

EMC PowerPath for Linux

Version 6.3

Installation and Administration Guide

302-004-486

02

Copyright © 1998-2018 Dell Inc. or its subsidiaries. All rights reserved.

Published February 2018

Dell 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.” DELL 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 DELL SOFTWARE DESCRIBED IN THIS PUBLICATION REQUIRES AN APPLICABLE SOFTWARE LICENSE.

Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners.
Published in the USA.

EMC Corporation
Hopkinton, Massachusetts 01748-9103
1-508-435-1000 In North America 1-866-464-7381
www.EMC.com

CONTENTS

Tables		7
Chapter 1	Getting started	9
	Getting started.....	10
Chapter 2	Installing PowerPath	11
	Before you install.....	12
	Storage System related steps.....	12
	Reviewing compatibility with the environment.....	12
	Review host system compatibility.....	13
	Review PowerPath licensing.....	13
	Review third-party path management software.....	13
	Review the installation language.....	14
	Install the RPM file.....	14
	Install from a CD.....	15
	Install in a Xen environment.....	15
	Install PowerPath when VSF is already installed.....	17
	License after the installation.....	17
	Post-installation steps.....	18
Chapter 3	Administering in a Boot from SAN environment	19
	Before configuring for a root device.....	20
	Configuring PowerPath in a RHEL 6.x or OL 6.x Boot from SAN setup in a BIOS host.....	20
	Disabling dm-multipath on a RHEL 6.x or OL 6.x host.....	23
	Configuring PowerPath in a RHEL 7.x, OL 7.x or SLES 12 SP3 Boot from SAN setup in a BIOS host.....	23
	Disabling dm-multipath on a RHEL 7.x, OL 7.x or SLES 12 SP3 host..	26
	Configuring PowerPath in a SLES 11 SP4 Boot from SAN setup in a BIOS host.....	26
	Disabling dm-multipath on a SLES 11 SP4 host.....	30
	Upgrading PowerPath in a RHEL 6.x or OL 6.x Boot from SAN setup in a BIOS host.....	31
	Upgrade using the <code>emcpbfs_linux</code> script	31
	Upgrade using manual steps	31
	Upgrading PowerPath in a RHEL 7.x, OL 7.x or SLES 12 SP3 Boot from SAN setup in a BIOS host.....	33
	Upgrade using the <code>emcpbfs_linux</code> script	33
	Upgrade using manual steps	33
	Upgrading PowerPath in a SLES 11 SP4 Boot from SAN setup in a BIOS host	35
	Upgrade using the <code>emcpbfs_linux</code> script	35
	Upgrade using manual steps.....	36
	Upgrading the Linux kernel in a Boot from SAN setup in a BIOS host.....	37
	Upgrade using <code>emcpbfs_linux</code> script	37
	Upgrade using manual steps	38

	Uninstalling PowerPath in a RHEL 6.x or OL 6.x Boot from SAN setup in a BIOS host	40
	Uninstalling PowerPath in a RHEL 7.x, OL 7.x or SLES 12 SP3 Boot from SAN setup in a BIOS host.....	42
	Uninstall using <code>emcpbfs_linux_script</code>	42
	Uninstall using manual steps.....	43
	Uninstalling PowerPath in a SLES 11 SP4 Boot from SAN setup in a BIOS host.....	44
	Uninstall using <code>emcpbfs_linux_script</code>	44
	Uninstall using manual steps.....	45
	Configure PowerPath in a RHEL 6.x or OL 6.x Boot from SAN setup in an UEFI host.....	45
	Configure PowerPath in a RHEL 7.x, OL 7.x or SLES 12 SP3 Boot from SAN setup in an UEFI host.....	48
	Upgrading PowerPath in a RHEL 6.x or OL 6.x Boot from SAN setup in an UEFI host.....	51
	Upgrade using the <code>emcpbfs_linux</code> script	51
	Upgrade using manual steps	52
	Upgrading PowerPath in a RHEL 7.x, OL 7.x or SLES 12 SP3 Boot from SAN setup in an UEFI host.....	53
	Upgrade using the <code>emcpbfs_linux</code> script	53
	Upgrade using manual steps	54
	Uninstalling PowerPath in a RHEL 6.x or OL 6.x Boot from SAN setup in an UEFI host	55
	Uninstall using <code>emcpbfs_linux_script</code>	55
	Uninstall using manual steps.....	56
	Uninstalling PowerPath in a RHEL 7.x, OL 7.x or SLES 12 SP3 Boot from SAN setup in an UEFI host.....	57
	Uninstall using <code>emcpbfs_linux_script</code>	57
	Uninstall using manual steps.....	58
Chapter 4	Managing paths to a logical device	61
	Dynamic LUN addition and removal.....	62
	Adding a LUN.....	62
	Removing a LUN.....	62
	Resizing a LUN.....	63
	Before you add or remove paths.....	63
	Adding new paths to a PowerPath logical device.....	64
	Removing some of the paths to PowerPath logical devices.....	65
Chapter 5	Coexistence with other path management software	69
	PowerPath and third-party path management software.....	70
	PowerPath co-management behavior.....	70
Chapter 6	Maintaining PowerPath	73
	Configure PowerPath for event monitoring.....	74
	Monitored events.....	74
	Configuring the Management Daemon.....	75
	Editing the event.....	76
	Changing the language of installation messages.....	76
	Saving new configurations.....	77
	Configuring LVM2 support.....	77
	Root file system not mounted on a logical volume.....	78

	Root file system mounted on a logical volume.....	79
	Root file system mounted on a logical volume residing on a local disk not managed by PowerPath.....	79
	Hot swapping an HBA.....	81
	Removing an HBA.....	81
	Adding an HBA.....	82
	Configuring for R1/R2 boot failover support.....	82
	R1/R2 supported configurations.....	82
	R1/R2 boot procedure.....	83
	Configure the network interface on the R1 host.....	84
	SCSI-2 reservation support in PowerPath for Linux	84
	Enable scsi name on NetApp arrays.....	84
	Removing devices from PowerPath control.....	84
	Expand the file system on a pseudo-device.....	85
	Supporting more than 256 storage devices.....	85
	/etc/fstab options.....	86
Chapter 7	Upgrading PowerPath	89
	Preparing for upgrade.....	90
	Upgrading PowerPath.....	90
	Upgrade PowerPath from a CD.....	91
	Upgrade PowerPath from a compressed archive.....	92
	Upgrade PowerPath across Linux platforms.....	93
	Upgrade the Linux kernel in a PowerPath installed host.....	93
	Upgrade PowerPath after upgrading from one major Linux release to the next	93
	Upgrade PowerPath before upgrading to the latest SLES Service Pack.....	94
	Upgrade PowerPath before upgrading RHEL update on Linux.....	95
	Upgrade PowerPath in a SteelEye LifeKeeper cluster environment...	95
Chapter 8	Removing PowerPath	97
	Before you remove PowerPath.....	98
	Uninstall PowerPath.....	98
	Uninstall PowerPath.....	98
	Saved configuration files.....	99
	Remove the PowerPath files.....	99
Chapter 9	Troubleshooting PowerPath	101
	Recover from a failed PowerPath installation.....	102
	Troubleshoot a PowerPath upgrade.....	102
	PowerPath custom settings not preserved.....	102
	Resolve missing or deleted files.....	103
	The /etc/fstab.....	104
	Problems booting the host.....	104
	Boot issue on RHEL6.x, SLES11 SP4, RHEL7.x, SLES 12 SP3, OL6.x and OL7.x boot-from-SAN when emcpbfs_linux is used.....	104
	/etc/opt/emcpower/emcpbfs_linux config.....	104
	/etc/opt/emcpower/emcpbfs_linux remove.....	104
	/etc/opt/emcpower/emcpbfs_linux upgrade.....	105

Chapter 10	Files changed by PowerPath	107
	Linux files modified by PowerPath installation.....	108
	Files created by PowerPath installation.....	108

TABLES

1	Getting started tasks.....	10
2	Software with disabled PowerPath support.....	70
3	Present or enabled software and multipathing behavior.....	70
4	Commands to disable support.....	71
5	Monitored events.....	74
6	Supported languages.....	77

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.

- [Getting started](#) 10

Getting started

Table 1 Getting started tasks

Do you want to	Look at
install PowerPath?	Installing PowerPath on page 11 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 on page 19 for administering PowerPath in a BFS environment, for configuring PowerPath logical device, managing PowerPath, or removing PowerPath from the host.
upgrade PowerPath?	Upgrading PowerPath on page 89 for pre-requisites and procedures to upgrade PowerPath.
troubleshoot PowerPath?	Troubleshooting PowerPath on page 101 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.

Note

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

- [Before you install](#)..... 12
- [Install the RPM file](#)..... 14
- [Install from a CD](#)..... 15
- [Install in a Xen environment](#)..... 15
- [Install PowerPath when VSF is already installed](#)..... 17
- [License after the installation](#)..... 17
- [Post-installation steps](#)..... 18

Before you install

Before installing, understand the expected environment requirements as per the <https://elabnavigator.emc.com/elc/elc/home> and prepare the host system for a successful PowerPath installation.

Note

Do not use any custom scripts to install, upgrade, or uninstall the PowerPath package.

Storage System related steps

Plan your installation based on your storage system:

Procedure

- CLARiiON® storage systems: Contact CLARiiON support to plan your installation. PowerPath installation is an integral part of a CLARiiON setup and configuration procedure.
- Symmetrix® storage systems: Ensure Symmetrix is working properly before PowerPath installation.
- 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.
- VPLEX Storage Systems: Ensure VPLEX storage array is working properly before PowerPath installation.
- XtremIO Storage Systems: Ensure XtremIO 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 Host Connectivity Guide and the CLARiiON storage-system support website provides additional information.

Reviewing compatibility with the environment

Procedure

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 that the PowerPath requirements are met according to the *PowerPath Family for Linux Release Notes*.
3. Ensure that the PowerPath requirements are met as per the <https://elabnavigator.emc.com/elc/elc/home>.
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 the <https://elabnavigator.emc.com/elc/elc/home>. Improper settings could cause unwanted I/O delays.

Review host system compatibility

Procedure

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. [Recover from a failed PowerPath installation](#) on page 102 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 [Disabling dm-multipath on a SLES 11 SP4 host](#) on page 30, [Disabling dm-multipath on a RHEL 6.x or OL 6.x host](#) on page 23, or [Disabling dm-multipath on a RHEL 7.x, OL 7.x or SLES 12 SP3 host](#) on page 26 for instructions on how to disable native multipathing.

When installing on SLES 12 SP3, ensure that 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. Configure LVM,
See the section [Configuring LVM2 support](#) on page 77.
7. PowerPath uses 120 as the major number for the pseudo devices (emcpowerX) it creates. Hence, ensure that major number 120 is not used by any other driver.

Review PowerPath licensing

Procedure

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 .
 - A separate product license is not required for PowerPath Migration Enabler. The PowerPath multipathing license includes Migration Enabler functionality.

Review third-party path management software

Procedure

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.

Review the installation language

Procedure

- Select your installation language.

For example, select U.S. English.

```
En_US.utf8
```

The [Changing the language of installation messages](#) on page 76 section contains procedure for changing the installation language.

Install the RPM file

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

Before you begin

Note

All the procedures that are given for RHEL 6 are also applicable to Oracle Linux (OL) 6.x and Oracle VM (OVM) 3.x.x 3.x.

Procedure

1. Log in as root.
2. Download the PowerPath archive for the environment from [EMC Online Support](#).
3. Untar the PowerPath archive.

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

Where:

- *<version>* is the product version. For example, 6.3.0.00.00.
- *<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 SLES12SP3 or RHEL6.

6. Type the 24-character alphanumeric sequence from the License Key Card, and follow the onscreen instructions.

Install from a CD

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

Procedure

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

```
mkdir /cdrom
```

4. Mount the PowerPath CD on `/cdrom`.

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

5. Change directories to the operating system directory.

- For SLES,

```
cd /cdrom/LINUX/pp<version>/sles
```
- For RHEL and Asianux,

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

6. Copy the RPM package into a temporary folder.

7. Install PowerPath.

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

Where:

- `<version>` is the product version. For example, 6.3.0.00.00.
- `<build>` is the software build number.
- `<platform>` is the Linux distribution, for example SLES12SP3 or RHEL6.

Install 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.

Procedure

1. Log in as root.
2. Verify the following Xen tools and kernel RPMs are installed on the host.

If not present, then install them before installing PowerPath.

For more information about Xen kernels, refer to the vendor documentation for the operating system.

3. Verify that the system is running the Xen kernel.

```
uname -a
```

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

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

6. Create a Xen virtual machine.

For more information, refer to the vendor documentation for the operating system.

7. 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.

8. Locate the following line in the configuration file.

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

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

For example, add the following line:

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

Where:

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

10. Start the virtual machine.

```
xm create -c <vm_name>
```

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.

For example,

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.

Install PowerPath when VSF is already installed

Installing PowerPath after Veritas Storage Foundation (VSF) 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.

Procedure

1. Log in as root.
2. Exclude EMC arrays from DMP path management.
 - For Symmetrix arrays,
`vxddladm excludearray libname=libvxemc.so`
 - For VNX OE and CLARiiON arrays,
`vxddladm excludearray libname=libvxCLARiiON.so`
 - For VPLEX arrays,
`vxddladm excludearray libname=libvxInvista.so`
 - For XtremIO devices:
`vxddladm excludearray libname=libvxXtremio.so`

3. Run **vxctl enable**.

Verify that above EMC arrays are excluded from DMP path management. Run "vxdisk list". The list of devices under VSF is displayed.

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

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

6. Ensure that PowerPath claims the devices.

```
vxctl enable
```

Run "vxdisk list". The list of devices under VSF is displayed.

License 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.

Procedure

1. Register the PowerPath license.

```
emcpreg -install
```

a. Type **y**

b. Type 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 onscreen instructions.

2. Start the PowerPath service.

Use the script only to load and unload PowerPath modules.

- For RHEL 6, OL6 or SLES 11 SP4
`/etc/init.d/PowerPath start`
 - For RHEL7, OL7 or SLES 12 SP3
`systemctl start PowerPath.service`
-

Note

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

3. Change the load-balancing and failover policy if needed.

```
powermt set policy
```

By default, PowerPath sets the optimal policy based on the license installed.

4. Ensure that the PowerPath configuration settings are persisted across host restarts.

```
powermt save
```

Results

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.

Procedure

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 75.

2. Verify that all the PowerPath capabilities, including the load balancing and failover policies, are given as expected for your license.

```
powermt display dev=all
```

CHAPTER 3

Administering in a Boot from SAN environment

You can configure a PowerPath pseudo device (emcpowerX) 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. Depending on whether the host has BIOS or UEFI firmware, follow the relevant section to configure PowerPath to work in a Boot from SAN setup.

• Before configuring for a root device.....	20
• Configuring PowerPath in a RHEL 6.x or OL 6.x Boot from SAN setup in a BIOS host.....	20
• Configuring PowerPath in a RHEL 7.x, OL 7.x or SLES 12 SP3 Boot from SAN setup in a BIOS host.....	23
• Configuring PowerPath in a SLES 11 SP4 Boot from SAN setup in a BIOS host.....	26
• Upgrading PowerPath in a RHEL 6.x or OL 6.x Boot from SAN setup in a BIOS host.....	31
• Upgrading PowerPath in a RHEL 7.x, OL 7.x or SLES 12 SP3 Boot from SAN setup in a BIOS host.....	33
• Upgrading PowerPath in a SLES 11 SP4 Boot from SAN setup in a BIOS host.....	35
• Upgrading the Linux kernel in a Boot from SAN setup in a BIOS host.....	37
• Uninstalling PowerPath in a RHEL 7.x, OL 7.x or SLES 12 SP3 Boot from SAN setup in a BIOS host.....	42
• Uninstalling PowerPath in a SLES 11 SP4 Boot from SAN setup in a BIOS host.....	44
• Configure PowerPath in a RHEL 6.x or OL 6.x Boot from SAN setup in an UEFI host.....	45
• Configure PowerPath in a RHEL 7.x, OL 7.x or SLES 12 SP3 Boot from SAN setup in an UEFI host.....	48
• Upgrading PowerPath in a RHEL 6.x or OL 6.x Boot from SAN setup in an UEFI host.....	51
• Upgrading PowerPath in a RHEL 7.x, OL 7.x or SLES 12 SP3 Boot from SAN setup in an UEFI host.....	53
• Uninstalling PowerPath in a RHEL 6.x or OL 6.x Boot from SAN setup in an UEFI host.....	55
• Uninstalling PowerPath in a RHEL 7.x, OL 7.x or SLES 12 SP3 Boot from SAN setup in an UEFI host.....	57

Before configuring for a root device

- For SLES 11 SP4, 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.

⚠ 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 77 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 6.x or OL 6.x Boot from SAN setup in a BIOS host

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

Procedure

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 boot from SAN settings 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. Add 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 `ext3` 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 23.
 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 11 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, using 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 ramdisk 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 the following 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 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 filter device nodes are not listed in the output.

- f. Build the ramdisk image to include PowerPath modules.

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

This generates a new ramdisk image named:

```
initramfs-PP-<kernel_version>.<arch>.img
```

- g. Add a new boot entry with new ramdisk 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 ramdisk 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 ramdisk 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 ramdisk image.
17. Update the ramdisk 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

Procedure

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 ramdisk image so that dm-multipath is excluded from ramdisk image.


```
dracut /boot/initramfs-wo-DM-$(uname -r).img $(uname -r).
```
4. Add a new boot entry with the new ramdisk image file created in previous step in the `/boot/grub/menu.lst` file.
5. Restart the host with the ramdisk 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, OL 7.x or SLES 12 SP3 Boot from SAN setup in a BIOS host

To configure a PowerPath root device using the LVM on a RHEL 7.x, OL 7.x or SLES 12 SP3 host.

Procedure

1. Ensure that the fibre cable channel 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 BIOS on the first HBA and disable it on all others.
3. Ensure that 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 SP3, 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 500 MB (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 addition LUNs to the host.
7. Install and configure PowerPath.
8. Complete the configuration of the boot LUN under PowerPath.
 - To automatically complete the configuration.
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/emcpoweral /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 filter device nodes are not listed in the output.

f. Build the ramdisk image to include 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 SP3:

```
#dracut /boot/initrd-PP-$(uname -r) $(uname -r)
```

This generates a new ramdisk image named:

For RHEL 7.x/OL 7.x:

```
initramfs-PP-<kernel_version>.<arch>.img
```

For SLES 12 SP3

```
initrd-PP-<kernel_version>
```

g. Add a new boot entry with ramdisk image file name in `/etc/grub.d/40_custom`.

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

```
#!/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 ramdisk image automatically.

For RHEL 7.x/OL 7.x:

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

```
#grub2-set-default 'PowerPath - Red Hat Enterprise Linux Server'
```

For SLES 12 SP3:

```
#grub2-mkconfig -o /boot/grub2/grub.cfg
#grub2-set-default 'PowerPath - SLES 12-SP3'
```

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

Disabling dm-multipath on a RHEL 7.x, OL 7.x or SLES 12 SP3 host

Procedure

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 ramdisk image so that dm-multipath is excluded from ramdisk image.

```
dracut /boot/initramfs-wo-DM-$(uname -r).img $(uname -r)
```
4. Add a new boot entry with the ramdisk image file created in previous step in the `/etc/grub.d/40_custom` file, and then update the `grub.cfg` file.
5. Restart the host with the new ramdisk 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 in a BIOS host

To configure a PowerPath Boot from SAN device using the LVM on a SLES 11 SP4 host:

Procedure

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 on a SLES 11 SP4 host](#) on page 30.
 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 11 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 ramdisk 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                                0      0
defaults                            0      0
/dev/sles11/
root                                /
ext3                                acl,user_xattr                    1      1
/dev/disk/by-id/scsi-3600601605c7019002820bd32b767de11-
part1
ext3                                /boot                            1      2
acl,user_xattr
proc                                /
proc                                proc                                0      0
defaults                            0      0
sysfs                               /
sys                                sysfs                               0      0
noauto                              0      0
debugfs                             /sys/kernel/
debug                                debugfs                            0      0
noauto                              0      0
usbfs                               /proc/bus/
usb                                usbfs                               0      0
noauto                              0      0
devpts                              /dev/
pts                                devpts                             0      0
mode=0620,gid=5
```

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                                0      0
defaults                            0      0
/dev/sles11/
root                                /
ext3                                acl,user_xattr                    1      1
#/dev/disk/by-id/scsi-3600601605c7019002820bd32b767de11-
part1
ext3                                /boot                            1      2
acl,user_xattr
/dev/emcpowera1                     /
boot                                ext3                             1      2
acl,user_xattr
proc                                /proc                             0      0
proc                                defaults                         0      0
sysfs                               /sys
```

```

sysfs          noauto          0      0
debugfs        /sys/kernel/
debug          debugfs
noauto         0      0
usbfs          /proc/bus/
usb            usbfs          0      0
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 ramdisk image with the PowerPath feature to include PowerPath modules.

```

mkinitrd -i /boot/initrd-pp-3.0.101-63.1-default -k
/boot/vmlinuz-3.0.101-63.1-default -f powerpath

```

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

```

SUSE Linux Enterprise Server 11 SP4 - 3.0.101-63.1
root (hd0,0)
kernel /vmlinuz-3.0.101-63.1-default root=/dev/system/root
resume=/dev/system/swap splash=silent crashkernel=256M-:128M
showopts
initrd /initrd-pp-3.0.101-63.1-default

SUSE Linux Enterprise Server 11 SP4 - 3.0.101-63.1
root (hd0,0)
kernel /vmlinuz-3.0.101-63.1-default root=/dev/system/root
resume=/dev/system/swap splash=silent crashkernel=256M-:128M
showopts
initrd /initrd-3.0.101-63.1-default

```

15. Reboot the host, enable the BIOS on the other HBAs and select LUN 0 on those HBAs. After boot, the new ramdisk 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 ramdisk 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.

- a. Take a backup of the existing ramdisk image.
 - b. Update the ramdisk image by repeating step e in the manual steps detailed in the previous procedure.
-

Disabling dm-multipath on a SLES 11 SP4 host

Procedure

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 ramdisk image so that dm-multipath is excluded from `initrd`.

```
mkinitrd -i /boot/initrd-wo-DM-<kernel> -k /boot/vmlinuz-
<kernel>
```

For example,

```
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 ramdisk image file created in the previous step in `/boot/grub/menu.lst` file.
5. Restart the host with the new ramdisk image and ensure dm-multipath does not have any devices that are configured.

```
multipath -ll
```

Should not return any dm devices.

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

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

- [Upgrade using the emcpbfs_linux script](#) on page 31
- [Upgrade using manual steps](#) on page 31

Upgrade using the emcpbfs_linux script

Procedure

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 ramdisk image in `/boot/grub/menu.lst` and sets the proper `'default'` value.

Upgrade using manual steps

Procedure

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 ramdisk image without the PowerPath modules as the default entry.

5. Reboot the host.

6. Run

```
/etc/init.d/PowerPath start to start 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.

```
global_filter = [ "a./." ]
```

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

11. Delete the old ramdisk image containing PowerPath modules that was created in [Configuring PowerPath in a RHEL 6.x or OL 6.x Boot from SAN setup in a BIOS host](#) on page 20.

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

12. Build the ramdisk image to include upgraded PowerPath modules.

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

This generates a new ramdisk image file named:

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

13. The `/boot/grub/menu.lst` already contains an entry with ramdisk image created during initial setup, see [Configuring PowerPath in a RHEL 6.x or OL 6.x Boot from SAN setup in a BIOS host](#) on page 20.

Make it the default entry. Add a boot entry with ramdisk image containing PowerPath modules, if it is missing.

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

Upgrading PowerPath in a RHEL 7.x, OL 7.x or SLES 12 SP3 Boot from SAN setup in a BIOS host

To upgrade PowerPath in a Boot from SAN setup on a RHEL 7.x, OL 7.x and SLES 12 SP3 host, use one of the following methods:

- [Upgrade using the emcpbfs_linux script](#) on page 33
- [Upgrade using manual steps](#) on page 33

Upgrade using the `emcpbfs_linux` script

Procedure

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 ramdisk image in `/boot/grub/menu.lst` and sets the proper `'default'` value.

Upgrade using manual steps

Procedure

1. Disable the 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./" ]
```

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

For RHEL 7.x/OL 7.x:

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

For SLES 12 SP3:

```
#grub2-set-default 'SLES 12-SP3'
```

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 ramdisk image containing PowerPath modules that was created in [Configuring PowerPath in a RHEL 7.x, OL 7.x or SLES 12 SP3 Boot from SAN setup in a BIOS host](#) on page 23.

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

12. Build the ramdisk image to include upgraded PowerPath modules.

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

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

For SLES 12 SP3:

```
#dracut /boot/initrd-PP-$(uname -r) $(uname -r)
```

13. The `/etc/grub.d/40_custom` already contains an entry with ramdisk image created during initial setup, see [Configuring PowerPath in a RHEL 7.x, OL 7.x or SLES 12 SP3 Boot from SAN setup in a BIOS host](#) on page 23.

Make it the default entry. Add a boot entry with ramdisk image containing PowerPath modules, if it is missing.

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

Upgrading PowerPath in a SLES 11 SP4 Boot from SAN setup in a BIOS host

To upgrade PowerPath in a Boot from SAN setup on a SLES 11 SP4 host, use one of the following methods:

- [Upgrade using the `emcpbfs_linux` script](#) on page 35
- [Upgrade using manual steps](#) on page 36

Upgrade using the `emcpbfs_linux` script

Procedure

1. Run PowerPath upgrade.

`/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. 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.

Note

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

Upgrade using manual steps

Procedure

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.

For example,

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

```
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 acl,user_xattr 1 2
/dev/emcpowera1 /boot ext3 acl,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 ramdisk image containing PowerPath modules that was created in step 12 in [Configuring PowerPath in a SLES 11 SP4 Boot from SAN setup in a BIOS host](#) on page 26.

```
# rm -f /boot/initrd-pp-3.0.101-63.1-default
```

12. Create a new ramdisk image to include upgraded PowerPath modules.

```
mkinitrd -i /boot/initrd-pp-3.0.101-63.1-default -k
/boot/vmlinuz-3.0.101-63.1-default -f powerpath
```

13. The `/boot/grub/menu.lst` already contains an entry with ramdisk image created during initial setup, see [Configuring PowerPath in a SLES 11 SP4 Boot from SAN setup in a BIOS host](#) on page 26.

Make it the default entry. Add a boot entry with ramdisk image containing PowerPath modules, if it is missing.

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

Upgrading the Linux kernel in a Boot from SAN setup in a BIOS host

Before you begin

Note

Before upgrading the kernel on RHEL 6, if the root file system 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, Oracle Linux, or SLES 12 SP3 host, PowerPath automatically restarts 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 SLES 11 SP4, then complete one of the following steps to ensure that PowerPath functions normally after the upgrade:

- [Upgrade using emcpbfs_linux script](#) on page 37
- [Upgrade using manual steps](#) on page 38

Upgrade using `emcpbfs_linux` script

Procedure

1. Run PowerPath upgrade.

`/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. 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.

Note

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

Upgrade using manual steps

Procedure

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

For example,

```
/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
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. Verify that the filter device nodes are not listed in the output.

4. Upgrade the kernel as described in Oracle or Red Hat documentation and then restart the host.
5. Start PowerPath.

- For RHEL 6.x, OL 6.x or SLES 11 SP4
`/etc/init.d/PowerPath start`
- For RHEL 7.x, OL 7.x or SLES 12 SP3
`systemctl start PowerPath.service`

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:

- Copy contents of `/etc/emc/archive` to another folder such as `/tmp/pp_config_backup`.
- Uninstall the new version of PowerPath.
- Delete contents of `/etc/emc/archive` folder since that now contains configuration files updated by the new PowerPath version.
- Install the old PowerPath version.
- Copy the backed up contents from step a to `/etc/emc/archive` folder.
- Start PowerPath.

While starting, PowerPath restores the config files from `/etc/emc/archive` and the old settings are 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. Unmount and mount `/boot` so that it is mounted on pseudo device.

```
umount /boot/
```

```
mount /boot/
```

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

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

9. Build the ramdisk image to include upgraded PowerPath modules.

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

This generates a new ramdisk image file named:

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

10. Add a new boot entry which has the new ramdisk image built in step 9 in `/boot/grub/menu.lst` and make it the default entry.
11. Reboot the host so that the new ramdisk image containing updated PowerPath modules take effect.

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

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](#) on page 40
- [Uninstall using manual steps](#) on page 41

Uninstall using emcpbfs_linux_script

Before you begin

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`.

Procedure

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 ramdisk 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

Procedure

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 ramdisk image.

If the original ramdisk image 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 ramdisk image file:

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

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

Uninstalling PowerPath in a RHEL 7.x, OL 7.x or SLES 12 SP3 Boot from SAN setup in a BIOS host

Use one of the following to uninstall PowerPath from a Boot from SAN setup on RHEL 7.x, OL 7.x or SLES 12 SP3 host:

- [Uninstall using emcpbfs_linux_script](#) on page 42
- [Uninstall using manual steps](#) on page 43

Uninstall using `emcpbfs_linux_script`

Before you begin

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`.

Procedure

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 ramdisk 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

Procedure

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 ramdisk 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 ramdisk image file.

For example on RHEL7.x/OL7.x:

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

For example on SLES 12 SP3:

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

5. 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.
```

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

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

7. Set the original boot entry as the default entry.

For RHEL 7.x/OL 7.x:

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

For SLES 12 SP3:

```
# grub2-set-default 'SLES 12 SP3'
```

8. Reboot the host and then uninstall PowerPath.

Uninstalling PowerPath in a SLES 11 SP4 Boot from SAN setup in a BIOS host

Uninstall PowerPath using one of the following methods:

- [Uninstall using emcpbfs_linux_script](#) on page 44
- [Uninstall using manual steps](#) on page 45

Uninstall using `emcpbfs_linux_script`

Before you begin

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`.

Procedure

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 remove` adds a new boot entry for the new ramdisk 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

Procedure

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 ramdisk 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 ramdisk image created in the previous step in `/boot/grub/menu.lst` and make it the default entry.
8. Reboot the host.
9. Uninstall PowerPath.

Configure PowerPath in a RHEL 6.x or OL 6.x Boot from SAN setup in an UEFI host

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

Procedure

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 boot from SAN settings 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. Add LUN 0 as a boot device in the first HBA's configuration page.
5. Boot the host and start RHEL 6.x or OL 6.x installation.
6. Select GPT as the partitioning method and create a custom partition setup on the target storage device with three partitions as shown below:
 - First as a `/boot` partition formatted as an `ext3` or `ext4` file system with at least 500 MB in size.
 - Second as a `/boot/efi`, EFI system partition (ESP) formatted as a `vfat` file system with at least 200 MB in size.
 - Third 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` or `ext4` and the other as swap and designate the mount points for each volume.

9. Ensure GRUB is the boot loader selected and uses GUID partition table (GPT).
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 23.

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 11 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, using one of the following methods:
 - Using the `emcpbfs_linux` script.
-

Note

In a multi-boot environment ensure that the default value in the `/boot/efi/EFI/redhat/grub.conf` 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 ramdisk 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 the following manual steps.
 - a. If necessary, save a backup of the `/etc/fstab` file. Then, edit `/etc/fstab` to mount `/boot/efi` by pseudo-device name (`dev/emcpowera`) :

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

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

```
#umount /boot/efi

#mount /boot/efi
```

- 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 filter device nodes are not listed in the output.

- f. Build the ramdisk image to include PowerPath modules.

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

This generates a new ramdisk image named:

```
initramfs-PP-<kernel_version>.<arch>.img
```

- g. Add a new boot entry with ramdisk image name in `/boot/efi/EFI/redhat/grub.conf` 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=latacyrheb-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 boot from SAN settings for the other HBAs and add LUN 0 on those HBAs.

After boot, the new ramdisk 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 ramdisk image needs to be updated. Either run `emcpgfs_linux` script with config option after deleting the `/etc/emc/.emcpgfs_linux_done` file or execute the following steps manually.

16. Take a backup of the existing ramdisk image.
17. Update the ramdisk image by following step f in the manual steps detailed in the previous procedure.

Configure PowerPath in a RHEL 7.x, OL 7.x or SLES 12 SP3 Boot from SAN setup in an UEFI host

To configure a PowerPath root device using the LVM on a RHEL 7.x, OL 7.x or SLES 12 SP3 host.

Procedure

1. Ensure that the fibre cable channel 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 boot from SAN settings on the first HBA and disable it on all others.
3. Ensure that 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 SP3, ensure native multipathing is not automatically activated during the OS installation as there is not 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/efi` as basic partition of 500 MB (vfat), and Logical Volumes for `/` (root) and swap on Volume Group created from remaining space of the disk using GPT as the partitioning method.
6. After OS Installation is complete, add additional paths to the boot LUN and attach addition LUNs to the host.
7. Install and configure PowerPath.
8. Complete the configuration of the boot LUN under PowerPath.

- To automatically complete the configuration.

```
Run /etc/opt/emcpower/emcpbfs_linux config PowerPath
```

- To manually complete the configuration.

- a. Edit `/etc/fstab` to mount `/boot/efi` by pseudo-device name (for example, `/dev/emcpowerX`).

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

- b. Unmount and mount `/boot/efi` so that it is mounted on pseudo device.

```
#umount /boot/efi
```

```
#mount /boot/efi
```

- 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 `lvm diskscan` to ensure that the filter field is working correctly. Verify that the filter device nodes are not listed in the output.

- f. Build the ramdisk image to include 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 SP3:

```
#dracut /boot/initrd-PP-$(uname -r) $(uname -r)
```

This generates a new ramdisk image named:

For RHEL 7.x/OL 7.x:

```
initramfs-PP-<kernel_version>.<arch>.img
```

For SLES 12 SP3:

```
initrd-PP-<kernel_version>
```

- g. Add a new boot entry with ramdisk image file name in `/etc/grub.d/40_custom`.

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

```
#!/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
    linuxefi /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
    initrdefi /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 ramdisk image automatically.

For RHEL 7.x/OL 7.x:

```
#grub2-mkconfig -o /boot/efi/EFI/redhat/grub.cfg
```

```
#grub2-set-default 'PowerPath RedHat Enterprise Linux Server'
```

For SLES 12 SP3:

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

```
#grub2-set-default 'PowerPath - SLES 12 SP3'
```

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

Upgrading PowerPath in a RHEL 6.x or OL 6.x Boot from SAN setup in an UEFI host

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

- [Upgrade using the emcpbfs_linux script](#) on page 51
- [Upgrade using manual steps](#) on page 52

Upgrade using the `emcpbfs_linux` script

Procedure

1. Run PowerPath upgrade.

`/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/efi'` 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 ramdisk image in `/boot/efi/EFI/redhat/grub.conf` and sets the proper 'default' value.

Upgrade using manual steps

Procedure

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/efi vfat
defaults 1 2
/dev/emcpowera1 /boot/efi vfat 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/efi/EFI/redhat/grub.conf` and make the entry that contains the ramdisk image without the PowerPath modules as the default entry.
5. Reboot the host.
6. After the host boots, run PowerPath.

```
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/efi` by pseudo (emcpower) devices.

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

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

```
#umount /boot/efi
```

```
#mount /boot/efi
```

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 ramdisk image containing PowerPath modules that was created in step 12 in [Configure PowerPath in a RHEL 6.x or OL 6.x Boot from SAN setup in an UEFI host](#) on page 45.

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

12. Build the ramdisk image to include upgraded PowerPath modules.

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

This generates a new ramdisk image file named:

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

13. The `/boot/efi/EFI/redhat/grub.conf` already contains an entry with ramdisk image created during initial setup, see [Configure PowerPath in a RHEL 6.x or OL 6.x Boot from SAN setup in an UEFI host](#) on page 45.
Make it the default entry. Add a boot entry with ramdisk image containing PowerPath modules, if it is missing.
14. Reboot the host so that the new ramdisk image containing updated PowerPath modules takes effect.

Upgrading PowerPath in a RHEL 7.x, OL 7.x or SLES 12 SP3 Boot from SAN setup in an UEFI host

To upgrade PowerPath in a Boot from SAN setup on a RHEL 7.x, OL 7.x or SLES 12 SP3 host, use one of the following methods:

- [Upgrade using the `emcpbfs_linux` script](#) on page 53
- [Upgrade using manual steps](#) on page 54

Upgrade using the `emcpbfs_linux` script

Procedure

1. Run PowerPath upgrade.

`/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/efi'` 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.

Note

`emcpbfs_linux config/remove` adds a new boot entry for the new ramdisk image in `/boot/efi/EFI/redhat/grub.cfg` on RHEL 7.x/OL 7.x or `/boot/grub2/grub.cfg` on SLES 12 SP3 file and sets the proper 'default' value.

Upgrade using manual steps

Procedure

1. Disable the 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/efi vfat
defaults 1 2
/dev/emcpowera1 /boot/efi vfat 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/*/" ]
```

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

For RHEL 7.x/OL 7.x:

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

For SLES 12 SP3:

```
# grub2-set-default 'SLES 12 SP3'
```

5. Reboot the host.
6. After the host boots, run PowerPath.

```
rpm -Uvh <new_powerpath_package> to upgrade PowerPath.
```

7. Run

```
systemctl start PowerPath.service to start PowerPath
```

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

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

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

```
#umount /boot/efi
```

```
#mount /boot/efi
```

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 ramdisk image containing PowerPath modules that was created in [Configure PowerPath in a RHEL 7.x, OL 7.x or SLES 12 SP3 Boot from SAN setup in an UEFI host](#) on page 48.

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

12. Build the ramdisk image to include upgraded PowerPath modules.

- For example on RHEL 7.x or OL 7.x:

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

- For SLES 12 SP3

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

13. The `/etc/grub.d/40_custom` already contains an entry with ramdisk image created during initial setup, see [Configure PowerPath in a RHEL 7.x, OL 7.x or SLES 12 SP3 Boot from SAN setup in an UEFI host](#) on page 48.

Make it the default entry. Add a boot entry with ramdisk image containing PowerPath modules, if it is missing.

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

Uninstalling PowerPath in a RHEL 6.x or OL 6.x Boot from SAN setup in an UEFI host

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](#) on page 55
- [Uninstall using manual steps](#) on page 56

Uninstall using `emcpbfs_linux_script`

Before you begin

Note

In a multi-boot environment ensure that the default value in the `/boot/efi/EFI/redhat/grub.conf` is set properly with the current boot kernel before running `emcpbfs_linux config/remove`.

Procedure

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/efi'` entry only in the `/etc/fstab` file.

3. Reboot the host.

Note

`emcpbfs_linux config/remove` adds a new boot entry for the new ramdisk image in `boot/efi/EFI/redhat/grub.conf` 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

Procedure

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/efi` by UUID.

```
UUID=110223d8-f5f3-4058-812f-16d9e9ee7eab /boot/efi vfat
defaults 1 2
#/dev/emcpower1 /boot/efi vfat defaults 1 2
```

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

If the original ramdisk image 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 ramdisk image file:

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


- c. Add a new boot entry for the new ramdisk image created `boot/efi/EFI/redhat/grub.conf`.
 - d. Reboot the host and select the boot loader entry with ramdisk image in the previous step.
6. Uninstall PowerPath.

Uninstalling PowerPath in a RHEL 7.x, OL 7.x or SLES 12 SP3 Boot from SAN setup in an UEFI host

Use one of the following to uninstall PowerPath from a Boot from SAN setup on RHEL 7.x, OL 7.x or SLES 12 SP3 host:

- [Uninstall using `emcpbfs_linux_script`](#) on page 57
- [Uninstall using manual steps](#) on page 58

Uninstall using `emcpbfs_linux_script`

Before you begin

Note

In a multi-boot environment ensure that the default value in the `/boot/efi/EFI/redhat/grub.cfg` on RHEL 7.x/OL7.x or `/boot/grub2/grub.cfg` on SLES 12 SP3 file is set properly with the current boot kernel before running `emcpbfs_linux_config/remove`.

Procedure

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/efi'` entry only in the `/etc/fstab` file.

3. Reboot the host.

Note

`emcpbfs_linux config/remove` adds a new boot entry for the new ramdisk image in `/boot/efi/EFI/redhat/grub.cfg` on RHEL 7.x/OL7.x or `/boot/grub2/grub.cfg` on SLES 12 SP3 file 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**Procedure**

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/efi` by UUID.

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

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

4. If the original ramdisk 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 ramdisk image file.

For example on RHEL7.x/OL7.x:

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

For example on SLES 12 SP3:

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

5. 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.
```

6. Recreate the `grub.cfg` file to remove the PowerPath boot entry.

For RHEL 7.x/OL 7.x:

```
# grub2-mkconfig -o /boot/efi/EFI/redhat/grub.cfg
```

For SLES 12 SP3:

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

7. Set the original boot entry as the default entry.

For RHEL 7.x/OL 7.x:

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

For SLES 12 SP3:

```
# grub2-set-default 'SLES 12 SP3'
```

8. Reboot the host and then uninstall PowerPath.

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](#)..... 62
- [Before you add or remove paths](#)..... 63
- [Adding new paths to a PowerPath logical device](#)..... 64
- [Removing some of the paths to PowerPath logical devices](#)..... 65

Dynamic LUN addition and removal

PowerPath supports dynamically adding and removing LUNs in all supported operating systems.

Note

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

Follow this procedure to add a new LUN dynamically to PowerPath. All I/O to the SCSI devices should be quiesced before this procedure as it involves rescan of the SCSI buses.

Procedure

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>**, **<#>**, or **hba=<#> tgt=<#>**.

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 specified, 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

Follow this procedure to remove a LUN that is already configured and managed in PowerPath.

Procedure

1. Identify the device to be removed.

```
/etc/opt/emcpower/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.

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

Where *<pseudo-list>* can be `emcpowerX emcpowerY`.

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

Resizing a LUN

Follow this procedure to expand the size of a LUN already configured and managed in PowerPath.

Procedure

1. Perform online expansion of the LUN on the array by following the steps listed in the relevant array documentation.
2. For all the native devices associated with the LUN, run the following command.

```
echo 1 > /sys/block/sd*/device/rescan
```
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 fail if the affected logical devices are in use (either just marked alive or with I/O in progress). This is noted 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 the 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 to a PowerPath logical device

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

Before you begin

Note

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

⚠ CAUTION

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

Procedure

1. Confirm the current configuration.
`powermt display`
2. Confirm the configuration of the logical devices to which new paths are added.
`powermt display dev=all`
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 *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 the vendor documentation.
 - c. Attach cables.
 - d. Adjust zoning on the 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.
 - Automatically scan and configure the additional paths.
`/etc/opt/emcpower/emcplun_linux scan <option> [-noprompt]`
Where <option> can be `hba`, `iscsi`, `open-fcoe`, `all`,
`aid=<array id>`, `<#>`, `hba=<#>` `tgt=<#>`.
 - 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. Configure the new paths under PowerPath control.


```
powermt config
```

- c. Inspect the new PowerPath configuration.

```
powermt display dev=all
```

The new paths should be displayed with the alive state. Run `powermt restore` to test all paths. Scan operating system error logs to ensure that no errors are logged against the new paths.

- d. Correct any issues that are detected.

- e. Save the new configuration.

```
powermt save
```

Removing some of the paths to PowerPath logical devices

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

CAUTION

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.

Procedure

1. Stop all applications that use the native device that is 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 this 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 is interrupted. (If a PowerPath device is open, `powermt remove` does not remove that device or the last path to that device).

3. Confirm the configuration of the logical devices from which paths are removed.

```
powermt display dev=all
```

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. Identify the PowerPath HBA number that is associated with the paths to be removed, and identify the native device that is associated with the paths to be removed.

```
powermt display dev=all
```

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

8. Obtain bus, TID, and LUN.

```
pp_inq -bt1
```

Save the output of the command.

9. If using character I/O (sg) devices on Linux, run `sg_map -i -x` and then save the output.
10. If a native device that points to the path to be removed is in use, modify all applications that are configured to use that device, so that they can use another device.
11. Set the mode to standby for all paths to be removed.

```
powermt set mode=standby
```

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 issues 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. Confirm that the paths that are listed as dead are the ones you intended to remove.

```
powermt restore
```

Note down all the native device names that the paths that are listed as dead are the ones you intended to remove.

14. Remove the logical device, specifying both the `hba` and the `dev` arguments.

```
powermt remove
```

CAUTION

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

⚠ CAUTION

Ensure that no applications try to use the native device that is being removed. If an application accesses a native path 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 that are associated with the paths that were removed (the native devices names that are noted in Step 11).
-

Note

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. Remove the stale native devices from the kernel.

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

17. Inspect the new PowerPath configuration.

- a. Check the path state.

```
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 that are associated with the affected logical devices should reflect the 'alive' state.

18. Correct any issues that are detected before saving the PowerPath configuration or using the new logical devices.
19. Save the new configuration.

```
powermt save
```


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](#)..... 70
- [PowerPath co-management behavior](#) 70

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). 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.

Table 2 Software with disabled PowerPath support

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 co-management behavior

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

Because PowerPath is not supported by 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>`.

Table 3 Present or enabled software and multipathing behavior

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

Because PowerPath and the third-party software cannot co-manage devices, do not initiate co-management by running 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 remain disabled across restart.

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>`.

Table 4 Commands to disable support

Before you install or enable this software	Run this command
IBM SDD	<code>powermt unmanage class=ess</code>
HP StorageWorks XP AutoPath	<code>powermt unmanage class=hpxp</code>

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.

• Configure PowerPath for event monitoring	74
• Configuring the Management Daemon	75
• Changing the language of installation messages	76
• Saving new configurations	77
• Configuring LVM2 support	77
• Hot swapping an HBA	81
• Configuring for R1/R2 boot failover support	82
• SCSI-2 reservation support in PowerPath for Linux	84
• Enable scsi name on NetApp arrays	84
• Removing devices from PowerPath control	84
• Expand the file system on a pseudo-device	85
• Supporting more than 256 storage devices	85
• /etc/fstab options	86

Configure PowerPath for event monitoring

The PowerPath management daemon provides an event for multipathing events listed in the `emcp_mond.conf` file. This includes the ability for the host to be connected to a PowerPath Management Appliance.

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 1 minute.
- All paths to a volume are dead for 5 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 specified delay time elapses, a trap is generated to report the event. If canceling 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.

The following table lists the events, event descriptions, and event IDs and time thresholds. It also shows whether for each case an SNMP trap is issued.

Table 5 Monitored events

Monitored event	Event cause	Event ID	SNMP Trap
Path is Dead	Path is dead for 1 minute.	0x2,0x4,60	Yes
	Path is dead and then recovers within 1 minute.	0x2,0x4,60	No
All Paths Dead	All paths are dead for 5 minutes.	0x20,0x40,300	Yes
	All paths are dead and then recover within 5 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 76.

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 canceling event, and the delay time in seconds for the event. Commas delimit these three event parameters. 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

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. Configure `snmpd` persistence across restart.
 - 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 that is already located in `/etc/snmp/snmp.conf`, the previous file is overwritten.

- b. Verify that the `snmpd` service is enabled by using `chkconfig` command; if not, then enable the service.
- c. Start the `snmpd` daemon.

```
/etc/init.d/snmpd start
```

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

5. Check the status.

```
ps -ef | grep snmpd
```

For example,

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

6. Restart the management daemon.

- On RHEL7, OL7 or SLES 12 SP3,
`systemctl restart EmcpMond.service`
- On all other kernels,
`/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.

Procedure

1. Open the `/etc/opt/emcpower/emcp_mond.conf` file.
2. Change the time threshold.

For example, to change the time threshold 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 management daemon.
 - On RHEL7, OL7 or SLES 12 SP3,
`systemctl restart EmcpMond.service`
 - On all other kernels,
`/etc/init.d/emcp_mond.sh stop`
`/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.

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 6 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	
Portugese (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.

Before you begin

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 [Configure PowerPath in a RHEL 6.x or OL 6.x Boot from SAN setup in an UEFI host](#) on page 45 or [Configure PowerPath in a RHEL 7.x, OL 7.x or SLES 12 SP3 Boot from SAN setup in an UEFI host](#) on page 48.

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.

Procedure

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	<pre>filter = ["a/emcpower.*/", "r/sd.*/", "r/disk.*/"]</pre>
SLES 12 SP3	<pre>global_filter = ["a/emcpower.*/", "r/sd.*/", "r/disk.*/"]</pre>
RHEL 6.x	<pre>filter = ["a/emcpower.*/", "r/sd.*/", "r/disk.*/"]</pre>

Kernel	Filter
RHEL 7.x	<pre>global_filter = ["a/emcpower.*/", "r/sd.*/", "r/disk.*/"]</pre>
Oracle Linux 6.x	<pre>filter = ["a/emcpower.*/", "r/sd.*/", "r/disk.*/"]</pre>
Oracle Linux 7.x	<pre>global_filter = ["a/emcpower.*/", "r/sd.*/", "r/disk.*/"]</pre>

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 [Configure PowerPath in a RHEL 6.x or OL 6.x Boot from SAN setup in an UEFI host](#) on page 45 or [Configure PowerPath in a RHEL 7.x, OL 7.x or SLES 12 SP3 Boot from SAN setup in an UEFI host](#) on page 48.

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:

Procedure

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.

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 6.x	<pre>filter = ["a /dev./*/by-id/scsi-SSEAGATE_ST373207LC_3KT4PMAN.* ", "a/emcpower.* ", "r/sd.* ", "r/disk.* "]</pre>
SLES SP4	<pre>filter = ["a /dev./*/by-id/scsi-SSEAGATE_ST373207LC_3KT4PMAN.* ", "a/emcpower.* ", "r/sd.* ", "r/disk.* "]</pre>
Oracle Linux 6.x	<pre>filter = ["a /dev./*/by-id/scsi-SSEAGATE_ST373207LC_3KT4PMAN.* ", "a/emcpower.* ", "r/sd.* ", "r/disk.* "]</pre>
RHEL 7.x	<pre>global_filter = ["a /dev./*/by-id/scsi-SSEAGATE_ST373207LC_3KT4PMAN.* ", "a/emcpower.* ", "r/sd.* ", "r/disk.* "]</pre>
SLES 12 SP3	<pre>global_filter = ["a /dev./*/by-id/scsi-SSEAGATE_ST373207LC_3KT4PMAN.* ", "a/emcpower.* ", "r/sd.* ", "r/disk.* "]</pre>
Oracle Linux 7.x	<pre>global_filter = ["a /dev./*/by-id/scsi-SSEAGATE_ST373207LC_3KT4PMAN.* ", "a/emcpower.* ", "r/sd.* ", "r/disk.* "]</pre>

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:

Procedure

1. Verify that the Fujitsu drivers `shfchp` and `pciehp` 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:

Procedure

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®. 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.

Before you begin

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:

Procedure

- 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.

Results

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

Before you begin

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.

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.

Procedure

1. Enable `report_scsi_name` by entering the following NetApp command:

```
igroup set igroup_name report_scsi_name yes
```

2. 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:

```
system1> 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

Procedure

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.

Before you begin

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

Storage device	Support steps:
On RHEL 6.x	<ol style="list-style-type: none"> 1. Create the <code>lpfc.conf</code> file in the <code>/etc/modprobe.d/</code> directory. 2. Edit the file and set the <code>lpfc_max_luns</code> parameter to 512.

Storage device	Support steps:
	<ol style="list-style-type: none"> 3. Recreate the <code>initramfs</code> by running the <code>dracut -f</code> command. 4. Reboot the host. <p>For more information, refer to the EMC Host Connectivity Guide for Linux or the vendor website.</p>
On RHEL 7.x	<ol style="list-style-type: none"> 1. Edit the <code>/etc/modprobe.d/lpfc.conf</code> by adding the text: <code>options scsi_mod max_luns=16383 max_report_luns=16384 options lpfc lpfc_max_luns=16384</code> 2. Recreate the <code>initrd</code>. 3. Point the <code>initrd</code> to the new <code>initrd</code> when booting.
SLES 12.x	<ol style="list-style-type: none"> 1. Edit the <code>/etc/modprobe.d/lpfc.conf</code> by adding the text: <code>options scsi_mod max_luns=16383 max_report_luns=16384 options lpfc lpfc_max_luns=16384</code> 2. Recreate the <code>initrd</code>. 3. Point the <code>initrd</code> to the new <code>initrd</code> when booting.
OL 6.x	<ol style="list-style-type: none"> 1. Create the <code>lpfc.conf</code> file in the <code>/etc/modprobe.d/</code> directory. 2. Edit the file and set the <code>lpfc_max_luns</code> parameter to 512. 3. Recreate the <code>initramfs</code> by running the <code>dracut -f</code> command. 4. Reboot the host.
OL 7.x	<ol style="list-style-type: none"> 1. Edit the <code>/etc/modprobe.d/lpfc.conf</code> by adding the text: <code>options scsi_mod max_luns=16383 max_report_luns=16384 options lpfc_max_luns=16384</code> 2. Recreate the <code>initrd</code>. 3. Point the <code>initrd</code> to the new <code>initrd</code> when booting.

/etc/fstab options

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
```

RHEL7, OL7 and SLES 12 SP3

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](#).....90
- [Upgrading PowerPath](#).....90

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 that is 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 file system 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 that are 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 that are 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 to [Changing the language of installation messages](#) on page 76 for information about using a localized version of the PowerPath installer when upgrading PowerPath.
- Refer to [Troubleshooting PowerPath](#) on page 101 if you encounter any problems while upgrading PowerPath.

Note

Do not use any custom scripts to install, upgrade, or uninstall the PowerPath package.

Upgrading PowerPath

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

[Upgrade PowerPath from a CD](#) on page 91

[Upgrade PowerPath from a compressed archive](#) on page 92

[Upgrade PowerPath across Linux platforms](#) on page 93

[Upgrade PowerPath before upgrading RHEL update on Linux on page 95](#)

Upgrade PowerPath from a CD

Upgrade to the latest version of PowerPath using the installation CD.

Procedure

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

```
mkdir /cdrom
```

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

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

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

- For RHEL

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

- For SLES

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

5. Log in as root.

6. Upgrade to the new PowerPath version.

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

Where:

- *Version* is the product version. For example, 6.3.0.00.00.
- *Build* is the software build number.
- *Platform* is the Linux distribution, for example SLES12SP3 or RHEL6.

7. Start PowerPath.

- For RHEL 6.x, OL 6.x or SLES 11 SP4

```
/etc/init.d/PowerPath start
```

- For RHEL 7.x, OL 7.x or SLES 12 SP3

```
systemctl start PowerPath.service
```

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 configuration files updated by the new PowerPath version.
- d. Install the old PowerPath version.
- e. Copy the backed up contents from step a to `/etc/emc/archive` folder.

f. Start PowerPath.

While starting, PowerPath restores the config files from `/etc/emc/archive` and the old settings are restored.

8. If the PowerPath configuration includes any custom policies or settings, run `powermt load` to load these settings.

Upgrade PowerPath from a compressed archive

As a root user, upgrade PowerPath with the PowerPath archive.

Procedure

1. Log in as root.

2. Download the PowerPath archive from [EMC Online Support](#).

From [EMC Online Support](#), click **Support by Product** and then search for PowerPath Linux downloads.

3. Untar the PowerPath archive.

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

4. Upgrade to the new PowerPath version.

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

Where:

- *Version* is the product version. For example, 6.3.0.00.00.
- *Build* is the software build number.
- *Platform* is the Linux distribution, for example SLES12SP3 or RHEL6.

5. Start PowerPath.

- For RHEL 6.x, OL 6.x or SLES 11 SP4

```
/etc/init.d/PowerPath start
```

- For RHEL 7.x, OL 7.x or SLES 12 SP3

```
systemctl start PowerPath.service
```

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 configuration files updated by the new PowerPath version.
- d. Install the old PowerPath version.
- e. Copy the backed up contents from step a to `/etc/emc/archive` folder.
- f. Start PowerPath.

While starting, PowerPath restores the config files from `/etc/emc/archive` and the old settings are restored.

Upgrade PowerPath across Linux platforms

Before you begin

Note

If you want to upgrade PowerPath in a boot from SAN set up or upgrade the Linux kernel, use the information listed in [Administering in a Boot from SAN environment](#) on page 19.

Refer to the following procedures to upgrade PowerPath before or after upgrading the Linux platform:

Procedure

- [Upgrade the Linux kernel in a PowerPath installed host](#) on page 93
- [Upgrade PowerPath after upgrading from one major Linux release to the next](#) on page 93
- [Upgrade PowerPath before upgrading to the latest SLES Service Pack](#) on page 94
- [Upgrade PowerPath before upgrading RHEL update on Linux](#) on page 95
- [Upgrade PowerPath in a SteelEye LifeKeeper cluster environment](#) on page 95

Upgrade the Linux kernel in a PowerPath installed host

To upgrade the Linux kernel in a PowerPath installed setup:

Procedure

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.

Upgrade PowerPath after upgrading from one major Linux release to the next

Upgrade PowerPath after upgrading from major Linux release to next major Linux release, for example, from RHEL 6.5 to RHEL 7.0.

Procedure

1. Perform one of the following steps:
 - For RHEL6.x, OL6.x or SLES11 SP4, rename `/etc/init.d/PowerPath` to `/etc/init.d/PowerPath.old`
 - For RHEL 7.x, OL 7.x or SLES 12 SP3


```
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 ensures that the root file system is mounted as `ext4` on the next reboot.

3. Upgrade to the new PowerPath version.

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

Where:

- *Version* is the product version. For example, 6.3.0.00.00.
- *Build* is the software build number.
- *Platform* is the Linux distribution, for example SLES12SP3 or RHEL6.

4. Start PowerPath.

- For RHEL 6.x, OL 6.x or SLES 11 SP4

```
/etc/init.d/PowerPath start
```

- For RHEL 7.x, OL 7.x or SLES 12 SP3

```
systemctl start PowerPath.service
```

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:

- Copy contents of `/etc/emc/archive` to another folder such as `/tmp/pp_config_backup`.
- Uninstall the new version of PowerPath.
- Delete contents of `/etc/emc/archive` folder since that now contains configuration files updated by the new PowerPath version.
- Install the old PowerPath version.
- Copy the backed up contents from step a to `/etc/emc/archive` folder.
- Start PowerPath.

While starting, PowerPath restores the config files from `/etc/emc/archive` and the old settings are restored.

5. After you have checked to make sure that the upgrade is complete, delete `/etc/init.d/PowerPath.old`.

Upgrade PowerPath before upgrading to the latest SLES Service Pack

To upgrade PowerPath before upgrading Linux SLES on a host:

Before you begin

Procedure

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.

Upgrade PowerPath before upgrading RHEL update on Linux

To upgrade PowerPath before upgrading RHEL on a host:

Procedure

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, 6.3.0.00.00.
- **<build>** is the software build number.
- **<platform>** is the Linux distribution, for example SLES12SP3 or RHEL6.

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

Upgrade 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.

Before you begin**Note**

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

Procedure

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, 6.3.0.00.00.
 - `<build>` is the software build number.
 - `<platform>` is the Linux distribution, for example SLES12SP3 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.

Results

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](#)..... 98
- [Uninstall PowerPath](#)..... 98

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 that are 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.

Note

Do not use any custom scripts to install, upgrade, or uninstall the PowerPath package.

Uninstall PowerPath

Procedure

1. To uninstall a host booting from a local disk:
Uninstall PowerPath using the steps in [Uninstall PowerPath](#) on page 98.
2. To uninstall a host booting from a SAN setup:
Follow the relevant steps given in Chapter 3 [Administering in a Boot from SAN environment](#) on page 19.

Uninstall PowerPath

Procedure

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

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

4. Remove the software.

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

Results

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.

Remove the 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.

Procedure

1. Remove the PowerPath license and all supporting files. :

```
/sbin/emcp_cleanup
```

`emcp_cleanup` issues the following 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, and then manually delete `/sbin/emcp_cleanup`.

CHAPTER 9

Troubleshooting PowerPath

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

• Recover from a failed PowerPath installation	102
• Resolve missing or deleted files	103
• The /etc/fstab	104
• Problems booting the host	104
• Boot issue on RHEL6.x, SLES11 SP4, RHEL7.x, SLES 12 SP3, OL6.x and OL7.x boot-from-SAN when emcpbfs_linux is used	104

Recover from a failed PowerPath installation

If PowerPath installation fails, correct error conditions reported by the installation program. Install PowerPath following the instructions in [Install the RPM file](#) on page 14.

Procedure

1. Determine if the PowerPath modules are loaded.

```
lsmod | grep emc
```

If the modules are loaded, type `/etc/init.d/PowerPath stop`, or for RHEL7, OL7 or SLES 12 SP3 type `systemctl stop PowerPath` to unload them.

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

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

3. Check the `/etc/opt/emcpower/log` file for more 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
```

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

Troubleshoot 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 the custom PowerPath settings are lost during the upgrade, check the following for the possible source of the error.

Procedure

1. Check if the configuration files are corrupt.

If the PowerPath configuration files are corrupt, the custom configuration is no longer available after you upgrade. If an application was configured to use PowerPath pseudo devices, you must reconfigure that application to use the PowerPath pseudo devices after upgrading (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 the PowerPath configuration.

2. Check if the PowerPath configuration that is recorded in the `powermt_custom.xml` file is outdated and does not match the current PowerPath configuration.

If you change the PowerPath configuration but do not run `powermt save`, and then upgrade PowerPath, the custom configuration that is 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 the PowerPath configuration.

3. Check if the file system where the `powermt_custom.xml` file is stored or the file system where PowerPath is being upgraded, is full.

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

4. Check if connectivity problems on the SAN causes the host not to see some devices.

Resolve 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.

Procedure

1. Stop PowerPath.
 - For RHEL 6.x, OL 6.x or SLES 11 SP4
`/etc/init.d/PowerPath stop`
 - RHEL7, OL7 or SLES 12 SP3,
`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.

2. Remove PowerPath.

```
rpm -qa | grep EMCpower.LINUX

rpm -ev EMCpower.LINUX-<version>-<build>
```

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

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

CAUTION

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

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

4. Reinstall PowerPath.

Refer to [Install the RPM file](#) on page 14.

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 start correctly.

Procedure

- 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 SP3, 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 SP3 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 SP3 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 SP3 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.

CHAPTER 10

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](#)..... 108
- [Files created by PowerPath installation](#)..... 108

Linux files modified by PowerPath installation

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

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/

RHEL6, OVM 3.x.x, and OL6

- `PowerPath`
- `PowerPathPost`

SLES 11 SP4

- `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 SP4

- S04powermig.sles

/etc/rc.d/init.d

RHEL6, OVM 3.x.x, OL6

- PowerPath
- PowerPathPost
- emcp_mond.sh

/etc

RHEL6, OVM 3.x.x, 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 SP3, RHEL 7, RHEL 6, 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-6.3.x.00.00

- .pp_build_version
- .pp_vendor
- .pp_version
- EULA.pdf
- PowerPathPost
- PowerPath
- PowerPathPost-start.service (SLES 12 SP3, RHEL7 and OL7)
- PowerPathPost-stop.service (SLES 12 SP3, RHEL7 and OL7)
- PowerPath.service (SLES 12 SP3, 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-6.3.x.00.00/bin

- .cmds
- .drivers_base
- .drivers_ext
- .drivers_mgr
- .sharedlibs
- .sharedlibs32
- .staticlibs

/etc/opt/emcpower/EMCpower.LINUX-6.3.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-6.3.x.00.00/bin/driver

SLES 12 SP3

- sles12sp3default_x8664
- sles12sp3default_x8664/emcp.ko
- sles12sp3default_x8664/emcpdm.ko
- sles12sp3default_x8664/emcpgpx.ko
- sles12sp3default_x8664/emcpioc.ko
- sles12sp3default_x8664/emcpmpx.ko

SLES 11 SP4

- sles11sp4default_x8664/emcpdm.ko
- sles11sp4default_x8664/emcpgpx.ko
- sles11sp4default_x8664/emcpioc.ko
- sles11sp4default_x8664/emcp.ko
- sles11sp4default_x8664/emcpmpx.ko
- sles11sp4xen_x8664/emcpdm.ko
- sles11sp4xen_x8664/emcpgpx.ko
- sles11sp4xen_x8664/emcpioc.ko
- sles11sp4xen_x8664/emcp.ko
- sles11sp4xen_x8664/emcpmpx.ko

RHEL7

- rhel7_x8664/emcp.ko
- rhel7_x8664/emcpdm.ko
- rhel7_x8664/emcpgpx.ko

- rhel7_x8664/emcpioc.ko
- rhel7_x8664/emcpmpx.ko

RHEL6

- rhel6_x8664/emcp.ko
- rhel6_x8664/emcpdm.ko
- rhel6_x8664/emcpgpx.ko
- rhel6_x8664/emcpioc.ko
- rhel6_x8664/emcpmpx.ko

OL6

- ol6_uek3_x8664/emcp.ko
- ol6_uek3_x8664/emcpdm.ko
- ol6_uek3_x8664/emcpgpx.ko
- ol6_uek3_x8664/emcpioc.ko
- ol6_uek3_x8664/emcpmpx.ko
- ol6_uek4_r2_x8664/emcp.ko
- ol6_uek4_r2_x8664/emcpdm.ko
- ol6_uek4_r2_x8664/emcpgpx.ko
- ol6_uek4_r2_x8664/emcpioc.ko
- ol6_uek4_r2_x8664/emcpmpx.ko

OL7

- 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/emcpmpx.ko
- ol7_uek4_r2_x8664/emcp.ko
- ol7_uek4_r2_x8664/emcpdm.ko
- ol7_uek4_r2_x8664/emcpgpx.ko
- ol7_uek4_r2_x8664/emcpioc.ko
- ol7_uek4_r2_x8664/emcpmpx.ko

/etc/opt/emcpower/EMCpower.LINUX-6.3.x.00.00/bin/lib

- emcp_mond_netsnmp.so
- libemcp.so
- libemcp_serv_rtl.so
- libemcp_core.so
- libemcp_lam.so
- libemcp_lic_rtl.so
- libemcp_mp_rtl.so
- libemcp_mpapi_rtl.so

- libemcphp.so
- libpn.so

/etc/opt/emcpower/EMCpower.LINUX-6.3.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
- libemcphp.so
- libpn.so

/etc/opt/emcpower/EMCpower.LINUX-6.3.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.6.3.x.00.00/man

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

/etc/udev/rules.d

- 05-udev-pp.rules
- 70-persistent-ipoib.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
- libemcpmp.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
- emcpmpx.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 SP3 and OL7

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

/usr/lib/PowerPath

RHEL7, SLES 12 SP3 and OL7

- PowerPath
- PowerPathPost
- emcp_mond.sh

/usr/share/man/man1

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

SLES 12 SP3, 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-6.3.0.00.00/90powerpath
- 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-6.3.x.00.00/90powerpath
- check
- emcp_pre-trigger.sh
- emcp_settled.sh
- install
- 64-emcp.rules
- module-setup.sh

Files changed by PowerPath