```
/etc/lvm/lvm.conf_bkp_emcpbfs_linux_upgrade_Jul152013_14h43m46s
```

**Note:** The **emcpbfs\_linux** script also comments out original entry in the **/etc/fstab** and **/etc/lvm/lvm.conf**, if required these can also be used for recovering the configuration.

*For example*

```
/etc/fstab
```

```
###emcpbfs_linux###Thu Jul 19 19:54:40 IST 2012###
```

```
#/dev/disk/by-id/scsi-360000970000192601710533033464437-part1 /boot
ext3 acl,user_xattr 1 2
```

```
#####
```

For RHEL 6.x and OL 6.x:

```
/etc/lvm/lvm.conf
```

```
###emcpbfs_linux###Thu Jul 19 19:54:40 IST 2012###
```

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

```
#####
```

For RHEL 7.x, SLES 12 SP2, and OL 7.x use global filter and comment out the global\_filter:

```
/etc/lvm/lvm.conf
```

```
global_filter = ["a/emcpower.*/", "r/sd.*/", "r/disk.*/"]
```

Boot from OS CD or DVD in rescue mode, mount **/boot** and **/**, then manually move the backup files to original location and reboot the host. If required, select non-PowerPath **initramfs** during host boot up.

## System hangs

System hangs may be caused by a number of reasons. This section provides workarounds to some scenarios that you may encounter.

## Unsupported HBA drivers

If your system hangs after you install PowerPath or reboot after installation the first time, 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 EMC Online Support site.

---

The driver version distributed with your Linux version 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. Rebuild the RAM disk being used.
4. Reboot 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. Reboot the host with the storage attached.

## Other causes

If your system hangs even though the correct HBA drivers have been used, check the Issue Tracker application on EMC Online Support site for other possible causes.

# APPENDIX A

## Files changed by PowerPath

The configuration and registration files are affected because of installing a newer version of PowerPath.

- ◆ [Linux files modified by PowerPath installation.....](#) 96
- ◆ [Files created by PowerPath installation.....](#) 96

## Linux files modified by PowerPath installation

Configuration and system files are modified when PowerPath is installed on a Linux host.

### RHEL 5

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

### RHEL 6

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

### RHEL 7

- ◆ /etc/modprobe.d/powerpath.conf
- ◆ /lib/modules/\*/modules.dep — For each /lib/modules subdirectory that gets a set of PowerPath drivers installed.

## Files created by PowerPath installation

When PowerPath is installed on a Linux host, the installation creates a set of files.

### emcp\_devicesDB.dat

- ◆ emcp\_devicesDB.idx
- ◆ emcp\_registration
- ◆ modprobe.conf.pp
- ◆ powermt\_custom.xml
- ◆ powermt\_custom.xml.lock

### /etc/emc/

- ◆ powerkmd.custom
- ◆ mpaa.excluded
- ◆ mpaa.lams
- ◆ powerpath\_pmi.txt
- ◆ powerpath\_pmi.txt.lock

- ◆ ppme
- ◆ archive

/etc/init.d/

### **RHEL5, RHEL6, RHEL 7, OVM 3.x.x, OL6, and OL7**

- ◆ PowerPath
- ◆ PowerPathPost

### **SLES 11 SP4 and SLES 12 SP2**

- ◆ PowerPath
- ◆ emcp\_mond.sh
- ◆ boot.powerpath
- ◆ emcp\_mond.sles
- ◆ PowerPathPost
- ◆ PowerPathPost.sles

/etc/init.d/boot.d/

### **SLES only**

- ◆ K18boot.powerpath
- ◆ S04boot.powerpath

/etc/rc.d/rc3.d

### **SLES 11 SP4**

- ◆ K01powerfcoe.sles
- ◆ K01poweriscsi.sles
- ◆ K01powermig.sles
- ◆ S03powerfcoe.sles
- ◆ S04powermig.sles

/etc/rc.d/rc5.d

### **SLES 11 SP4**

- ◆ K01powerfcoe.sles
- ◆ K01poweriscsi.sles
- ◆ K01powermig.sles
- ◆ S03powerfcoe.sles
- ◆ S03poweriscsi.sles

- ◆ S04powermig.sles

/etc/rc.d/rc2.d

#### **SLES 11 SP4**

- ◆ K01powermig.sles
- ◆ S04powermig.sles

/etc/rc.d/rcS.d

#### **SLES 11 P4**

- ◆ S04powermig.sles

/etc/rc.d/init.d

#### **RHEL5, RHEL6, OVM 3.x.x, OL6**

- ◆ PowerPath
- ◆ PowerPathPost
- ◆ emcp\_mond.sh

/etc

#### **RHEL5, RHEL6, OVM 3.1.1, OL6**

- ◆ rc.d/rc3.d/S24Powerlscsi
- ◆ rc.d/rc3.d/S24PowerFcoe
- ◆ rc.d/rc3.d/S99PowerMig
- ◆ rc.d/rc3.d/K99PowerMig
- ◆ rc.d/rc3.d/S51emcp\_mond
- ◆ rc.d/rc5.d/S24Powerlscsi
- ◆ rc.d/rc5.d/S24PowerFcoe
- ◆ rc.d/rc5.d/S99PowerMig
- ◆ rc.d/rc5.d/K99PowerMig
- ◆ rc.d/rc5.d/S51emcp\_mond

/etc/opt/emcpower

- ◆ \_\_emcp\_db\_global\_lock
- ◆ \_\_emcp\_db\_lock
- ◆ .os\_version
- ◆ .pp\_version
- ◆ EULA.pdf

- ◆ bin
- ◆ emcp\_mond-sample.conf
- ◆ emcpbfs\_linux
- ◆ emcpmgr
- ◆ log
- ◆ patch.log
- ◆ powercf
- ◆ pp\_reset\_dm\_deps.sh
- ◆ server.pem
- ◆ snmpd.conf
- ◆ emcplun\_linux
- ◆ 80-udev-pp-xio.rules

/etc/opt/emcpower/.tmp

**SLES 12 SP2, RHEL 7, RHEL 6, RHEL 5, OL6, OL 7, OVM 3.x.x**

- ◆ rollbackfiles.tar
- ◆ rollbackfiles <XXXXX>.tar
- ◆ undolist
- ◆ undolist <XXXXX>

**SLES 11 SP4**

- ◆ rollbackfiles.tar
- ◆ rollbackfiles <XXXXX>.tar
- ◆ setup-mkdumprd.sh
- ◆ undolist
- ◆ undolist <XXXXX>

/etc/opt/emcpower/bin

- ◆ envsubst
- ◆ gettext
- ◆ gettext.sh

/etc/opt/emcpower/EMCpower.LINUX-6.2.x.00.00

- ◆ .pp\_build\_version
- ◆ .pp\_vendor
- ◆ .pp\_version

- ◆ EULA.pdf
- ◆ PowerPathPost
- ◆ PowerPath
- ◆ PowerPathPost-start.service (SLES 12 SP2, RHEL7 and OL7)
- ◆ PowerPathPost-stop.service (SLES 12, RHEL7 and OL7)
- ◆ PowerPath.service (SLES 12 SP2, RHEL7 and OL7)
- ◆ boot-powerpath.sh
- ◆ boot.powerpath
- ◆ emcp\_mond.sh
- ◆ emcp\_mond.sles (SLES only)
- ◆ enable
- ◆ modprobe.conf.pp
- ◆ powerfcoc.sles (SLES only)
- ◆ powerpathpost.sles (SLES only)
- ◆ pp\_reset\_dm\_deps.sh
- ◆ pp\_udev.sh
- ◆ pp\_udev\_xio.sh
- ◆ setup-powerpath.sh
- ◆ udev-pp.rules
- ◆ udev-pp-xio.rules

### **RHEL6 and OL6**

- ◆ 04-poweriscsi

/etc/opt/emcpower/EMCpower.LINUX-6.2.x.00.00/bin

- ◆ .cmds
- ◆ .drivers\_base
- ◆ .drivers\_ext
- ◆ .drivers\_mgr
- ◆ .sharedlibs
- ◆ .sharedlibs32
- ◆ .staticlibs

/etc/opt/emcpower/EMCpower.LINUX-6.2.x.00.00/bin/cmds

- ◆ emcp\_cleanup
- ◆ emcp\_mond

- ◆ emcp\_mond\_edit
- ◆ emcp\_mond-sample.conf
- ◆ emcpadm
- ◆ emcpmgr
- ◆ emcpmigd
- ◆ emcpreg
- ◆ powercf
- ◆ powermig
- ◆ powermt
- ◆ pp\_inq
- ◆ server.pem
- ◆ snmpd.conf
- ◆ powermt\_display\_dead

/etc/opt/emcpower/EMCpower.LINUX-6.2.x.00.00/bin/driver

- ◆ sles12default\_x8664
- ◆ sles12default\_x8664/emcp.ko
- ◆ sles12default\_x8664/emcpdm.ko
- ◆ sles12default\_x8664/emcpgpx.ko
- ◆ sles12default\_x8664/emcpioc.ko
- ◆ sles12default\_x8664/emcpmpx.ko
- ◆ sles12xen\_x8664
- ◆ sles12xen\_x8664/emcp.ko
- ◆ sles12xen\_x8664/emcpdm.ko
- ◆ sles12xen\_x8664/emcpgpx.ko
- ◆ sles12xen\_x8664/emcpioc.ko
- ◆ sles12xen\_x8664/emcpmpx.ko

#### **SLES 11 SP4**

- ◆ sles11sp3default\_x8664
- ◆ sles11sp3default\_x8664/emcp.ko
- ◆ sles11sp3default\_x8664/emcpdm.ko
- ◆ sles11sp3default\_x8664/emcpgpx.ko
- ◆ sles11sp3default\_x8664/emcpioc.ko
- ◆ sles11sp3default\_x8664/emcpmpx.ko
- ◆ sles11sp3xen\_x8664

- ◆ sles11sp3xen\_x8664/emcp.ko
- ◆ sles11sp3xen\_x8664/emcpdm.ko
- ◆ sles11sp3xen\_x8664/emcpgpx.ko
- ◆ sles11sp3xen\_x8664/emcpioc.ko
- ◆ sles11sp3xen\_x8664/emcpx.ko

### **RHEL7**

- ◆ rhel7\_x8664/emcp.ko
- ◆ rhel7\_x8664/emcpdm.ko
- ◆ rhel7\_x8664/emcpgpx.ko
- ◆ rhel7\_x8664/emcpioc.ko
- ◆ rhel7\_x8664/emcpx.ko

### **RHEL6**

- ◆ rhel6\_x8664/emcp.ko
- ◆ rhel6\_x8664/emcpdm.ko
- ◆ rhel6\_x8664/emcpgpx.ko
- ◆ rhel6\_x8664/emcpioc.ko
- ◆ rhel6\_x8664/emcpx.ko

### **RHEL5**

- ◆ rhel5\_x8664/emcp.ko
- ◆ rhel5\_x8664/emcpdm.ko
- ◆ rhel5\_x8664/emcpgpx.ko
- ◆ rhel5\_x8664/emcpioc.ko
- ◆ rhel5\_x8664/emcpx.ko
- ◆ rhel5xen\_x8664/emcp.ko
- ◆ rhel5xen\_x8664/emcpdm.ko
- ◆ rhel5xen\_x8664/emcpgpx.ko
- ◆ rhel5xen\_x8664/emcpioc.ko
- ◆ rhel5xen\_x8664/emcpx.ko

### **OL7 UEK3 R2**

- ◆ ol7\_uek3\_r2\_x8664/emcp.ko
- ◆ ol7\_uek3\_r2\_x8664/emcpdm.ko
- ◆ ol7\_uek3\_r2\_x8664/emcpgpx.ko
- ◆ ol7\_uek3\_r2\_x8664/emcpioc.ko

- ◆ ol7\_uek3\_r2\_x8664/emcpx.ko

### **OL6 UEK2 R2**

- ◆ ol6\_uek2\_r2\_x8664/emcp.ko
- ◆ ol6\_uek2\_r2\_x8664/emcpdm.ko
- ◆ ol6\_uek2\_r2\_x8664/emcpgpx.ko
- ◆ ol6\_uek2\_r2\_x8664/emcpioc.ko
- ◆ ol6\_uek2\_r2\_x8664/emcpx.ko

### **OL6 UEK2 R3**

- ◆ ol6\_uek2\_r3\_x8664/emcp.ko
- ◆ ol6\_uek2\_r3\_x8664/emcpdm.ko
- ◆ ol6\_uek2\_r3\_x8664/emcpgpx.ko
- ◆ ol6\_uek2\_r3\_x8664/emcpioc.ko
- ◆ ol6\_uek2\_r3\_x8664/emcpx.ko

### **OL6 UEK2 R4**

- ◆ ol6\_uek2\_r4\_x8664/emcp.ko
- ◆ ol6\_uek2\_r4\_x8664/emcpdm.ko
- ◆ ol6\_uek2\_r4\_x8664/emcpgpx.ko
- ◆ ol6\_uek2\_r4\_x8664/emcpioc.ko
- ◆ ol6\_uek2\_r4\_x8664/emcpx.ko

### **OL6 UEK2 R5**

- ◆ ol6\_uek2\_r5\_x8664/emcp.ko
- ◆ ol6\_uek2\_r5\_x8664/emcpdm.ko
- ◆ ol6\_uek2\_r5\_x8664/emcpgpx.ko
- ◆ ol6\_uek2\_r5\_x8664/emcpioc.ko
- ◆ ol6\_uek2\_r5\_x8664/emcpx.ko

### **OL6 UEK3**

- ◆ ol6\_uek3\_x8664/emcp.ko
- ◆ ol6\_uek3\_x8664/emcpdm.ko
- ◆ ol6\_uek3\_x8664/emcpgpx.ko
- ◆ ol6\_uek3\_x8664/emcpioc.ko
- ◆ ol6\_uek3\_x8664/emcpx.ko

### **OL6 UEK3 R3**

- ◆ ol6\_uek3\_r3\_x8664/emcp.ko

- ◆ ol6\_uek3\_r3\_x8664/emcpdm.ko
- ◆ ol6\_uek3\_r3\_x8664/emcpgpx.ko
- ◆ ol6\_uek3\_r3\_x8664/emcpioc.ko
- ◆ ol6\_uek3\_r3\_x8664/emcpx.ko

#### **OL6 UEK3 R4**

- ◆ ol6\_uek3\_r4\_x8664/emcp.ko
- ◆ ol6\_uek3\_r4\_x8664/emcpdm.ko
- ◆ ol6\_uek3\_r4\_x8664/emcpgpx.ko
- ◆ ol6\_uek3\_r4\_x8664/emcpioc.ko
- ◆ ol6\_uek3\_r4\_x8664/emcpx.ko

#### **OL6 UEK3 R5**

- ◆ ol6\_uek3\_r5\_x8664/emcp.ko
- ◆ ol6\_uek3\_r5\_x8664/emcpdm.ko
- ◆ ol6\_uek3\_r5\_x8664/emcpgpx.ko
- ◆ ol6\_uek3\_r5\_x8664/emcpioc.ko
- ◆ ol6\_uek3\_r5\_x8664/emcpx.ko

#### **OL6 UEK3 R1**

- ◆ ol6\_uek3\_r1\_x8664/emcp.ko
- ◆ ol6\_uek3\_r1\_x8664/emcpdm.ko
- ◆ ol6\_uek3\_r1\_x8664/emcpgpx.ko
- ◆ ol6\_uek3\_r1\_x8664/emcpioc.ko
- ◆ ol6\_uek3\_r1\_x8664/emcpx.ko

#### **OL6 UEK R1**

- ◆ ol6\_uek1\_r1\_x8664/emcp.ko
- ◆ ol6\_uek1\_r1\_x8664/emcpdm.ko
- ◆ ol6\_uek1\_r1\_x8664/emcpgpx.ko
- ◆ ol6\_uek1\_r1\_x8664/emcpioc.ko
- ◆ ol6\_uek1\_r1\_x8664/emcpx.ko

#### **OL6 UEK R4**

- ◆ ol6\_uek1\_r4\_x8664/emcp.ko
- ◆ ol6\_uek1\_r4\_x8664/emcpdm.ko
- ◆ ol6\_uek1\_r4\_x8664/emcpgpx.ko
- ◆ ol6\_uek1\_r4\_x8664/emcpioc.ko

- ◆ ol6\_uek1\_r4\_x8664/emcpx.ko

/etc/opt/emcpower/EMCpower.LINUX-6.2.x.00.00/bin/lib

- ◆ emcp\_mond\_netsnmp.so
- ◆ libemcp.so
- ◆ libemcp\_core.so
- ◆ libemcp\_lam.so
- ◆ libemcp\_lic\_rtl.so
- ◆ libemcp\_mp\_rtl.so
- ◆ libemcp\_mpapi\_rtl.so
- ◆ libemcpx.so
- ◆ libpn.so

/etc/opt/emcpower/EMCpower.LINUX-6.2.x.00.00/bin/lib32

- ◆ libemcp.so
- ◆ libemcp\_core.so
- ◆ libemcp\_lam.so
- ◆ libemcp\_lic\_rtl.so
- ◆ libemcp\_mp\_rtl.so
- ◆ libemcp\_mpapi\_rtl.so
- ◆ libemcpx.so
- ◆ libpn.so

/etc/opt/emcpower/EMCpower.LINUX-6.2.x.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

#### /etc/opt/emcpower/EMCpower.6.2.x.00.00/man

- ◆ .man
- ◆ emcpadm.1
- ◆ emcpreg.1
- ◆ emcpupgrade.1
- ◆ powermig.1
- ◆ powermt.1

#### /etc/udev/rules.d

- ◆ 05-udev-pp.rules
- ◆ 80-udev-pp-xio.rules

#### /lib

- ◆ libemcp.so
- ◆ libemcp\_core.so
- ◆ libemcp\_lam.so
- ◆ libemcp\_lic\_rtl.so
- ◆ libemcp\_mp\_rtl.so
- ◆ libemcp\_mpapi\_rtl.so
- ◆ libemcump.so
- ◆ libpn.so

#### /lib/udev

- ◆ pp\_udev.sh
- ◆ pp\_udev\_xio.sh

#### /lib/modules/<kernel\_version>/powerpath

- ◆ emcpdm.ko
- ◆ emcpgpx.ko
- ◆ emcpioc.ko
- ◆ emcp.ko
- ◆ emcpx.ko

Where <kernel\_version> is the version of the kernel currently running on a host.

## /lib64

- ◆ emcp\_mond\_netsnmp.so
- ◆ libemcp.so
- ◆ libemcp\_core.so
- ◆ libemcp\_lam.so
- ◆ libemcp\_lic\_rtl.so
- ◆ libemcp\_mp\_rtl.so
- ◆ libemcp\_mpapi\_rtl.so
- ◆ libemcpmp.so
- ◆ libpn.so

## /sbin

- ◆ emcp\_cleanup
- ◆ emcp\_mond
- ◆ emcp\_mond\_edit
- ◆ emcpadm
- ◆ emcpmigd
- ◆ emcpreg
- ◆ powermig
- ◆ powermt
- ◆ powermt\_display\_dead

## /usr/lib/systemd/system

### **RHEL7, SLES 12 SP2 and OL7**

- ◆ PowerPath.service
- ◆ PowerPathPost-start.service
- ◆ PowerPathPost-stop.service

## /usr/lib/PowerPath

### **RHEL7, SLES 12 SP2 and OL7**

- ◆ PowerPath
- ◆ PowerPathPost
- ◆ emcp\_mond.sh

## /usr/share/man/man1

- ◆ powermig.1
- ◆ powermt.1
- ◆ emcpadm.1
- ◆ emcprep.1
- ◆ emcpupgrade.1

### **SLES 12 SP2, RHEL7.x and OL7.x**

- ◆ /usr/lib/dracut/modules.d/90powerpath
- ◆ 64-emcp.rules
- ◆ emcp\_pre-trigger.sh
- ◆ emcp\_settled.sh
- ◆ module-setup.sh
- ◆ /etc/opt/emcpower/EMCpower.LINUX-6.2.0.00.00/90powerpath
- ◆ 64-emcp.rules
- ◆ check
- ◆ emcp\_pre-trigger.sh
- ◆ emcp\_settled.sh
- ◆ install
- ◆ module-setup.sh

## /lib/mkinitrd/scripts

- ◆ boot-powerpath.sh
- ◆ setup-powerpath.sh

## /lib/mkinitrd/boot

- ◆ 21-powerpath.sh

## /lib/mkinitrd/setup

- ◆ 71-powerpath.sh

### **RHEL6.x and OL6.x**

- ◆ /usr/share/dracut/modules.d/90powerpath
- ◆ /etc/opt/emcpower/EMCpower.LINUX-6.2.x.00.00/90powerpath
- ◆ check
- ◆ emcp\_pre-trigger.sh
- ◆ emcp\_settled.sh
- ◆ install

- ◆ 64-emcp.rules
- ◆ module-setup.sh

Files changed by PowerPath