

# Dell EMC PowerPath for Linux

Version 6.5

## Installation and Administration Guide

302-005-499

01

January 2019

Copyright © 1998-2019 Dell Inc. or its subsidiaries. All rights reserved.

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 Technologies, Dell, EMC, 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.

Dell EMC  
Hopkinton, Massachusetts 01748-9103  
1-508-435-1000 In North America 1-866-464-7381  
[www.DellEMC.com](http://www.DellEMC.com)

# CONTENTS

<b>Tables</b>		<b>7</b>
<b>Chapter 1</b>	<b>Getting started</b>	<b>9</b>
	Getting started.....	10
<b>Chapter 2</b>	<b>Installing PowerPath</b>	<b>11</b>
	Before you install.....	12
	Storage System related steps.....	12
	Reviewing compatibility with the environment.....	12
	Review host system compatibility.....	12
	Review PowerPath licensing.....	13
	Review third-party path management software.....	13
	Review the installation language.....	14
	Install the RPM file.....	14
	Install PowerPath when VSF is already installed.....	14
	License after the installation.....	15
	Post-installation steps.....	16
<b>Chapter 3</b>	<b>Administering in a Boot from SAN environment</b>	<b>19</b>
	Before configuring for a root device.....	20
	Configuring PowerPath in an RHEL 6.x or OL 6.x Boot from SAN setup in a BIOS host.....	20
	Configure PowerPath in an RHEL 7.x, OL 7.x, SLES 12 SP4, or SLES 15 Boot from SAN setup in a BIOS host.....	22
	Upgrading PowerPath in a RHEL 6.x or OL 6.x Boot from SAN setup in a BIOS host.....	25
	Upgrade using the <code>emcpbfs_linux</code> script .....	25
	Upgrade using manual steps.....	25
	Upgrade PowerPath in an RHEL 7.x, OL 7.x, SLES 12 SP4, or SLES 15 Boot from SAN setup in a BIOS host.....	27
	Upgrade using the <code>emcpbfs_linux</code> script.....	27
	Upgrade using manual steps.....	28
	Upgrade the Linux kernel in a Boot from SAN setup in a BIOS/UEFI host.....	30
	Upgrade using <code>emcpbfs_linux</code> script .....	30
	Upgrade using manual steps.....	31
	Uninstalling PowerPath in a RHEL 6.x or OL 6.x Boot from SAN setup in a BIOS host .....	34
	Uninstall using <code>emcpbfs_linux_script</code> .....	34
	Uninstall using manual steps.....	34
	Uninstall PowerPath in an RHEL 7.x, OL 7.x, SLES 12 SP4, or SLES 15 Boot from SAN setup in a BIOS host.....	35
	Uninstall using <code>emcpbfs_linux_script</code> .....	35
	Uninstall using manual steps.....	36
	Configure PowerPath in an RHEL 6.x or OL 6.x Boot from SAN setup in an UEFI host.....	38
	Configure PowerPath in an RHEL 7.x, OL 7.x, SLES 12 SP4, or SLES 15 Boot from SAN setup in an UEFI host.....	40

	Upgrading PowerPath in a RHEL 6.x or OL 6.x Boot from SAN setup in an UEFI host.....	43
	Upgrade using the <code>emcpbfs_linux</code> script .....	43
	Upgrade using manual steps .....	44
	Upgrade PowerPath in an RHEL 7.x, OL 7.x, SLES 12 SP4, or SLES 15 Boot from SAN setup in an UEFI host.....	45
	Upgrade using the <code>emcpbfs_linux</code> script .....	45
	Upgrade using manual steps .....	46
	Uninstalling PowerPath in a RHEL 6.x or OL 6.x Boot from SAN setup in an UEFI host .....	48
	Uninstall using <code>emcpbfs_linux_script</code> .....	48
	Uninstall using manual steps.....	48
	Uninstall PowerPath in an RHEL 7.x, OL 7.x, SLES 12 SP4, or SLES 15 Boot from SAN setup in an UEFI host.....	49
	Uninstall using <code>emcpbfs_linux_script</code> .....	49
	Uninstall using manual steps.....	50
<b>Chapter 4</b>	<b>Managing paths to a logical device</b>	<b>53</b>
	Dynamic LUN addition and removal.....	54
	Adding a LUN.....	54
	Removing a LUN.....	55
	Resizing a LUN.....	55
	Before you add or remove paths.....	55
	Adding new paths to a PowerPath logical device.....	56
	Removing some of the paths to PowerPath logical devices.....	57
<b>Chapter 5</b>	<b>PowerPath features for PowerMax array</b>	<b>61</b>
	PowerPath Host Registration.....	62
	Overview.....	62
	Initial Registration and Storage Configuration.....	62
	Modes of operation.....	62
	Settings.....	63
	Log messages.....	63
	FX Licensing.....	64
	PowerPath Device IO Statistics.....	64
	Overview.....	64
	Settings.....	65
	PowerPath Device in Use Reporting.....	65
	Overview.....	65
	Settings.....	65
<b>Chapter 6</b>	<b>Coexistence with other path management software</b>	<b>67</b>
	PowerPath and third-party path management software.....	68
	PowerPath co-management behavior.....	68
<b>Chapter 7</b>	<b>Maintaining PowerPath</b>	<b>71</b>
	Configure remote management service.....	72
	Configure event monitoring with SNMP traps.....	73
	Monitored events.....	73
	Configure SNMP traps.....	74
	Changing the language of installation messages.....	77
	Saving new configurations.....	78
	Configuring LVM2 support.....	78

	Root file system not mounted on a logical volume.....	79
	Root file system mounted on a logical volume.....	80
	Root file system mounted on a logical volume residing on a local disk not managed by PowerPath.....	80
	<b>Configuring PowerPath on cluster nodes .....</b>	<b>82</b>
	Hot swapping an HBA.....	83
	Removing an HBA.....	83
	Adding an HBA.....	84
	Configuring for R1/R2 boot failover support.....	84
	R1/R2 supported configurations.....	84
	R1/R2 boot procedure.....	85
	Configure the network interface on the R1 host.....	86
	SCSI-2 reservation support in PowerPath for Linux .....	86
	Enable scsi name on NetApp arrays.....	86
	Removing devices from PowerPath control.....	87
	Expand the file system on a pseudo-device.....	87
	Supporting more than 256 storage devices.....	87
	/etc/fstab options.....	88
<b>Chapter 8</b>	<b>Upgrading PowerPath .....</b>	<b>89</b>
	Preparing for upgrade.....	90
	Upgrading PowerPath.....	90
	Upgrade PowerPath from a compressed archive.....	90
	Upgrade PowerPath across Linux platforms.....	92
	Upgrade the Linux kernel in a PowerPath installed host.....	92
	Upgrade PowerPath after upgrading from one major Linux release to the next .....	92
	Upgrade PowerPath before upgrading to the latest SLES Service Pack..	93
	Upgrade PowerPath before upgrading RHEL update on Linux.....	94
	Upgrade PowerPath in a SteelEye LifeKeeper cluster environment.....	94
<b>Chapter 9</b>	<b>Removing PowerPath .....</b>	<b>97</b>
	Before you remove PowerPath.....	98
	Uninstall PowerPath.....	98
	Uninstall PowerPath.....	98
	Saved configuration files.....	99
	Remove the PowerPath files.....	99
<b>Chapter 10</b>	<b>Troubleshooting PowerPath .....</b>	<b>101</b>
	Recover from a failed PowerPath installation.....	102
	Troubleshoot a PowerPath upgrade.....	103
	PowerPath custom settings not preserved.....	103
	Resolve missing or deleted files.....	103
	The /etc/fstab.....	104
	Problems booting the host.....	104
	Boot issue on RHEL6.x, RHEL7.x, SLES 12 SP4, SLES 15, 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.....	105
	/etc/opt/emcpower/emcpbfs_linux upgrade.....	105

<b>Chapter 11</b>	<b>Files changed by PowerPath</b>	<b>107</b>
	Linux files modified by PowerPath installation.....	108
	Files created by PowerPath installation.....	108
<b>Appendix A</b>	<b>Retpoline Kernels</b>	<b>117</b>
	PowerPath for Linux 6.5 and retpoline.....	118
	PowerPath for Linux 6.5 and non-retpoline (older) kernels.....	118
<b>Appendix B</b>	<b>Dynamic Linking of OpenSSL Libraries</b>	<b>121</b>
	Overview.....	122
	Changes.....	122
	Expected behavior.....	122
	With OpenSSL 1.0.x.....	122
	With no OpenSSL library or incompatible version.....	122
	PowerPath for Linux and SLES 15.....	123
<b>Appendix C</b>	<b>Disable native multipathing</b>	<b>125</b>
	Disabling dm-multipath on an RHEL 6.x or OL 6.x host.....	126
	Disabling dm-multipath on an RHEL 7.x or OL 7.x host.....	126
<b>Appendix D</b>	<b>Booting with duplicate PVs on SLES BFS hosts</b>	<b>129</b>
	SLES hosts and LVM.....	130
	SLES BFS hosts and duplicate PVs.....	130
	Configure LVM to accept duplicate PVs.....	131
	Troubleshooting.....	131
<b>Index</b>		<b>133</b>

# TABLES

1	Getting started tasks.....	10
2	Software with disabled PowerPath support.....	68
3	Present or enabled software and multipathing behavior.....	68
4	Commands to disable support.....	69
5	Monitored events.....	73
6	Supported languages.....	78
7	Linux Kernel Filter Fields.....	79
8	Kernel Filter.....	81
9	Operating System Support Steps.....	88





# 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?	<a href="#">Installing PowerPath</a> on page 11 for quick installation with a compressed archive from <a href="#">Dell EMC Online Support</a> .
administer PowerPath?	<a href="#">Administering in a Boot from SAN environment</a> 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?	<a href="#">Upgrading PowerPath</a> on page 89 for pre-requisites and procedures to upgrade PowerPath.
troubleshoot PowerPath?	<a href="#">Troubleshooting PowerPath</a> 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 [Dell 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.

- [Before you install](#)..... 12
- [Install the RPM file](#).....14
- [Install PowerPath when VSF is already installed](#).....14
- [License after the installation](#)..... 15
- [Post-installation steps](#)..... 16

## Before you install

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

 **Note:** Ensure that the storage system logical devices support PowerPath. The *Host Connectivity Guide* and the storage system support website provides additional information.

## Storage System related steps

Plan your installation based on your storage system:

### Procedure

- DELL SC: Ensure that DELL SC storage array is working properly before PowerPath installation.
- CLARiiON® storage systems: Contact CLARiiON support to plan your installation. PowerPath installation is a part of a CLARiiON setup and configuration procedure.
- Symmetrix® storage systems: Ensure that Symmetrix is working properly before PowerPath installation.
- VNX™ OE storage systems: Contact VNX support to plan your installation. PowerPath installation is a part of a VNX OE setup and configuration procedure.
- VPLEX Storage Systems: Ensure that VPLEX storage array is working properly before PowerPath installation.
- XtremIO Storage Systems: Ensure that XtremIO storage array is working properly before PowerPath installation.
- Third-party storage systems: Ensure that the third-party array is working properly before installing PowerPath.

## Reviewing compatibility with the environment

### Procedure

1. Review the Linux kernel updates as per the Dell 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/eln/elhome>.
4. Use EMC Grab utilities, Serviceability feature of PowerPath Management Appliance, and EMC Reporting tool and verify that your environment meets the requirements.
5. Ensure that HBA BIOS and fibre drivers are configured as per the <https://elabnavigator.emc.com/eln/elhome>. 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. For Boot from SAN configurations, ensure that the file system layout of the Operating System installation is according to the layout specified in [Administering in a Boot from SAN environment](#) on page 19.
4. PowerPath for Linux 6.5 works only with retpoline enabled kernels. See [Retpoline Kernels](#) on page 117.
5. Ensure any version of EMC Unisphere® and Navisphere® Application Transparent Failover (ATF) is removed from the host by EMC Professional Services.
6. 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 an RHEL 6.x or OL 6.x host](#) on page 126 or [Disabling dm-multipath on an RHEL 7.x or OL 7.x host](#) on page 126 for instructions on how to disable native multipathing.

When installing on SLES 12 SP4 or SLES 15, ensure that native multipathing is not automatically activated.

7. Configure LVM,
 

See the section [Configuring LVM2 support](#) on page 78.
8. 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.
9. PowerPath needs a compatible version of OpenSSL library to be installed on the host if the remote management feature is enabled. See [Dynamic Linking of OpenSSL Libraries](#) on page 121.

## Review PowerPath licensing

### About this task

There are two ways to license PowerPath:

### Procedure

1. FX array bundle:

When the host is attached to a PowerMax array with an FX bundle, then PowerPath is automatically licensed. See [FX Licensing](#) on page 64.

2. License key:

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

**Note:** In either case, 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

### About this task

PowerPath is not supported with Linux DM-MPIO. If DM-MPIO manages any array devices, PowerPath does not start.

In case, there is a storage vendor specific path management software for a third-party storage array, it can coexist with PowerPath if the following conditions are met:

1. Ensure that the third party storage system class type is in unmanaged state. 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. The same HBAs cannot be used for both PowerPath and the third party path management software.

## Review the installation language

### About this task

To review the procedure for changing the installation language, see [Changing the language of installation messages](#) on page 77.

## Install the RPM file

From the [Dell 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.

### Procedure

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

```
[root@hostname ~]# tar -xzf DelleMCPower.LINUX-<version>.<build>.tar.gz
```

Where:

- *<version>* is the product version. For example, 6.5.0.00.00.
  - *<build>* is the software build number.
4. Copy the RPM package into a temporary folder.
  5. Install PowerPath.

```
[root@hostname ~]# rpm -ivh DelleMCPower.LINUX-  
<version>.<build>.<platform>.x86_64.rpm
```

Where *<platform>* is the Linux distribution, for example SLES12SP4 or RHEL6.

## Install PowerPath when VSF is already installed

Installing PowerPath after Veritas Storage Foundation (VSF) requires reconfiguring Veritas Volume Manager to recognize the pseudo devices.

### About this task

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 SC devices,  
`vxddladm excludearray libname=libvxSc.so`
  - For For Unity devices,  
`vxddladm excludearray libname=libvxUnity.so`
  - 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=libvxVplex.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.

```
[root@hostname ~]# rpm -ivh DelleMCPower.LINUX-<version>-<build>.<platform>.x86_64.rpm
```

6. Ensure that PowerPath claims the devices.

```
[root@hostname ~]# vxctl enable
```

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

7. Run the command:

```
[root@hostname ~]# vxdisk scandisks
```

## 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.

```
[root@hostname ~]# emcprep -install
```

**Note:** This step is needed only if licensing is done with a license key and not required if an FX based license is used.

- 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 on-screen instructions.

## 2. Start the PowerPath service.

Use the script only to load and unload PowerPath modules.

- For RHEL 6 or OL6

```
[root@hostname ~]# /etc/init.d/PowerPath start
```

- For RHEL7, OL7, SLES 12 SP4, or SLES 15

```
[root@hostname ~]# systemctl start PowerPath.service
```

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

```
[root@hostname ~]# powermt set policy={ad|co|lb|li|re|rr|so|si}
[dev=<device>|all] [class=<class>|all]
```

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

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

```
[root@hostname ~]# 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. Configure the remote management service if the host needs to be connected to a PowerPath Management Appliance.

For steps on configuring, see [Configure remote management service](#) on page 72.

2. Configure the management daemon to send SNMP traps when path events occur, if necessary.



For steps on configuring, see [Configure SNMP traps](#) on page 74.

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

```
[root@hostname ~]# 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 an RHEL 6.x or OL 6.x Boot from SAN setup in a BIOS host](#)..... 20
- [Configure PowerPath in an RHEL 7.x, OL 7.x, SLES 12 SP4, or SLES 15 Boot from SAN setup in a BIOS host](#).....22
- [Upgrading PowerPath in a RHEL 6.x or OL 6.x Boot from SAN setup in a BIOS host](#)..... 25
- [Upgrade PowerPath in an RHEL 7.x, OL 7.x, SLES 12 SP4, or SLES 15 Boot from SAN setup in a BIOS host](#).....27
- [Upgrade the Linux kernel in a Boot from SAN setup in a BIOS/UEFI host](#)..... 30
- [Uninstalling PowerPath in a RHEL 6.x or OL 6.x Boot from SAN setup in a BIOS host](#) ..... 34
- [Uninstall PowerPath in an RHEL 7.x, OL 7.x, SLES 12 SP4, or SLES 15 Boot from SAN setup in a BIOS host](#)..... 35
- [Configure PowerPath in an RHEL 6.x or OL 6.x Boot from SAN setup in an UEFI host](#)..... 38
- [Configure PowerPath in an RHEL 7.x, OL 7.x, SLES 12 SP4, or SLES 15 Boot from SAN setup in an UEFI host](#)..... 40
- [Upgrading PowerPath in a RHEL 6.x or OL 6.x Boot from SAN setup in an UEFI host](#).....43
- [Upgrade PowerPath in an RHEL 7.x, OL 7.x, SLES 12 SP4, or SLES 15 Boot from SAN setup in an UEFI host](#)..... 45
- [Uninstalling PowerPath in a RHEL 6.x or OL 6.x Boot from SAN setup in an UEFI host](#) .....48
- [Uninstall PowerPath in an RHEL 7.x, OL 7.x, SLES 12 SP4, or SLES 15 Boot from SAN setup in an UEFI host](#)..... 49

## Before configuring for a root device

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 78 provides information on modifying the `/etc/lvm/lvm.conf` file.


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

## Configuring PowerPath in an 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 an 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 that 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 to make 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 to prevent the 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 operating system installation is complete.

4. Configure the boot LUN in the HBA BIOS by following the procedure in HBA vendor documentation.
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:

 **Note:** Ensure that the partition layout type is MBR and not GPT.

- The `/boot` partition that is formatted as an `ext3`, `ext4`, or `xfs` 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`, `ext4`, or `xfs` and the other as `swap` and designate the mount points for each volume.

9. Ensure that GRUB is the boot loader is selected and that the boot loader location specifies the MBR of the wanted 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, see [Disabling dm-multipath on an RHEL 6.x or OL 6.x host](#) on page 126.
12. If the operating system 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.
14. To complete the configuration of the boot LUN under PowerPath, using one of the following methods:

- **Using the `emcpbfs_linux` script:**

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

Run the following command to configure PowerPath:

```
[root@hostname~] # /etc/opt/emcpower/emcpbfs_linux config
```

- i **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 `/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
```

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

```
[root@hostname ~]# umount /boot/
[root@hostname ~]# 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 a new ramdisk image to include PowerPath modules.

```
[root@hostname ~]# dracut /boot/initramfs-PP-$(uname -r).img $(uname -r)
```

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

```
Red Hat Enterprise Linux (2.6.32-754.el6.x86_64)
root (hd0,0)
kernel /vmlinuz-2.6.32-754.el6.x86_64 ro root=/dev/mapper/vg_lib1044-
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_lib1044/lv_swap rd_LVM LV=vg_lib1044/lv_root rd_NO_DM
initrd /initramfs-PP-2.6.32-754.el6.x86_64.img

Red Hat Enterprise Linux (2.6.32-754.el6.x86_64)
root (hd0,0)
kernel /vmlinuz-2.6.32-754.el6.x86_64 ro root=/dev/mapper/vg_lib1044-
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_lib1044/lv_swap rd_LVM LV=vg_lib1044/lv_root rd_NO_DM
initrd /initramfs-2.6.32-754.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, and 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 run the following steps manually:
- Take a backup of the existing ramdisk image.
- Update the ramdisk image by following step f in the manual steps detailed in the previous procedure.

## Configure PowerPath in an RHEL 7.x, OL 7.x, SLES 12 SP4, or SLES 15 Boot from SAN setup in a BIOS host

To configure a PowerPath root device using the LVM on an RHEL 7.x, OL 7.x, SLES 12 SP4, or SLES 15 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 SP4 or SLES 15, ensure that native multipathing is not automatically activated. Because, there is no way to disable native multipathing after the operating system installation is complete.
- During configuration, it is advisable to have only the boot LUN connected to the host to make it easier to identify the boot disk during installation. It is also advisable to have only one path from the array to the host when installing the operating system. It prevents the operating system 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 operating system installation is complete.
- To review the procedure for disabling dm-multipath on an RHEL 7.x or OL 7.x, see [Disabling dm-multipath on an RHEL 7.x or OL 7.x host](#) on page 126.

4. At installation, select **Basic Storage Devices** as the type of device for installation.

**Note:** Ensure that the partition layout type is MBR and not GPT.

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 operating system installation is complete, add additional paths to the boot LUN and attach additional LUNs to the host.
7. Install and configure PowerPath.
8. Complete the configuration of the boot LUN under PowerPath.

- To automatically complete the configuration.

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/emcpower1 /boot ext4 defaults 1 2
```

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

```
[root@hostname ~]# umount /boot/
[root@hostname ~]# mount /boot/
```

- c. Setup the LVM filter in `/etc/lvm/lvm.conf`.


```
#global_filter = [ "a./.*" ]
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 a new ramdisk image to include PowerPath modules:

 **Note:** On RHEL7 and OL7, the ramdisk image is named as `initramfs`. On SLES12 and SLES15, it is named as `initrd`.

On RHEL 7.x/OL 7.x:

```
[root@hostname ~]# dracut /boot/initramfs-PP-$(uname -r).img $(uname -r)
```

On SLES 12 SP4/SLES 15:

```
[root@hostname ~]# dracut /boot/initrd-PP-$(uname -r) $(uname -r)
```

g. Add a boot entry with the new 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-3.10.0-327.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=latacyrheb-sun16
    vconsole.keymap=us rhgb quiet LANG=en_US.UTF-8
    initrd16 /initramfs-PP-3.10.0-327.el7.x86_64.img
```

h. Update the new entry in the `grub.cfg` file and make this entry as default so that system boots with the new ramdisk image automatically.

On RHEL 7.x/OL 7.x:

```
[root@hostname ~]# grub2-mkconfig -o /boot/grub2/grub.cfg

[root@hostname ~]# grub2-set-default 'PowerPath - Red Hat Enterprise
Linux Server'
```



On SLES 12 SP4/SLES 15:

```
[root@hostname ~]# grub2-mkconfig -o /boot/grub2/grub.cfg
[root@hostname ~]# grub2-set-default 'PowerPath - SLES 12-SP4'
```

9. Reboot the host to have the new ramdisk image that has PowerPath modules to take effect.

## 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 25
- [Upgrade using manual steps](#) on page 25

### 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, it indicates that `emcpbfs_linux config` has not been used during PowerPath configuration.

See [Upgrade using manual steps](#) on page 25 to upgrade PowerPath.


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

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

3. Reboot the host.
4. Upgrade and start PowerPath.

 **Note:** To review the procedure for upgrading PowerPath, see [Upgrade PowerPath from a compressed archive](#) on page 90.

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/emcpower1 /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. Upgrade and start PowerPath.

 **Note:** To review the procedure for upgrading PowerPath, see [Upgrade PowerPath from a compressed archive](#) on page 90.

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

```
#UUID=110223d8-f5f3-4058-812f-16d9e9ee7eab /boot ext4 defaults 1 2
/dev/emcpower1 /boot ext4 defaults 1 2
```

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

```
[root@hostname ~]# umount /boot/
[root@hostname ~]# mount /boot/
```

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

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

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

```
[root@hostname ~]# rm -f /boot/initramfs-PP-$(uname -r).img
```

11. Build a new ramdisk image to include upgraded PowerPath modules.

```
[root@hostname ~]# dracut /boot/initramfs-PP-$(uname -r).img $(uname -r)
```

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

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

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

## Upgrade PowerPath in an RHEL 7.x, OL 7.x, SLES 12 SP4, or SLES 15 Boot from SAN setup in a BIOS host

To upgrade PowerPath in a Boot from SAN setup on an RHEL 7.x, OL 7.x, SLES 12 SP4, or SLES 15 host, use one of the following methods:

- [Upgrade using the `emcpbfs\_linux` script](#) on page 27
- [Upgrade using manual steps](#) on page 28

### Upgrade using the `emcpbfs_linux` script

#### About this task

#### CAUTION

On SLES12 SP4 and SLES15, LVM does not activate logical volumes if their physical volume (PV) appears on multiple devices. In case of BFS hosts with the root filesystem on a logical volume, if the host sees multiple PVs on boot, the host does not boot up as the root logical volume and swap volume cannot be activated.

This happens when the host is booted and PowerPath or any other multipathing software is not configured on the host. To enable the host to boot properly, change the `allow_changes_with_duplicate_pvs` parameter in the `devices{}` section in `/etc/lvm/lvm.conf` to 1. If the parameter is not present, add the same. Ensure that these changes are made before upgrading or uninstalling PowerPath, else the host may not come up on reboot.


See [Bootting with duplicate PVs on SLES BFS hosts](#) on page 129 for more information.


#### 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, it indicates that `emcpbfs_linux config` has not been used during PowerPath configuration.

See [Upgrade using manual steps](#) on page 28 to upgrade PowerPath.
2. Ensure that all other entries in the `/etc/fstab` file that refer to PowerPath pseudo (`emcpower`) devices are commented.
 

 **Note:** `emcpbfs_linux` script modifies `'/boot'` entry only in the `/etc/fstab` file.
3. Reboot the host.
4. Upgrade and start PowerPath.
 

 **Note:** To review the procedure for upgrading PowerPath, see [Upgrade PowerPath from a compressed archive](#) on page 90.
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

### About this task

#### CAUTION

On SLES12 SP4 and SLES15, LVM does not activate logical volumes if their physical volume (PV) appears on multiple devices. In case of BFS hosts with the root filesystem on a logical volume, if the host sees multiple PVs on boot, the host does not boot up as the root logical volume and swap volume cannot be activated.

This happens when the host is booted and PowerPath or any other multipathing software is not configured on the host. To enable the host to boot properly, change the `allow_changes_with_duplicate_pvs` parameter in the `devices{}` section in `/etc/lvm/lvm.conf` to 1. If the parameter is not present, add the same. Ensure that these changes are made before upgrading or uninstalling PowerPath, else the host may not come up on reboot.

See [Booting with duplicate PVs on SLES BFS hosts](#) on page 129 for more information.

### Procedure

1. Disable the following PowerPath services from starting on host reboot.

```
[root@hostname ~]# systemctl disable PowerPath.service
[root@hostname ~]# systemctl disable PowerPathPost-start.service
[root@hostname ~]# 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/emcpower1 /boot ext4 defaults 1 2
```

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

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

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

On RHEL 7.x/OL 7.x:

```
[root@hostname ~]# grub2-set-default 'Red Hat Enterprise Linux Server'
```

On SLES 12 SP4/SLES 15:

```
[root@hostname ~]# grub2-set-default 'SLES 12-SP4'
```

5. Reboot the host.

## 6. Upgrade and start PowerPath.

**Note:** To review the procedure for upgrading PowerPath, see [Upgrade PowerPath from a compressed archive](#) on page 90.

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

```
#UUID=110223d8-f5f3-4058-812f-16d9e9ee7eab /boot ext4 defaults 1 2
/dev/emcpower1 /boot ext4 defaults 1 2
```

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

```
[root@hostname ~]# umount /boot/
[root@hostname ~]# mount /boot/
```

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

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

10. Delete the old ramdisk image containing PowerPath modules that was created in [Configure PowerPath in an RHEL 7.x, OL 7.x, SLES 12 SP4, or SLES 15 Boot from SAN setup in a BIOS host](#) on page 22.

```
[root@hostname ~]# rm -f /boot/initramfs-PP-$(uname -r).img
```

## 11. Build a new ramdisk image to include upgraded PowerPath modules.

On RHEL 7.x or OL 7.x:

```
[root@hostname ~]# dracut /boot/initramfs-PP-$(uname -r).img $(uname -r)
```

On SLES 12 SP4/SLES 15:

```
[root@hostname ~]# dracut /boot/initrd-PP-$(uname -r) $(uname -r)
```

12. The `/etc/grub.d/40_custom` already contains an entry with ramdisk image created during initial setup, see [Configure PowerPath in an RHEL 7.x, OL 7.x, SLES 12 SP4, or SLES 15 Boot from SAN setup in a BIOS host](#) on page 22.

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

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

# Upgrade the Linux kernel in a Boot from SAN setup in a BIOS/UEFI host

## Before you begin

When upgrading a kernel on a host by installing rpm or using an online option such as Yum, Zypper, or Yast on an RHEL, Oracle Linux, SLES 12 SP4, or SLES 15 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, and then complete one of the following steps to ensure that PowerPath functions normally after the upgrade:

- [Upgrade using emcpbfs\\_linux script](#) on page 30
- [Upgrade using manual steps](#) on page 31

## Upgrade using `emcpbfs_linux` script

### About this task

When upgrading a kernel on a host by installing rpm or using an online option such as Yum, Zypper, or Yast on an RHEL, Oracle Linux, SLES 12 SP4, or SLES 15 host, PowerPath automatically restarts without requiring any extra steps. In such cases, do only the steps that are required to update the kernel as specified in the operating system documentation. Other steps that are mentioned here is not required to be performed.


### CAUTION

On SLES12 SP4 and SLES15, LVM does not activate logical volumes if their physical volume (PV) appears on multiple devices. In case of BFS hosts with the root filesystem on a logical volume, if the host sees multiple PVs on boot, the host does not boot up as the root logical volume and swap volume cannot be activated.

This happens when the host is booted and PowerPath or any other multipathing software is not configured on the host. To enable the host to boot properly, change the `allow_changes_with_duplicate_pvs` parameter in the `devices{}` section in `/etc/lvm/lvm.conf` to 1. If the parameter is not present, add the same. Ensure that these changes are made before upgrading or uninstalling PowerPath, else the host may not come up on reboot.

See [Bootting with duplicate PVs on SLES BFS hosts](#) on page 129 for more information.

### Procedure

1. Run `/etc/opt/emcpower/emcpbfs_linux remove` to prepare the host for kernel upgrade.  
If the `ERROR: PowerPath is not configured on Boot from SAN. message` is displayed, it indicates that `emcpbfs_linux config` has not been used during PowerPath configuration.  
See [Upgrade using manual steps](#) on page 31 to upgrade PowerPath.
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 operating system vendor documentation and reboot the host.
4. Start PowerPath.

- On RHEL 6.x or OL 6.x

```
[root@hostname ~]# /etc/init.d/PowerPath start
```

- On RHEL 7.x, OL 7.x, SLES 12 SP4, or SLES 15

```
[root@hostname ~]# systemctl start PowerPath.service
```

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

## Upgrade using manual steps

### About this task

When upgrading a kernel on a host by installing rpm or using an online option such as Yum, Zypper, or Yast on an RHEL, Oracle Linux, SLES 12 SP4, or SLES 15 host, PowerPath automatically restarts without requiring any extra steps. In such cases, do only the steps required to update the kernel as specified in the operating system documentation. Other steps mentioned here is not required to be performed.

### ⚠ CAUTION

On SLES12 SP4 and SLES15, LVM does not activate logical volumes if their physical volume (PV) appears on multiple devices. In case of BFS hosts with the root filesystem on a logical volume, if the host sees multiple PVs on boot, the host does not boot up as the root logical volume and swap volume cannot be activated.

This happens when the host is booted and PowerPath or any other multipathing software is not configured on the host. To enable the host to boot properly, change the `allow_changes_with_duplicate_pvs` parameter in the `devices{}` section in `/etc/lvm/lvm.conf` to 1. If the parameter is not present, add the same. Ensure that these changes are made before upgrading or uninstalling PowerPath, else the host may not come up on reboot.

See [Booting with duplicate PVs on SLES BFS hosts](#) on page 129 for more information.

### 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/emcpower1 /boot ext4
defaults 1 2
/dev/mapper/vg_lib1048-Swap swap swap
defaults 0 0
tmpfs /dev/shm tmpfs
defaults 0 0
devpts /dev/pts devpts
gid=5,mode=620 0 0
sysfs /sys sysfs
defaults 0 0
proc /proc proc
defaults 0 0
```

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

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

3. Rebuild the LVM2 cache and verify that 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 operating system vendor documentation and then restart the host.

5. Start PowerPath.

- On RHEL 6.x or OL 6.x

```
[root@hostname ~]# /etc/init.d/PowerPath start
```

- On RHEL 7.x, OL 7.x, SLES 12 SP4, or SLES 15

```
[root@hostname ~]# systemctl start PowerPath.service
```

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/emcpower1 /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.

```
[root@hostname ~]# umount /boot/
[root@hostname ~]# 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 a new ramdisk image to include PowerPath modules.

On RHEL 6.x/7.x and OL 6.x/7.x:

```
[root@hostname ~]# dracut /boot/initramfs-PP-$(uname -r).img $(uname -r)
```

On SLES 12 SP4/SLES 15:

```
[root@hostname ~]# dracut /boot/initrd-PP-$(uname -r) $(uname -r)
```

## 10. Add a new boot entry which has the new ramdisk image built in step 9.

- On RHEL6.x/OL6.x, add a new entry with the new ramdisk image name in `/boot/grub/menu.lst` created in the previous step and make it the default entry.

In `/boot/grub/menu.lst` and make it the default entry.

```
Red Hat Enterprise Linux (2.6.32-754.el6.x86_64)
root (hd0,0)
kernel /vmlinuz-2.6.32-754.el6.x86_64 ro root=/dev/mapper/vg_lib1044-
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_lib1044/lv_swap rd_LVM LV=vg_lib1044/lv_root rd_NO_DM
initrd /initramfs-PP-2.6.32-754.el6.x86_64.img
```

```
Red Hat Enterprise Linux (2.6.32-754.el6.x86_64)
root (hd0,0)
kernel /vmlinuz-2.6.32-754.el6.x86_64 ro root=/dev/mapper/vg_lib1044-
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_lib1044/lv_swap rd_LVM LV=vg_lib1044/lv_root rd_NO_DM
initrd /initramfs-2.6.32-754.el6.x86_64.img
```

- On RHEL7.x/OL7.x/SLES 12 SP4/SLES 15, add a boot entry with the new ramdisk image file name in `/etc/grub.d/40_custom`. For example,

Update the new entry in the `grub.cfg` file, and make this entry as default so that system boots with the new ramdisk image automatically.

```
[root@hostname ~]# grub2-mkconfig -o /boot/grub2/grub.cfg
[root@hostname ~]# grub2-set-default 'PowerPath - Red Hat Enterprise
Linux Server'
```

On SLES 12 SP4/SLES 15:

```
[root@hostname ~]# grub2-mkconfig -o /boot/grub2/grub.cfg
[root@hostname ~]# grub2-set-default 'PowerPath - SLES 12-SP4'
```

## 11. Reboot the host so that the new ramdisk image containing 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 34
- [Uninstall using manual steps](#) on page 34

### Uninstall using `emcpbfs_linux_script`

#### Before you begin

**Note:** In a multiboot 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 prepare PowerPath for uninstallation.

If the ERROR: PowerPath is not configured on Boot-From-SAN message is displayed, it indicates that `emcpbfs_linux_config` has not been used during PowerPath configuration.

See [Uninstall using manual steps](#) on page 34 to uninstall PowerPath.

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

**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 so that PowerPath does not automatically start on reboot.
2. Edit the `/etc/lvm/lvm.conf` file to restore the filter and accept SCSI devices.

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

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. Perform the following steps to re-create a new ramdisk image without PowerPath modules and reboot the host using the new ramdisk image without PowerPath modules:
  - a. Move dracut configuration files for PowerPath out of `/usr/share/dracut/modules.d/` by entering the following:

```
[root@hostname ~]# mv /usr/share/dracut/modules.d/90powerpath ~/dracut/
```

- b. Generate the ramdisk image file:

```
[root@hostname ~]# dracut /boot/initramfs-$(uname -r).img $(uname -r)
```

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

## Uninstall PowerPath in an RHEL 7.x, OL 7.x, SLES 12 SP4, or SLES 15 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, SLES 12 SP4, or SLES 15 host:

- [Uninstall using emcpbfs\\_linux\\_script](#) on page 35
- [Uninstall using manual steps](#) on page 36

### Uninstall using `emcpbfs_linux_script`

#### Before you begin

**Note:** In a multiboot 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`.

#### About this task

#### CAUTION

On SLES12 SP4 and SLES15, LVM does not activate logical volumes if their physical volume (PV) appears on multiple devices. In case of BFS hosts with the root filesystem on a logical volume, if the host sees multiple PVs on boot, the host does not boot up as the root logical volume and swap volume cannot be activated.

This happens when the host is booted and PowerPath or any other multipathing software is not configured on the host. To enable the host to boot properly, change the `allow_changes_with_duplicate_pvs` parameter in the `devices{}` section in `/etc/lvm/lvm.conf` to 1. If the parameter is not present, add the same. Ensure that these changes are made before upgrading or uninstalling PowerPath, else the host may not come up on reboot.

See [Booting with duplicate PVs on SLES BFS hosts](#) on page 129 for more information.


### Procedure

1. Run `/etc/opt/emcpower/emcpbfs_linux remove` to prepare for PowerPath uninstallation.


If the ERROR: PowerPath is not configured on Boot-From-SAN message is displayed, it indicates that `emcpbfs_linux config` has not been used during PowerPath configuration.

See [Uninstall using manual steps](#) on page 36 to uninstall PowerPath.

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

 **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

### About this task

#### CAUTION

On SLES12 SP4 and SLES15, LVM does not activate logical volumes if their physical volume (PV) appears on multiple devices. In case of BFS hosts with the root filesystem on a logical volume, if the host sees multiple PVs on boot, the host does not boot up as the root logical volume and swap volume cannot be activated.

This happens when the host is booted and PowerPath or any other multipathing software is not configured on the host. To enable the host to boot properly, change the `allow_changes_with_duplicate_pvs` parameter in the `devices{}` section in `/etc/lvm/lvm.conf` to 1. If the parameter is not present, add the same. Ensure that these changes are made before upgrading or uninstalling PowerPath, else the host may not come up on reboot.

See [Booting with duplicate PVs on SLES BFS hosts](#) on page 129 for more information.

### Procedure

1. Disable the following PowerPath services from starting on host reboot:

```
[root@hostname ~]# systemctl disable PowerPath.service
[root@hostname ~]# systemctl disable PowerPathPoststart.service
[root@hostname ~]# systemctl disable EmcpMond.service
```

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

```
UUID=b4caa11e-9aa1-4777-bd5a-012e827514b1 /boot/efi xfs defaults 1 2
#/dev/emcpowera1 /boot/efi xfs defaults 1 2
```

3. Remount or boot so that it is mounted on standard partition as original.

```
root@lnq196018 ~]# umount /boot/efi
[root@lnq196018 ~]# mount /boot/efi
```

4. Edit the `/etc/lvm/lvm.conf` file to change the `global_filter`.

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

5. Recreate the ramdisk image without PowerPath modules using the following steps:

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

```
[root@hostname ~]# mv /usr/lib/dracut/modules.d/90powerpath ~/dracut/
```

- b. Generate the new ramdisk image file.

On RHEL7.x/OL7.x:

```
[root@hostname ~]# dracut /boot/initramfs-$(uname -r).img $(uname -r)
```

On SLES 12 SP4/SLES 15:

```
[root@hostname ~]# dracut /boot/initrd-$(uname -r) $(uname -r)
```

- c. 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.
```

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

On RHEL7.x/OL7.x:

```
[root@hostname ~]# grub2-mkconfig -o /boot/efi/EFI/redhat/grub.cfg
```

On SLES 12 SP4/SLES 15:

```
[root@hostname ~]# grub2-mkconfig -o /boot/grub2/grub.cfg
```

- e. Set the new ramdisk image without PowerPath modules that are created in the previous step, as the default boot entry.

On RHEL 7.x/OL 7.x:

```
[root@hostname ~]# grub2-set-default 'Red Hat Enterprise Linux Server'
```

On SLES 12 SP4/SLES 15:

```
[root@hostname ~]# grub2-set-default 'SLES 12-SP4'
```

6. Reboot the host and then uninstall PowerPath.

## Configure PowerPath in an 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 an 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 that 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 to make 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. It prevents the operating system 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 operating system 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`, `ext4`, or `xf`s 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`, `ext4`, or `xf`s and the other as swap and designate the mount points for each volume.

9. Ensure that GRUB is the selected boot loader 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, see [Disabling dm-multipath on an RHEL 6.x or OL 6.x host](#) on page 126.

12. If the operating system 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` file is set properly with the current boot kernel before running `emcpbfs_linux config/remove`.

Run the following command to configure PowerPath:

```
[root@hostname~] # /etc/opt/emcpower/emcpbfs_linux_config
```

`emcpbfs_linux config` 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.

- **Using 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.

```
[root@hostname ~]# umount /boot/efi
[root@hostname ~]# 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 a new ramdisk image to include PowerPath modules.

```
[root@hostname ~]# dracut /boot/initramfs-PP-$(uname -r).img $(uname -r)
```

- g. Add a 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 (2.6.32-754.el6.x86_64)
root (hd0,0)
kernel /vmlinuz-2.6.32-754.el6.x86_64 ro root=/dev/mapper/vg_lib1044-
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_lib1044/lv_swap rd LVM LV=vg_lib1044/lv_root rd_NO_DM
initrd /initramfs-PP-2.6.32-754.el6.x86_64.img

Red Hat Enterprise Linux (2.6.32-754.el6.x86_64)
root (hd0,0)
kernel /vmlinuz-2.6.32-754.el6.x86_64 ro root=/dev/mapper/vg_lib1044-
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_lib1044/lv_swap rd LVM LV=vg_lib1044/lv_root rd_NO_DM
initrd /initramfs-2.6.32-754.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, 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 run the following steps manually.

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

## Configure PowerPath in an RHEL 7.x, OL 7.x, SLES 12 SP4, or SLES 15 Boot from SAN setup in an UEFI host

To configure a PowerPath root device using the LVM on an RHEL 7.x, OL 7.x, SLES 12 SP4, or SLES 15 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 SP4 or SLES 15, ensure that native multipathing is not automatically activated. Because, there is not way to disable native multipathing after the operating system installation is complete.
- During configuration, it is advisable to have only the boot LUN connected to the host to make it easier to identify the boot disk during installation. It is also advisable to have only one path from the array to the host when installing the operating system. It prevents the operating system 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 operating system installation is complete.
- To review the procedure for disabling dm-multipath on an RHEL 7.x or OL 7.x, see [Disabling dm-multipath on an RHEL 7.x or OL 7.x host](#) on page 126.

4. At installation, select **Basic Storage Devices** as the type of device for installation.
5. Proceed with installation.
  - For OL 7.x and RHEL7.x UEFI Boot From SAN:  
Create `/boot` as basic partition more than 500MB (ext4), a `/boot/efi` as basic partition more than 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.
  - For SLES UEFI Boot From SAN:  
Create 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 operating system 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.

```
[root@hostname ~]# umount /boot/efi
[root@hostname ~]# mount /boot/efi
```

c. Setup the LVM filter in `/etc/lvm/lvm.conf`.

```
#global_filter = [ "a./.*/" ]
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 a new ramdisk image to include PowerPath modules:

On RHEL 7.x/OL 7.x:

```
[root@hostname ~]# dracut /boot/initramfs-PP-$(uname -r).img $(uname -r)
```

On SLES 12 SP4/SLES 15:

```
[root@hostname ~]# dracut /boot/initrd-PP-$(uname -r) $(uname -r)
```

g. Add a 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-3.10.0-327.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-3.10.0-327.el7.x86_64.img
```

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

On RHEL 7.x/OL 7.x:

```
[root@hostname ~]# grub2-mkconfig -o /boot/efi/EFI/redhat/grub.cfg
[root@hostname ~]# grub2-set-default 'PowerPath RedHat Enterprise Linux
Server'
```

On SLES 12 SP4/SLES 15:

```
[root@hostname ~]# grub2-mkconfig -o /boot/grub2/grub.cfg
[root@hostname ~]# grub2-set-default 'PowerPath - SLES 12-SP4'
```

10. Reboot the host to have the new ramdisk image that has PowerPath modules to take 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 43
- [Upgrade using manual steps](#) on page 44

### 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, it indicates that `emcpbfs_linux config` has not been used during PowerPath configuration.

See [Upgrade using manual steps](#) on page 44 to upgrade PowerPath.

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

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

3. Reboot the host.
4. Upgrade and start PowerPath.

**Note:** To review the procedure for upgrading PowerPath, see [Upgrade PowerPath from a compressed archive](#) on page 90.

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

**Note:** `emcpbfs_linux config/remove` adds a 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/emcpower1 /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. Upgrade and start PowerPath.

 **Note:** To review the procedure for upgrading PowerPath, see [Upgrade PowerPath from a compressed archive](#) on page 90.

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

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

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

```
[root@hostname ~]# umount /boot/efi
[root@hostname ~]# mount /boot/efi
```

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

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

10. Delete the old ramdisk image containing PowerPath modules that was created in step 12 in [Configure PowerPath in an RHEL 6.x or OL 6.x Boot from SAN setup in an UEFI host](#) on page 38.

```
[root@hostname ~]# rm -f /boot/initramfs-PP-$(uname -r).img
```

11. Build a new ramdisk image to include upgraded PowerPath modules.

```
[root@hostname ~]# dracut /boot/initramfs-PP-$(uname -r).img $(uname -r)
```

- The `/boot/efi/EFI/redhat/grub.conf` already contains an entry with ramdisk image that is created during initial setup, see [Configure PowerPath in an RHEL 6.x or OL 6.x Boot from SAN setup in an UEFI host](#) on page 38.

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

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

## Upgrade PowerPath in an RHEL 7.x, OL 7.x, SLES 12 SP4, or SLES 15 Boot from SAN setup in an UEFI host

To upgrade PowerPath in a Boot from SAN setup on an RHEL 7.x, OL 7.x, SLES 12 SP4, or SLES 15 host, use one of the following methods:

- [Upgrade using the `emcpbfs\_linux` script](#) on page 45
- [Upgrade using manual steps](#) on page 46

### Upgrade using the `emcpbfs_linux` script

#### About this task

#### CAUTION

On SLES12 SP4 and SLES15, LVM does not activate logical volumes if their physical volume (PV) appears on multiple devices. In case of BFS hosts with the root filesystem on a logical volume, if the host sees multiple PVs on boot, the host does not boot up as the root logical volume and swap volume cannot be activated.

This happens when the host is booted and PowerPath or any other multipathing software is not configured on the host. To enable the host to boot properly, change the `allow_changes_with_duplicate_pvs` parameter in the `devices{}` section in `/etc/lvm/lvm.conf` to 1. If the parameter is not present, add the same. Ensure that these changes are made before upgrading or uninstalling PowerPath, else the host may not come up on reboot.

See [Bootting with duplicate PVs on SLES BFS hosts](#) on page 129 for more information.


#### Procedure

- 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, it indicates that `emcpbfs_linux config` has not been used during PowerPath configuration.

See [Upgrade using manual steps](#) on page 46 to upgrade PowerPath.

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

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

- Reboot the host.
- Upgrade and start PowerPath.

To review the procedure for upgrading PowerPath, see [Upgrade PowerPath from a compressed archive](#) on page 90.

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

**Note:** `emcpbfs_linux config/remove` adds a 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 SP4/SLES 15 file and sets the proper *default* value.

## Upgrade using manual steps

### About this task

#### CAUTION

On SLES12 SP4 and SLES15, LVM does not activate logical volumes if their physical volume (PV) appears on multiple devices. In case of BFS hosts with the root filesystem on a logical volume, if the host sees multiple PVs on boot, the host does not boot up as the root logical volume and swap volume cannot be activated.

This happens when the host is booted and PowerPath or any other multipathing software is not configured on the host. To enable the host to boot properly, change the `allow_changes_with_duplicate_pvs` parameter in the `devices{}` section in `/etc/lvm/lvm.conf` to 1. If the parameter is not present, add the same. Ensure that these changes are made before upgrading or uninstalling PowerPath, else the host may not come up on reboot.

See [Booting with duplicate PVs on SLES BFS hosts](#) on page 129 for more information.

### Procedure

1. Disable the following PowerPath services from starting on host reboot.

```
[root@hostname ~]# systemctl disable PowerPath.service
[root@hostname ~]# systemctl disable PowerPathPost-start.service
[root@hostname ~]# 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 and restore the default filter in `/etc/lvm/lvm.conf` such that LVM filter accepts all devices.

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

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

On RHEL 7.x/OL 7.x:

```
[root@hostname ~]# grub2-set-default 'Red Hat Enterprise Linux Server'
```

On SLES 12 SP4/SLES 15:

```
[root@hostname ~]# grub2-set-default 'SLES 12-SP4'
```

5. Reboot the host.
6. Upgrade and start PowerPath.

To review the procedure for upgrading PowerPath, see [Upgrade PowerPath from a compressed archive](#) on page 90.

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

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

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

```
[root@hostname ~]# umount /boot/efi
[root@hostname ~]# mount /boot/efi
```

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

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

10. Delete the old ramdisk image containing PowerPath modules that was created in [Configure PowerPath in an RHEL 7.x, OL 7.x, SLES 12 SP4, or SLES 15 Boot from SAN setup in an UEFI host](#) on page 40.

```
[root@hostname ~]# rm -f /boot/initramfs-PP-$(uname -r).img
```

11. Build a new ramdisk image to include upgraded PowerPath modules.

- On RHEL 7.x or OL 7.x:

```
[root@hostname ~]# dracut /boot/initramfs-PP-$(uname -r).img $(uname -r)
```

- On SLES 12 SP4 or SLES 15:

```
[root@hostname ~]# dracut /boot/initrd-$(uname -r) $(uname -r)
```

12. The `/etc/grub.d/40_custom` already contains an entry with ramdisk image that is created during initial setup, see [Configure PowerPath in an RHEL 7.x, OL 7.x, SLES 12 SP4, or SLES 15 Boot from SAN setup in an UEFI host](#) on page 40.

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

13. 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 48
- [Uninstall using manual steps](#) on page 48

### 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 prepare PowerPath for uninstallation.

If the `ERROR: PowerPath is not configured on Boot-From-SAN.` message is displayed, it indicates that `emcpbfs_linux config` has not been used during PowerPath configuration.

See [Uninstall using manual steps](#) on page 48 to uninstall PowerPath.

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

**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 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 so that PowerPath does not automatically start on reboot.
2. Edit the `/etc/lvm/lvm.conf` file to restore the filter and accept SCSI device nodes.

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



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/emcpowera1 /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 a new ramdisk image without the PowerPath modules.

Recreate the new ramdisk image without PowerPath modules using the following steps:

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

```
[root@hostname ~]# mv /usr/share/dracut/modules.d/90powerpath ~/dracut/
```

- b. Generate the ramdisk image file:

```
[root@hostname ~]# dracut /boot/initramfs-PP-$(uname -r).img $(uname -r)
```

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

## Uninstall PowerPath in an RHEL 7.x, OL 7.x, SLES 12 SP4, or SLES 15 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, SLES 12 SP4, or SLES 15 host:

- [Uninstall using emcpbfs\\_linux\\_script](#) on page 49
- [Uninstall using manual steps](#) on page 50

### 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 SP4/SLES 15 file is set properly with the current boot kernel before running `emcpbfs_linux config/remove`.

#### About this task

**CAUTION**  
On SLES12 SP4 and SLES15, LVM does not activate logical volumes if their physical volume (PV) appears on multiple devices. In case of BFS hosts with the root filesystem on a logical

volume, if the host sees multiple PVs on boot, the host does not boot up as the root logical volume and swap volume cannot be activated.

This happens when the host is booted and PowerPath or any other multipathing software is not configured on the host. To enable the host to boot properly, change the `allow_changes_with_duplicate_pvs` parameter in the `devices{}` section in `/etc/lvm/lvm.conf` to 1. If the parameter is not present, add the same. Ensure that these changes are made before upgrading or uninstalling PowerPath, else the host may not come up on reboot.

See [Booting with duplicate PVs on SLES BFS hosts](#) on page 129 for more information.


### Procedure

1. Run `/etc/opt/emcpower/emcpbfs_linux remove` to prepare PowerPath for uninstallation.


If the ERROR: PowerPath is not configured on Boot-From-SAN. message is displayed, it indicates that `emcpbfs_linux config` has not been used during PowerPath configuration.

See [Uninstall using manual steps](#) on page 50 to uninstall PowerPath.

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

 **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 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 SP4/SLES 15 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

### About this task

#### CAUTION

On SLES12 SP4 and SLES15, LVM does not activate logical volumes if their physical volume (PV) appears on multiple devices. In case of BFS hosts with the root filesystem on a logical volume, if the host sees multiple PVs on boot, the host does not boot up as the root logical volume and swap volume cannot be activated.

This happens when the host is booted and PowerPath or any other multipathing software is not configured on the host. To enable the host to boot properly, change the `allow_changes_with_duplicate_pvs` parameter in the `devices{}` section in `/etc/lvm/lvm.conf` to 1. If the parameter is not present, add the same. Ensure that these changes are made before upgrading or uninstalling PowerPath, else the host may not come up on reboot.

See [Booting with duplicate PVs on SLES BFS hosts](#) on page 129 for more information.

### Procedure

1. Disable the following PowerPath services from starting on host reboot:

```
[root@hostname ~]# systemctl disable PowerPath.service
[root@hostname ~]# systemctl disable PowerPathPost-start.service
[root@hostname ~]# systemctl disable EmcpMond.service
```

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

```
UUID=b4caa11e-9aa1-4777-bd5a-012e827514b1 /boot/efi xfs defaults 1 2
#/dev/emcpowera1 /boot/efi xfs defaults 1 2
```

3. Remount or boot so that it is mounted on standard partition as original.

```
root@lnq196018 ~]# umount /boot/efi
[root@lnq196018 ~]# mount /boot/efi
```

4. Edit the `/etc/lvm/lvm.conf` file to change the `global_filter`.

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

5. Perform the following steps to re-create a new ramdisk image without PowerPath modules and reboot the host using the new ramdisk image without PowerPath modules:

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

```
[root@hostname ~]# mv /usr/lib/dracut/modules.d/90powerpath ~/dracut/
```

b. Generate the new ramdisk image file.

On RHEL7.x/OL7.x:

```
[root@hostname ~]# dracut /boot/initramfs-$(uname -r).img $(uname -r)
```

On SLES 12 SP4/SLES 15:

```
[root@hostname ~]# dracut /boot/initrd-$(uname -r) $(uname -r)
```

c. 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.
```

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

On RHEL 7.x/OL 7.x:

```
[root@hostname ~]# grub2-mkconfig -o /boot/efi/EFI/redhat/grub.cfg
```

On SLES 12 SP4/SLES 15:

```
[root@hostname ~]# grub2-mkconfig -o /boot/grub2/grub.cfg
```

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

On RHEL 7.x/OL 7.x:

```
[root@hostname ~]# grub2-set-default 'Red Hat Enterprise Linux Server'
```

On SLES 12 SP4/SLES 15:

```
[root@hostname ~]# grub2-set-default 'SLES 12-SP4'
```

6. 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](#)..... 54
- [Before you add or remove paths](#)..... 55
- [Adding new paths to a PowerPath logical device](#)..... 56
- [Removing some of the paths to PowerPath logical devices](#)..... 57

## 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

#### About this task

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

#### Procedure

1. Add the new LUNs into the configuration.

**Note:** If PowerPath already manages some LUNs from the same array on which the new LUNs are added, PowerPath automatically triggers a bus scan and configuration of the new LUNs within two minutes.

If PowerPath does not manage any LUNs from the same array on which the new LUNs are added or if the LUNs are added on Dell SC arrays, continue to step 2 to detect and configure the newly added LUNs.

2. Scan for the LUNs by running the following script:

```
[root@hostname ~]# /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:

```
[root@hostname ~]# /etc/opt/emcpower/emcplun_linux scan hba
```

For scanning iSCSI software initiator, type:

```
[root@hostname ~]# /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 about 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 that needs to be removed.

```
[root@hostname ~]# /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.

```
[root@hostname ~]# /etc/opt/emcpower/emcplun_linux remove <pseudo-list>
[-noprompt]
```

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

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

## Resizing a LUN

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.

```
[root@hostname ~]# 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 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.

```
[root@hostname ~]# powermt display
```

2. Confirm the configuration of the logical devices to which new paths are added.

```
[root@hostname ~]# 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:**

```
[root@hostname ~]# /etc/opt/emcpower/emcplun_linux scan <option> [-noprmt]
```

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.

```
[root@hostname ~]# echo "---" > /sys/class/scsi_host/host<#>/scan
```

- b. Configure new paths under PowerPath control.

```
[root@hostname ~]# powermt config
```

- c. Inspect the new PowerPath configuration.

```
[root@hostname ~]# 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.

```
[root@hostname ~]# 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.

### About this task


 **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 associated with paths that 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.

```
[root@hostname ~]# 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.

```
[root@hostname ~]# powermt display dev=all
```

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

8. Obtain bus, TID, and LUN.

```
[root@hostname ~]# 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.

```
[root@hostname ~]# 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.

```
[root@hostname ~]# powermt restore
```

Note down all the native device names that are associated with the paths that are listed as dead.

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

```
[root@hostname ~]# powermt remove dev=<path>|<device>|all
```

**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.

```
[root@hostname ~]# echo 1 > /sys/block/<Native device>/device/delete
```

17. Inspect the new PowerPath configuration.

- a. Check the path state.

```
[root@hostname ~]# 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.

```
[root@hostname ~]# powermt save
```



# CHAPTER 5

## PowerPath features for PowerMax array

PowerPath for Linux 6.5 supports the following features when the host is connected to a PowerMax array running microcode 5978 and above.


- [PowerPath Host Registration](#)..... 62
- [FX Licensing](#)..... 64
- [PowerPath Device IO Statistics](#)..... 64
- [PowerPath Device in Use Reporting](#)..... 65

# PowerPath Host Registration

This section provides information about the PowerPath Host Registration feature.

## Overview

PowerPath Host Registration is a feature that enables a host to be registered with a storage array. This makes it easier to configure storage for a new host from an array. This section describes how PowerPath Host Registration works with a PowerMax array running microcode 5978 and above only.

 **Note:** All references to PowerMax refers only to PowerMax running microcode 5978 and above.

## Initial Registration and Storage Configuration

This section provides information about the initial registration and storage configuration of the host with a storage array.

When a new host is zoned to a PowerMax array, virtual devices that are known as LUNZ devices are presented to the host, one per logical path. These devices are automatically configured in PowerPath. However, `powermt display` does not show these devices. Once the LUNZ devices are configured in PowerPath, registration commands are automatically sent on each path to the array through the LUNZ devices.

The registration commands contain information such as hostname, UUID, operating system version, and PowerPath version. On receipt of these registration commands, the PowerMax array creates a host record and all the WWNs of the host's HBAs are grouped under this record. If AUTO IG creation feature is enabled on the array, an Initiator Group is created automatically.

After the Initiator Group is created (automatically or manually), other entities such as, the Storage Group, Port Group, and Masking view can be created to configure data LUNs for the host. On creation of these entities, PowerPath automatically triggers a bus scan on the host and then runs a `powermt config` command to configure the new LUNs on the host. The new LUNs are available for use by applications.

## Modes of operation

PowerPath Host Registration works in the following two modes:

### Automatic

PowerPath Host Registration is configured automatically in the following scenarios:

- On PowerPath start or host reboot, the registration is attempted within 5 minutes.
- Once every 24 hours after the initial registration, registration is again attempted. This attempt helps the array in being aware that the host is active.
- If any of the host information such as the hostname changes, a registration is attempted within 60 minutes.
- If the host is zoned to a new PowerMax array, the registration is attempted within 5 minutes.

## Manual

The `powermt update host_registration` command can be run at any time to send registration to all connected PowerMax arrays. This is useful in scenarios where the hostname is changed and the array needs to be updated immediately.

## Settings

This section provides information about the settings that you need to configure the PowerPath Host Registration.

PowerPath Host Registration is enabled by default for Symmetrix arrays, which can be seen with the `powermt display options` command.

```
[root@hostname ~]# powermt display options

Show CLARiiON LUN names:      true

Path Latency Monitor: Off

Performance Monitor: disabled

Autostandby:  IOs per Failure (iopf): enabled
               iopf aging period      : 1 d
               iopf limit              : 6000

Storage
System Class  Attributes
-----
Symmetrix    periodic autorestore = on
              reactive autorestore = on
              status = managed
              proximity based autostandby = off
              path retry timeout = 5
              all paths dead retry timeout = 45
              auto host registration = enabled
              device to array performance report = enabled
              device in use to array report = enabled
```

This setting can be changed with the following `powermt set` command:

```
[root@hostname ~]# powermt set auto_host_registration={on|off} class=<class>
```

There is also a setting on the PowerMax array which enables or disables accepting host registration commands. If the host registration is disabled on the array, but enabled on the host, registration is not attempted.

## Log messages

When a registration is attempted, either automatically or manually with the `powermt update host_registration` command, messages similar to the following are logged in the system log.

```
Jun 22 07:21:42 lnd198016 kernel: Info:Mpx:PowerPath Auto Host Registration to
Symmetrix-000197801413 is successfully done on Bus 2 Port FA 2d:04
Jun 22 07:21:42 lnd198016 kernel: Info:Mpx:PowerPath Auto Host Registration to
Symmetrix-000197801413 is successfully done on Bus 2 Port FA 2d:05
Jun 22 07:21:42 lnd198016 kernel: Info:Mpx:PowerPath Auto Host Registration to
Symmetrix-000197801413 is successfully done on Bus 5 Port FA 2d:05
Jun 22 07:21:42 lnd198016 kernel: Info:Mpx:PowerPath Auto Host Registration to
Symmetrix-000197801413 is successfully done on Bus 5 Port FA 2d:04
```

The following message is logged when an Initiator Group is created for the host. This message is logged every time after PowerPath is started.

```
Jun 22 07:21:42 lnd198016 kernel: Info:Mpx:Initiator Group name on array
000197801413 is lnd198016_ig.
```

The Solutions Enabler software is used to check the details of the host from the array after registration. In this PowerPath release, cluster details are not sent in the registration commands and hence the array does not display any cluster details even if the host is running a cluster software. Also, if there are virtual machines running on the host, their details are also not sent in the registration commands.

```
# ./symcfg -sid 413 list -ppreg -host host123

Symmetrix ID      : 000197801413

Host Name:        host123
OS Version:       OL7.5
OS Revision:      Linux - 4.14.35-1818.0.9
Hardware Vendor Name: Dell Inc.

PowerPath Version: 6.5.0
PowerPath Patch Level: 00.00-095
PowerPath License Info: xxxx-xxxx-xxxx-xxxx-xxxx-xxxx
Host Registration Time: 06/22/2018 16:33:09
Host Connectivity type: FC
Cluster Info:
  Cluster Name:
  Cluster Node Name:
WWNs:
  (1): 21000024FF2066E2
  (2): 21000024FF2066E3
VMs:
```

## FX Licensing

Licensing PowerPath for Linux usually involves acquiring a license key and entering it with the **emcprep** utility. If a PowerMax array is purchased with an FX bundle, all PowerPath hosts attached to it is automatically licensed by attaching data LUNs to the host.

The number of hosts that can thus be licensed is limited to 75 hosts per FX bundle. PowerPath is fully licensed not just for the PowerMax LUNs and paths, but for all LUNs and paths from other array types as well. PowerPath Migration Enabler functionality is also available through the FX bundle.

All **powermt** commands work normally without generating any license related warnings. To verify that the host is licensed through FX, you can use any of the following commands:

- [root@hostname ~]# **powermt check\_registration**  
Host is licensed through FX array.
- [root@hostname ~]# **emcprep -l**  
Host is licensed through FX array.

## PowerPath Device IO Statistics

This section provides information on the PowerPath Device IO Statistics feature.

### Overview

The PowerPath Device IO Statistics feature provides host IO performance collection and reporting to PowerMax. These statistics are collected per PowerPath managed PowerMax device and reported to array at one minute intervals. PowerMax populates these performance statistics in Unisphere providing a single pane of display with both host and array level statistics.



## Settings

This section provides information about the settings that you need to configure the PowerPath Device IO Statistics.

PowerPath Device IO Statistics is enabled by default for Symmetrix arrays. It can be seen with the `powermt display options` command.

```
[root@hostname ~]# powermt display options

    Show CLARiiON LUN names:      true

    Path Latency Monitor: Off

    Performance Monitor: disabled

    Autostandby:  IOs per Failure (iopf): enabled
                  iopf aging period   : 1 d
                  iopf limit          : 6000

    Storage
    System Class  Attributes
    -----
    Symmetrix     periodic autorestore = on
                  reactive autorestore = on
                  status = managed
                  proximity based autostandby = off
                  path retry timeout = 5
                  all paths dead retry timeout = 45
                  auto host registration = enabled
                  device to array performance report = enabled
                  device in use to array report = enabled
```

This setting can be changed with the following `powermt set` command:

```
[root@hostname ~]# powermt set device_perf_to_array_report={on|off}
[class=<class>|all]
```

## PowerPath Device in Use Reporting

This section provides information on the PowerPath Device in Use Reporting feature.

### Overview

PowerPath Device in Use Reporting feature provides a way to report if devices are used on the host. Once every 24 hours, for each PowerMax device, PowerPath sends details such as if the device is mounted and name of the process that used the device in the preceding 24 hours to the array.

The device usage statistics can be used by storage administrators to determine LUN usage on hosts as well as in array capacity and resource planning.

### Settings

This section provides information about the settings that you need to configure the PowerPath Device in Use.

PowerPath Device in Use is enabled by default for Symmetrix arrays. It can be seen with the `powermt display options` command.

```
[root@hostname ~]# powermt display options

    Show CLARiiON LUN names:      true
```

```
Path Latency Monitor: Off

Performance Monitor: disabled

Autostandby: IOs per Failure (iopf): enabled
              iopf aging period      : 1 d
              iopf limit              : 6000

Storage
System Class  Attributes
-----
Symmetrix    periodic autorestore = on
              reactive autorestore = on
              status = managed
              proximity based autostandby = off
              path retry timeout = 5
              all paths dead retry timeout = 45
              auto host registration = enabled
              device to array performance report = enabled
              device in use to array report = enabled
```

**This setting can be changed with the following powermt set command:**

```
[root@hostname ~]# powermt set dev_inuse_to_array_report={on|off}
[class=<class>|all]
```

# CHAPTER 6

## 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](#) .....68
- [PowerPath co-management behavior](#) ..... 68

## 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). See [Disable native multipathing](#) on page 125 for the steps to disable native multipathing.

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=&lt;class&gt;</code> Where <code>class</code> 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

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



# CHAPTER 7

## 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, hot swapping of an HBA, R1/R2 boot failover support, or SCSI-2 reservation support.

- [Configure remote management service](#)..... 72
- [Configure event monitoring with SNMP traps](#)..... 73
- [Changing the language of installation messages](#)..... 77
- [Saving new configurations](#).....78
- [Configuring LVM2 support](#)..... 78
- [Configuring PowerPath on cluster nodes](#) ..... 82
- [Hot swapping an HBA](#)..... 83
- [Configuring for R1/R2 boot failover support](#).....84
- [SCSI-2 reservation support in PowerPath for Linux](#) ..... 86
- [Enable scsi name on NetApp arrays](#).....86
- [Removing devices from PowerPath control](#)..... 87
- [Expand the file system on a pseudo-device](#)..... 87
- [Supporting more than 256 storage devices](#)..... 87
- [/etc/fstab options](#).....88

## Configure remote management service

PowerPath hosts can be remotely monitored by connecting it to the PowerPath Management Appliance (PPMA).

### About this task

With PPMA, the status of all the devices and paths managed by various PowerPath hosts can be viewed from a single window. Refer to the PowerPath Management Appliance documentation for more details. Refer to the PowerPath Management Appliance documentation for more details.

### Procedure

1. In order for the PowerPath hosts to be connected to the PowerPath Management Appliance, remote management service has to be enabled in the PowerPath management daemon with the following steps:

```
[root@hostname ~]# emcp_mond_edit

      Item                               Value
      ====                               =====
  1 Change remote management service status      Disabled
  2 Change management service heartbeat interval 10 minutes
  3 Change management service port                9083
  4 Single path gatekeeper devices                Allow
  5 Change SNMP trap generation status            Disabled
  6 Change "Path Dead" event                      Enabled, delay 60
sec
  7 Change "Path Alive" event                     Disabled, delay 60
sec
  8 Change "All Paths Dead" event                 Enabled, delay 300
sec
  9 Change "Not All Paths Dead" event             Disabled, delay 300
sec
 10 Change "Volume Dead" event                   Disabled, delay 300
sec
 11 Change "Max. Path Latency Exceeded" event    Disabled, delay 0
sec
 12 Write configuration file
 13 Exit

Enter number of desired choice [nn, q, ?]: 1

      Item                               Value
      ====                               =====
  1 Change remote management service status      Enabled
  2 Change management service heartbeat interval 10 minutes
  3 Change management service port                9083
  4 Single path gatekeeper devices                Allow
  5 Change SNMP trap generation status            Disabled
  6 Change "Path Dead" event                      Enabled, delay 60
sec
  7 Change "Path Alive" event                     Disabled, delay 60
sec
  8 Change "All Paths Dead" event                 Enabled, delay 300
sec
  9 Change "Not All Paths Dead" event             Disabled, delay 300
sec
 10 Change "Volume Dead" event                   Disabled, delay 300
sec
 11 Change "Max. Path Latency Exceeded" event    Disabled, delay 0
sec
 12 Write configuration file
 13 Exit
```



```

Enter number of desired choice [nn, q, ?]: 13

Item
====
1 Save configuration file, restart daemon and exit
2 Save configuration file and exit
3 Exit without saving configuration file
4 Back

Enter number of desired choice [1, q, ?]: 1
Writing configuration file.

Changed parameters:

Item                                     Changed Value
====                                     =====
remote management service                Enabled

```

**Note:** Change the management service port as required. Also ensure that the port configured for management service is opened in the firewall. Else, PowerPath Management Appliance will not be able to connect to the management daemon.

## Configure event monitoring with SNMP traps

The PowerPath management daemon provides a facility to generate SNMP traps for multipathing events that are listed in the `emcp_mond.conf` file.

### 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

**Table 5** Monitored events (continued)

Monitored event	Event cause	Event ID	SNMP Trap
	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 77.

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.

## Configure SNMP traps

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.

### Procedure

1. Copy over the sample `snmpd.conf` file from PowerPath install folder over the default `snmpd.conf`.

```
#cp /etc/opt/emcpower/snmpd.conf /etc/snmp/snmpd.conf
```

2. Edit `/etc/snmp/snmpd.conf` and add the following parameter to specify the destination to which the traps 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. It is the default port number of the host where the SNMP manager is listening.

For example,

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

3. Configure `snmpd` persistence across restart.
  - a. Use the `chkconfig` or `systemctl` command to verify that the `snmpd` service is enabled. If it is not enabled, enable the `snmpd` service.
  - b. Start the `snmpd` daemon.

On RHEL7, OL7, SLES 12 SP4, or SLES15:

```
[root@hostname ~]# systemctl start snmpd.service1
```

On all other kernels:

```
[root@hostname ~]# /etc/init.d/snmpd start
```

**Note:** It is vital that when snmp trap monitoring is enabled in PowerPath management daemon, the snmpd process should be running. If snmpd process needs to be stopped or disabled, SNMP trap generation should be first disabled in PowerPath management daemon. Otherwise, if snmpd process is stopped when SNMP trap generation is enabled, PowerPath management daemon may fail.

#### 4. Ensure that snmpd daemon is running.

```
[root@hostname ~]# ps -ef | grep snmpd
```

For example,

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

#### 5. Use the emcp\_mond\_edit utility to enable SNMP trap generation in PowerPath management daemon.

```
[root@hostname ~]# emcp_mond_edit

      Item                               Value
      ====                               =====
1 Change remote management service status Disabled
2 Change management service heartbeat interval 10 minutes
3 Change management service port 9083
4 Single path gatekeeper devices Allow
5 Change SNMP trap generation status Disabled
6 Change "Path Dead" event Enabled, delay 60
sec
7 Change "Path Alive" event Disabled, delay 60
sec
8 Change "All Paths Dead" event Enabled, delay 300
sec
9 Change "Not All Paths Dead" event Disabled, delay 300
sec
10 Change "Volume Dead" event Disabled, delay 300
sec
11 Change "Max. Path Latency Exceeded" event Disabled, delay 0
sec
12 Write configuration file
13 Exit[root@hostname ~]# emcp_mond_edit

      Item                               Value
      ====                               =====
1 Change remote management service status Disabled
2 Change management service heartbeat interval 10 minutes
3 Change management service port 9083
4 Single path gatekeeper devices Allow
5 Change SNMP trap generation status Disabled
6 Change "Path Dead" event Enabled, delay 60
sec
7 Change "Path Alive" event Disabled, delay 60
sec
8 Change "All Paths Dead" event Enabled, delay 300
sec
```

```

 9 Change "Not All Paths Dead" event          Disabled, delay 300
sec
10 Change "Volume Dead" event                Disabled, delay 300
sec
11 Change "Max. Path Latency Exceeded" event Disabled, delay 0
sec
12 Write configuration file
13 Exit[root@hostname ~]# emcp_mond_edit

```

```

      Item                               Value
      ====                               =====
 1 Change remote management service status Disabled
 2 Change management service heartbeat interval 10 minutes
 3 Change management service port             9083
 4 Single path gatekeeper devices             Allow
 5 Change SNMP trap generation status          Disabled
 6 Change "Path Dead" event                   Enabled, delay 60
sec
 7 Change "Path Alive" event                  Disabled, delay 60
sec
 8 Change "All Paths Dead" event              Enabled, delay 300
sec
 9 Change "Not All Paths Dead" event          Disabled, delay 300
sec
10 Change "Volume Dead" event                Disabled, delay 300
sec
11 Change "Max. Path Latency Exceeded" event Disabled, delay 0
sec
12 Write configuration file
13 Exit

```

Enter number of desired choice [nn, q, ?]: 5

```

      Item                               Value
      ====                               =====
 1 Change remote management service status Disabled
 2 Change management service heartbeat interval 10 minutes
 3 Change management service port             9083
 4 Single path gatekeeper devices             Allow
 5 Change SNMP trap generation status          Enabled
 6 Change "Path Dead" event                   Enabled, delay 60
sec
 7 Change "Path Alive" event                  Disabled, delay 60
sec
 8 Change "All Paths Dead" event              Enabled, delay 300
sec
 9 Change "Not All Paths Dead" event          Disabled, delay 300
sec
10 Change "Volume Dead" event                Disabled, delay 300
sec
11 Change "Max. Path Latency Exceeded" event Disabled, delay 0
sec
12 Write configuration file
13 Exit

```

Enter number of desired choice [nn, q, ?]: 13

```

      Item
      ====
 1 Save configuration file, restart daemon and exit
 2 Save configuration file and exit
 3 Exit without saving configuration file
 4 Back

```

Enter number of desired choice [1, q, ?]: 1  
Writing configuration file.

Changed parameters:

Item ====	Changed Value =====
SNMP trap generation	Enabled

**Note:** When snmp trap monitoring is enabled in PowerPath management daemon, ensure that snmpd process is running. If snmpd process needs to be stopped or disabled, SNMP trap generation should be first disabled in PowerPath management daemon. Otherwise, if snmpd process is stopped when SNMP trap generation is enabled, PowerPath management daemon may fail.

## Enabling or disabling of path dead or alive events

Use the `emcp_mond_edit` utility to configure the dead/alive path events.

### Editing the event

Use the `emcp_mond_edit` utility to configure the threshold for the events. Alternatively, to modify the threshold values manually, use the following procedure:

#### 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 seconds 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, SLES 12 SP4, or SLES15:

```
[root@hostname ~]# systemctl restart EmcpMond.service
```

- On all other kernels:

```
[root@hostname ~]# /etc/init.d/emcp_mond.sh stop
[root@hostname ~]# /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):

### About this task

```
[root@hostname ~]# export LANG=zh_CN.utf8
```

or

```
[root@hostname ~]# 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	
Portuguese (Brazilian)	pt_BR.utf8	
Spanish	es_ES.utf8	

## Saving new configurations

When new devices are added to a host or old devices removed from the host, the configuration is modified. Any change in the device configuration in the host requires you to carry out an explicit `powermt save` to save the changed configuration.

## Configuring LVM2 support

The LVM2 (logical volume manager) can recognize PowerPath as the single path to LUN and use the PowerPath pseudo devices. For LVM2 to recognize only the PowerPath pseudo devices,

modify the `/etc/lvm/lvm.conf` file in the host, and filter out the sd device nodes from the LVM2 internal cache.

### 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 [Disabling dm-multipath on an RHEL 6.x or OL 6.x host](#) on page 126 or [Disabling dm-multipath on an RHEL 7.x or OL 7.x host](#) on page 126.

### About this task

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:

**Table 7** Linux Kernel Filter Fields

Kernel	Filter
SLES 12 SP4	<pre>global_filter = [ "a/emcpower.*/", "r/sd.*/", "r/disk.*/" ]</pre>
SLES 15	<pre>global_filter = [ "a/emcpower.*/", "r/sd.*/", "r/disk.*/" ]</pre>
RHEL 6.x	<pre>filter = [ "a/emcpower.*/", "r/sd.*/", "r/disk.*/" ]</pre>
RHEL 7.x	<pre>global_filter = [ "a/emcpower.*/", "r/sd.*/", "r/disk.*/" ]</pre>

**Table 7** Linux Kernel Filter Fields (continued)

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

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

## Root file system mounted on a logical volume

### About this task

For details on the LVM filter that must be set, see:


- For BIOS host, [Configuring PowerPath in an RHEL 6.x or OL 6.x Boot from SAN setup in a BIOS host](#) on page 20 or [Configure PowerPath in an RHEL 7.x, OL 7.x, SLES 12 SP4, or SLES 15 Boot from SAN setup in a BIOS host](#) on page 22.
- For UEFI host, [Configure PowerPath in an RHEL 6.x or OL 6.x Boot from SAN setup in an UEFI host](#) on page 38 or [Configure PowerPath in an RHEL 7.x, OL 7.x, SLES 12 SP4, or SLES 15 Boot from SAN setup in an UEFI host](#) on page 40.

## 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`:

```
[root@hostname ~]# /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:

**Table 8** Kernel Filter

Kernel Filter	Filter
SLES 12 SP4	<pre>global_filter = [ "a /dev/.*/by-id/scsi-SSEAGATE_ST373207LC_3KT4PMAN.* ", "a/emcpower.* ", "r/sd.* ", "r/disk.* " ]</pre>
SLES 15	<pre>global_filter = [ "a /dev/.*/by-id/scsi-SSEAGATE_ST373207LC_3KT4PMAN.* ", "a/emcpower.* ", "r/sd.* ", "r/disk.* " ]</pre>
RHEL 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>
Oracle Linux 6.x	<pre>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 that are identified in steps 1 and 2 are listed in the command output.

## Configuring PowerPath on cluster nodes

### Before you begin

In cluster configurations, a set of LUNs is attached to every node in the cluster. To configure the cluster properly, ensure that the LUN to pseudo mapping is the same for all the LUNs in all the cluster nodes. It is also important that the LUN to pseudo mapping is made consistent on all the nodes before the cluster software is configured.

### About this task

To configure the same LUN to pseudo mapping easily with the help of the **emcpadm** utility, follow the procedure:

### Procedure

1. Export the LUN to pseudo name mapping of all the LUNs on any one node to a file:

```
[root@hostname ~]# emcpadm export_mappings -f powerpath.map
```

2. Copy the mapping file to all the other nodes of the cluster.
3. In each of the other nodes, perform the following steps to apply and validate the mappings:
  - a. Import the mappings from the mapping file.

```
[root@hostname ~]# emcpadm import_mappings -v -f powerpath.map
... Mapped emcpowera
... Mapped emcpowerd
... Mapped emcpowere
... Mapped emcpowerf
... Mapped emcpowerc
```

- b. Check if the mappings are consistent with the first node.

If the mappings are consistent, the *Remapping Status* column shows *no change*. In the following example, not all the LUN mappings are consistent:

```
[root@hostname ~]# emcpadm check_mappings -v -f powerpath.map
```

Imported Device Mapping	Remapping Status
emcpowera	no change
emcpowerc	no change
emcpowerd	no change
emcpowere	no change
emcpowerf	no change
emcpowerh already in use)	remaps: emcpowerb                      emcpowerh (* Device name
emcpowerg	no change

- c. Repeat steps a and b until the mappings are consistent. The *Remapping Status* column should show *no change* for all the devices similar to the following output:

```
[root@hostname ~]# emcpadm check_mappings -v -f powerpath.map
```

Imported Device Mapping	Remapping Status
emcpowera	no change
emcpowerc	no change
emcpowerd	no change
emcpowere	no change
emcpowerf	no change
emcpowerh already in use)	no change
emcpowerg	no change

```

-----
emcpowera      no change
emcpowerb      no change
emcpowerc      no change
emcpowerd      no change
emcpowere      no change
emcpowerf      no change
emcpowerg      no change

```

**Note:** When the mappings are imported, some of the pseudo device names or the device itself may be in use. In such scenarios, ensure that none of the devices are in use and repeat steps **a** and **b** until the mappings are consistent.

## 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 that is 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` is:

```
[root@hostname 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 number gives the corresponding PCI slot number.

For example, the available slots given from:

```
[root@hostname 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:

```
[root@hostname ~]# echo 0 > /sys/bus/pci/slots/<0196_0003>/power
```

Where: <0196\_0003> is the slot number that is 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:

```
[root@hostname ~]# 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 reenabling the HBA through PowerPath by running:


```
[root@hostname ~]# 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 Dell 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.

**Note:**

- 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.

- 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:

```
[root@hostname ~]# symaccess -sid <#> show view <view name>
```

- Each R1 device has only one mirror. (Concurrent SRDF is not supported.)
- Neither host has non-mirrored devices, BCVs, or gatekeepers.

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

## R1/R2 boot procedure

### Before you begin

### About this task

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 determines which Symmetrix volume IDs are valid (that is, the visible ones) and acts 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 operating system and PowerPath may not track that change.

Running `blockdev --rereadpt /dev/emcpowerX`, where `emcpowerX` is the R2 PowerPath pseudo device name, updates 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.

### About this task

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:

### About this task

- For a pseudo device, if one or more native device paths are removed from PowerPath control using `powermt remove dev=<path>|<device>`, all 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=<path>|<device> | class=<class>` 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 Unity, VNX, or CLARiiON devices, LUN nice name (s) and storage group (s) will 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. Run `blockdev --rereadpt /dev/<pseudo_device>` to re-read the partition table of the pseudo device.
3. Run `echo 1>/sys/block/<native_device>/device/rescan` to scan the native path.
4. 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

## Supporting more than 256 storage devices

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

### Before you begin

### About this task

- 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, see 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:

**Table 9** Operating System Support Steps

Operating System	Support steps:
RHEL 6.x and OL 6.x	<ol style="list-style-type: none"> <li>1. Create the <code>lpfc.conf</code> file in the <code>/etc/modprobe.d/</code> directory.</li> <li>2. Edit the file and set the <code>lpfc_max_luns</code> parameter to 512.</li> <li>3. Recreate the <code>initramfs</code> by running the <code>dracut -f</code> command.</li> <li>4. Reboot the host.</li> </ol> <p>For more information, see the EMC Host Connectivity Guide for Linux or the vendor website.</p>
RHEL 7.x, OL 7.x, SLES12.x, and SLES15	<ol style="list-style-type: none"> <li>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></li> <li>2. Recreate the <code>initrd</code>.</li> <li>3. Point the <code>initrd</code> to the new <code>initrd</code> when booting.</li> </ol>

## /etc/fstab options

### RHEL6.x and OL6.x

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/emcpower1 /mnt/emcpower1 ext3 _netdev 0 0
```

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

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

### RHEL7, OL7, SLES 12 SP4, and SLES 15

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/emcpower1 /mnt/emcpower1 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
```



# CHAPTER 8

## 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 Dell 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, PowerPath may report `no such file or directory` during the system initialization. Also, 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 Dell 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 PowerPath is listed. If the `lsof` command lists the processes along with opened PowerPath managed storage devices, 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. For information about using a localized version of the PowerPath installer when upgrading PowerPath, see [Changing the language of installation messages](#) on page 77.
- If you encounter any problems while upgrading PowerPath, see [Troubleshooting PowerPath](#) on page 101.

## Upgrading PowerPath

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

[Upgrade PowerPath from a compressed archive](#) on page 90

[Upgrade PowerPath across Linux platforms](#) on page 92

[Upgrade PowerPath before upgrading RHEL update on Linux](#) on page 94

### Upgrade PowerPath from a compressed archive

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

#### Procedure

1. Log in as root.

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

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

- Untar the PowerPath archive.

```
[root@hostname ~]# tar -xzf DelleMCPower.LINUX-<version>-<build>.tar.gz
```

- Upgrade to the new PowerPath version.

```
[root@hostname ~]# rpm -Uhv DelleMCPower.LINUX-<version>-<build>.<platform>.x86_64.rpm
```

Where:

- *Version* is the product version. For example, 6.5.0.00.00.
- *Build* is the software build number.
- *Platform* is the Linux distribution, for example SLES12SP4 or RHEL6.

- Start PowerPath.

- On RHEL 6.x or OL 6.x

```
[root@hostname ~]# /etc/init.d/PowerPath start
```

- On RHEL 7.x, OL 7.x, SLES 12 SP4, or SLES 15

```
[root@hostname ~]# systemctl start PowerPath.service
```

**Note:** If issues are observed during or after starting PowerPath and if the old PowerPath version needs to be restored, follow these steps 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 that are 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.

## 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.

### About this task

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 92
- [Upgrade PowerPath after upgrading from one major Linux release to the next](#) on page 92
- [Upgrade PowerPath before upgrading to the latest SLES Service Pack](#) on page 93
- [Upgrade PowerPath before upgrading RHEL update on Linux](#) on page 94
- [Upgrade PowerPath in a SteelEye LifeKeeper cluster environment](#) on page 94

## Upgrade the Linux kernel in a PowerPath installed host

### About this task

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 6.9 to RHEL 7.0.

### Procedure

1. Perform one of the following steps:
  - For RHEL6.x or OL6.x, rename `/etc/init.d/PowerPath` to `/etc/init.d/PowerPath.old`
  - For RHEL 7.x, OL 7.x, SLES 12 SP4, or SLES 15

```
[root@hostname ~]# systemctl disable PowerPath.service
[root@hostname ~]# systemctl disable PowerPathPost-start.service
[root@hostname ~]# 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. If you choose to migrate the file system to `ext4`, these options also enable migrating the existing `ext3` file system to `ext4` during the upgrade to RHEL 6.x. If you opt to use the `ext4migrate` boot option to migrate the `ext4`

file system, edit the `/etc/fstab` file manually after the upgrade. It ensures that the root file system is mounted as `ext4` on the next reboot.

### 3. Upgrade to the new PowerPath version.

```
[root@hostname ~]# rpm -Uhv DelleMCPower.LINUX-<version>-<build>.<platform>.x86_64.rpm
```

Where:

- *Version* is the product version. For example, 6.5.0.00.00.
- *Build* is the software build number.
- *Platform* is the Linux distribution, for example SLES12SP4 or RHEL6.

### 4. Start PowerPath.

- On RHEL 6.x or OL 6.x

```
[root@hostname ~]# /etc/init.d/PowerPath start
```

- On RHEL 7.x, OL 7.x, SLES 12 SP4, or SLES 15

```
[root@hostname ~]# systemctl start PowerPath.service
```

**Note:** If issues are observed during or after starting PowerPath and if the old PowerPath version needs to be restored, follow these steps 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 that are 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.

- After you have checked to ensure 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

- 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:

```
[root@hostname ~]# rpm -Uhv DelleMCPower.LINUX-<version>-<build>.<platform>.x86_64.rpm
```

Where:

- <version> is the product version. For example, 6.5.0.00.00.
  - <build> is the software build number.
  - <platform> is the Linux distribution, for example RHEL7.x 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 `lkbbackup -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 `lksstop -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 DelleMCPower.LINUX-<version>-<build>.<platform>.x86_64.rpm
```

Where:

- <version> is the product version; for example, 6.5.0.00.00.
- <build> is the software build number.
- <platform> is the Linux distribution, for example SLES12SP4 or RHEL6.

7. Reboot the host.
8. Verify the PowerPath configuration:
  - a. Run `powermt display`.
  - b. Run `powermt display dev=all`.
9. Run `lksstart` to start SteelEye LifeKeeper.
10. Repeat steps 1-9 for each node in the cluster.

### Results

- i** **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 9

## 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

### About this task

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

## Uninstall PowerPath

### About this task

#### 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.

```
[root@hostname ~]# rpm -qa | grep EMCpower.LINUX
```

4. Remove the software.

```
[root@hostname ~]# 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.

### About this task

- `/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 reinstall PowerPath, use `emcp_cleanup` to remove all these saved files from the host.

### Procedure

1. Remove the PowerPath license and all supporting files:

```
[root@hostname ~]# /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 10

## 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, RHEL7.x, SLES 12 SP4, SLES 15, OL6.x, and OL7.x boot-from-SAN when emcpbfs\\_linux is used](#).....104

## Recover from a failed PowerPath installation

### About this task

If PowerPath installation fails, correct error conditions are 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.

```
[root@hostname ~]# lsmod | grep emc
```

If the modules are loaded, type `/etc/init.d/PowerPath stop`, or for RHEL7, OL7, SLES 12 SP4, or SLES 15 type `systemctl stop PowerPath` to unload them.

2. Disable PowerPath from starting automatically on host boot.

- For RHEL6.x or OL6.x:

```
[root@hostname ~]# mv /etc/init.d/PowerPath /etc/init.d/PowerPath.saved
```

- For RHEL7, OL7, SLES 12 SP4, or SLES 15:

```
[root@hostname ~]# systemctl disable PowerPath.service
[root@hostname ~]# systemctl disable PowerPathPost-start.service
[root@hostname ~]# systemctl disable EmcpMond.service
```

3. Check the `/var/log/powerpath_install.log` file for more information about the failure.
4. Reboot the host.
5. Remove the PowerPath entry in the RPM database (if it exists).

Type the following commands to remove the entry:

```
[root@hostname ~]# rpm -qa | grep EMCpower.LINUX
[root@hostname ~]# rpm -ev EMCpower.LINUX-<version>-<build>
```

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

```
[root@hostname ~]# rm -rf /etc/opt/emcpower
```

## 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 it happens so, remove and then reinstall PowerPath.

#### Procedure


1. Stop PowerPath.
  - For RHEL 6.x or OL 6.x:

```
[root@hostname ~]# /etc/init.d/PowerPath stop
```

- RHEL7, OL7, SLES 12 SP4, or SLES 15:

```
[root@hostname ~]# systemctl stop 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.

```
[root@hostname ~]# rpm -qa | grep EMCpower.LINUX
[root@hostname ~]# 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.

```
[root@hostname ~]# rm -rf /etc/opt/emcpower
```

4. Reinstall PowerPath.

For the installation procedure, see [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, RHEL7.x, SLES 12 SP4, SLES 15, 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 that is listed as follows:

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

### About this task

```
/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/grub/menu.lst\_bkp\_emcpbfs\_linux\_config\_Jul112013\_14h20m36s

For RHEL 7.x, OL 7.x, SLES 12 SP4, or SLES 15,

- /boot/grub2/grub.cfg\_bkp\_emcpbfs\_linux\_config\_Nov262014\_15h29m07s

/etc/opt/emcpower/emcpbfs\_linux remove

**About this task**

/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/grub/menu.lst\_bkp\_emcpbfs\_linux\_remove\_Jul112013\_14h20m36s

For RHEL 7.x, OL 7.x, SLES 12 SP4, and SLES 15:

- /boot/grub2/grub.cfg\_bkp\_emcpbfs\_linux\_remove\_Nov262014\_14h39m50s

/etc/opt/emcpower/emcpbfs\_linux upgrade

**About this task**

/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, the script 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, OL 7.x, SLES 12 SP4, and SLES 15 use global filter and comment out the global\_filter:

```
/etc/lvm/lvm.conf
###emcpbfs_linux###Thu Jul 19 19:54:40 IST 2012###
# global_filter = [ "a|.*/" ]
#####
```

Boot from operating system CD or DVD in rescue mode, mount '/boot' and '/', 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 11

## 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`

### `/etc/init.d/boot.d/`

SLES only

- K18boot.powerpath
- S04boot.powerpath

#### **/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/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**

SLES15, SLES 12 SP4, RHEL 7, RHEL 6, OL 6, OL 7, OVM 3.x.x

- rollbackfiles.tar
- rollbackfiles <XXXXXX>.tar
- undolist
- undolist <XXXXXX>

#### **/etc/opt/emcpower/bin**

- envsubst
- gettext
- gettext.sh

### **/etc/opt/emcpower/EMCpower.LINUX-6.5.x.00.00**

- .pp\_build\_version
- .pp\_vendor
- .pp\_version
- EULA.pdf
- PowerPathPost
- PowerPath
- PowerPathPost-start.service (SLES15, SLES 12 SP4, RHEL7 and OL7)
- PowerPathPost-stop.service (SLES15, SLES 12 SP4, RHEL7 and OL7)
- PowerPath.service (SLES15, SLES 12 SP4, RHEL7 and OL7)
- boot-powerpath.sh
- boot.powerpath
- emcp\_mond.sh
- emcp\_mond.sles (SLES only)
- enable
- modprobe.conf.pp
- powerfcoe.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-6.5.x.00.00/bin**

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

### **/etc/opt/emcpower/EMCpower.LINUX-6.5.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-6.5.x.00.00/bin/driver](#)

#### SLES15

- sles15default\_x8664
- sles15default\_x8664/emcp.ko
- sles15default\_x8664/emcpdm.ko
- sles15default\_x8664/emcpgpx.ko
- sles15default\_x8664/emcpioc.ko
- sles15default\_x8664/emcpx.ko

#### SLES 12 SP4

- sles12sp4default\_x8664
- sles12sp4default\_x8664/emcp.ko
- sles12sp4default\_x8664/emcpdm.ko
- sles12sp4default\_x8664/emcpgpx.ko
- sles12sp4default\_x8664/emcpioc.ko
- sles12sp4default\_x8664/emcpx.ko

#### RHEL7

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

#### RHEL6

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

#### OL6

- 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/emcpx.ko

#### OL7

- 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/emcpx.ko

#### [/etc/opt/emcpower/EMCpower.LINUX-6.5.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
- libemcpxmp.so
- libpn.so

#### [/etc/opt/emcpower/EMCpower.LINUX-6.5.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
- libemcpxmp.so
- libpn.so

#### [/etc/opt/emcpower/EMCpower.LINUX-6.5.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.6.5.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

- emcpx.ko

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

#### **/lib64**

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

#### **/sbin**

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

#### **/usr/lib/systemd/system**

RHEL7, SLES 12 SP4, SLES15, and OL7

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

#### **/usr/lib/PowerPath**

RHEL7, SLES 12 SP4, SLES15, and OL7

- PowerPath
- PowerPathPost
- emcp\_mond.sh

#### **/usr/share/man/man1**

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

SLES15, SLES 12 SP4, 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-6.5.0.00.00/90powerpath
- 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-6.5.x.00.00/90powerpath
- check
- emcp\_pre-trigger.sh
- emcp\_settled.sh
- install
- 64-emcp.rules
- module-setup.sh

Files changed by PowerPath

# APPENDIX A

## Retpoline Kernels

To mitigate against kernel or cross-process memory disclosure (the Spectre attack), recent Linux kernels are being compiled with a new option. The `-mindirect-branch=thunk-extern` introduced to `gcc` to perform indirect calls through a so-called retpoline. Red Hat, SuSE, and Oracle have all released errata kernels that are built with the retpoline compiler for some of their existing OS versions. All new OS versions in the future will be released with retpoline kernels only.

For more information about retpoline compiler and kernels, see:

- <https://access.redhat.com/articles/3311301>
- <https://www.suse.com/support/kb/doc/?id=7022512>
- <https://blogs.oracle.com/linux/an-update-on-retpoline-enabled-kernels-for-oracle-linux-v2>
- [PowerPath for Linux 6.5 and retpoline](#)..... 118
- [PowerPath for Linux 6.5 and non-retpoline \(older\) kernels](#)..... 118

## PowerPath for Linux 6.5 and retpoline

PowerPath for Linux 6.5 can be installed only on retpoline kernels. It does not work on older kernels that have been built with a pre-retpoline compiler. It is important to check the PowerPath support matrix before installing or upgrading to PowerPath for Linux 6.5 version.

To verify whether a kernel has been built with a retpoline compiler, check if the running kernel has functions with the `_x86_indirect_thunk_` prefix. These are the functions used to implement retpoline itself. Check if those functions are present in the running kernel. A snippet of the actual output is copied below:

```
# cat /proc/kallsyms | grep __x86_indirect_thunk_
ffffffff81a04000 T __x86_indirect_thunk_rax
ffffffff81a04020 T __x86_indirect_thunk_rbx
ffffffff81a04040 T __x86_indirect_thunk_rcx
ffffffff81a04060 T __x86_indirect_thunk_rdx
ffffffff81a04080 T __x86_indirect_thunk_rsi
ffffffff81a040a0 T __x86_indirect_thunk_rdi
ffffffff81a040c0 T __x86_indirect_thunk_rbp
ffffffff81a040e0 T __x86_indirect_thunk_r8
```

**Note:** For more information about retpoline kernels, refer to the relevant operating system vendor's documentation. All queries specific to retpoline kernels should be directed to the operating system vendor.

## PowerPath for Linux 6.5 and non-retpoline (older) kernels

This section provides information about what is expected if PowerPath for Linux 6.5 is installed on a non-retpoline kernel.

### Fresh install

Installing PowerPath 6.5 fails on non-retpoline kernels with the following error:

```
[root@hostname ~]# rpm -ivh DellEMCPower.LINUX-6.5.0.00.00-067.RHEL6.x86_64.rpm
Preparing...##### [100%]

Error: This version of PowerPath will work only on retpoline enabled kernels.

error: %pre(EMCpower.LINUX-6.5.0.00.00-067.el6.x86_64) scriptlet failed, exit
status 1
error: install: %pre scriptlet failed (2), skipping
EMCpower.LINUX-6.5.0.00.00-067.el6
[root@hostname ~]#
```

### Upgrading PowerPath

Upgrading an older version of PowerPath to version 6.5 is allowed on an older non-retpoline kernel to enable customers who upgrade their kernel to a retpoline kernel along with a PowerPath upgrade. This condition is necessary because of the following scenario:

If only the kernel is upgraded to a retpoline kernel while an older version of PowerPath is installed and after the host is rebooted, IBRS mitigation is invoked because pre-6.5 PowerPath modules are not built with a retpoline compiler. Even if PowerPath is upgraded to version 6.5, the IBRS mitigation is active until a host reboot. To avoid the second reboot, upgrading PowerPath to 6.5 is allowed on a non-retpoline kernel, but the host must be upgraded to a retpoline kernel before PowerPath can be started.

The following warning message is displayed when PowerPath is upgraded from an older version to version 6.5 on a non-retpoline kernel:

```
[root@hostname ~]# rpm -Uvh DellEMCPower.LINUX-6.5.0.00.00-067.RHEL6.x86_64.rpm
Preparing...##### [100%]
1:EMCpower.LINUX##### [100%]
```

All trademarks used herein are the property of their respective owners.  
 NOTE: License registration is not required to manage the CLARiON AX series array.

\*\*\* IMPORTANT \*\*\*

Please check the following configurations before starting PowerPath:

- Add `_netdev` to `/etc/fstab` mount options for PowerPath pseudo devices.
- Ensure `netfs` service is started.
  - `netfs` service is needed to mount devices with `_netdev` option.
- Set LVM filter in `/etc/lvm/lvm.conf` according to PowerPath recommendation.
- Blacklist all devices in `/etc/multipath.conf` and stop `multipathd` service.
- Install PowerPath license(s) and ensure that policy is not set to

BasicFailover.

- If no license is available, ensure that only one HBA port is active in the host.

PowerPath supports only single-HBA configuration when unlicensed.

Warning: This version of PowerPath will work only on retpoline enabled kernels.

Refer to PowerPath Installation and Administration Guide for details.

If PowerPath is upgraded from an earlier version to 6.5 on a non-retpoline kernel, and then after the upgrade is completed, it should not be started as it fails to start. The kernel must be upgraded to a retpoline compiler and then the host must be rebooted. Once the host reboots, PowerPath 6.5 starts normally on the new retpoline kernel.

### Starting PowerPath after upgrade on a non-retpoline kernel

If PowerPath for Linux 6.5 is already installed, after an upgrade from an earlier version, on a non-retpoline kernel and if an attempt is made to start it, it fails with the following error:

```
[root@hostname ~]# /etc/init.d/PowerPath start
Starting PowerPath:
Error: This version of PowerPath will work only on retpoline enabled kernels.
failed
```





# APPENDIX B

## Dynamic Linking of OpenSSL Libraries

This section provides information about the linking of OpenSSL Libraries in PowerPath.

- [Overview](#) ..... 122
- [Changes](#) ..... 122
- [Expected behavior](#) ..... 122
- [PowerPath for Linux and SLES 15](#) ..... 123

## Overview

OpenSSL libraries are used by PowerPath in the remote management feature. When Dell EMC PowerPath Management Appliance (PPMA) host connects to a PowerPath host, the socket calls are secured by the OpenSSL library.

In earlier versions of PowerPath for Linux, the PowerPath host agent daemon, *emcp\_mond*, was statically linked to OpenSSL libraries. It did not matter if a compatible version of OpenSSL is installed on the host or not.


From PowerPath for Linux 6.5 version, the host agent daemon, *emcp\_mond* is dynamically linked to OpenSSL libraries. It is necessary that a compatible version of OpenSSL library is installed on the host in order for the remote management feature to work properly.

## Changes

From PowerPath for Linux 6.5, *emcp\_mond* is dynamically linked to OpenSSL libraries. And *emcp\_mond* uses the version of OpenSSL libs that are installed on host. Even if multiple versions of OpenSSL libraries are installed, *emcp\_mond* looks for and uses only the latest 1.0.x version of the libs.

There may be rare instances when OpenSSL library is not installed or an incompatible version is installed. In such cases, if *emcp\_mond* does not start, other features dependent on *emcp\_mond* such as, Auto Host Registration, is impacted.

To address this scenario, two versions of *emcp\_mond* binaries are included in the PowerPath package. The first version, *emcp\_mond* is used when a compatible version of OpenSSL libs is found to be installed on the host. The second version, *emcp\_mond\_no\_ssl* is used when OpenSSL library is not installed or an incompatible OpenSSL version is found to be installed.

 **Note:** PowerPath 6.5 supports any 1.0.x version of OpenSSL library only.

## Expected behavior

This section provides information on the expected behavior of PowerPath with or without installing OpenSSL.

### With OpenSSL 1.0.x

There is no change in how PowerPath works if a compatible version of OpenSSL is available on a host. `ps -ef` output should show `/sbin/emcp_mond` running when PowerPath is active.

### With no OpenSSL library or incompatible version

This section provides information on how PowerPath behaves when no OpenSSL library or an incompatible version is installed.

`ps -ef` output should show `/sbin/emcp_mond_no_ssl` running when PowerPath is active.

### PowerPath installed on host without OpenSSL

This section provides information about the warning that PowerPath generates when the PowerPath package is installed on a host which does not have OpenSSL library that is installed.

PowerPath generates the following warning:

```
[root@hostname ~]# rpm -ivh DellEMCPower.LINUX-6.5.0.00-067.RHEL7.x86_64.rpm
Preparing... ##### [100%]
Updating / installing...
1:EMCpower.LINUX-6.5.0.00-067.el##### [100%]
All trademarks used herein are the property of their respective owners.
NOTE:License registration is not required to manage the CLARiiON AX series
array.

*** IMPORTANT ***
Please check the following configurations before starting PowerPath:
  - Set LVM global_filter in /etc/lvm/lvm.conf according to PowerPath
  recommendation.
  - Blacklist all devices in /etc/multipath.conf and stop multipathd service.
  - Install PowerPath license(s) and ensure that policy is not set to
  BasicFailover.
  - If no license is available, ensure that only one HBA port is active in
  the host.
  PowerPath supports only single-HBA configuration when unlicensed.
Warning: Install OpenSSL library to use PowerPath remote management services.
Refer to PowerPath Installation and Administration Guide for details.
```

## PowerPath installed on host with incompatible OpenSSL library

This section provides information about the warning that PowerPath generates when the PowerPath package is installed on a host which has an incompatible OpenSSL library installed.

PowerPath generates the following warning.

```
[root@hostname ~]# rpm -ivh DellEMCPower.LINUX-6.5.0.00-067.RHEL7.x86_64.rpm
Preparing... ##### [100%]
Updating / installing...
1:EMCpower.LINUX-6.5.0.00-067.el##### [100%]
All trademarks used herein are the property of their respective owners.
NOTE:License registration is not required to manage the CLARiiON AX series
array.

*** IMPORTANT ***
Please check the following configurations before starting PowerPath:
  - Set LVM global_filter in /etc/lvm/lvm.conf according to PowerPath
  recommendation.
  - Blacklist all devices in /etc/multipath.conf and stop multipathd service.
  - Install PowerPath license(s) and ensure that policy is not set to
  BasicFailover.
  - If no license is available, ensure that only one HBA port is active in
  the host.
  PowerPath supports only single-HBA configuration when unlicensed.
Warning: Install supported version of OpenSSL library to use PowerPath remote
management services.
```

## PowerPath installed on host with incompatible or no OpenSSL library

This section provides information on the warning that `emcp_mond_edit` tool generates when the PowerPath package is installed on a host which has an incompatible or no OpenSSL library installed.

`emcp_mond_edit` tool generates the following warning.

```
WARNING:Please install OpenSSL libraries to enable remote management service
```

## PowerPath for Linux and SLES 15

This section provides information about PowerPath for Linux and SLES 15.

On the latest SLES 15 (SuSE Linux Enterprise Server 15) operating system, the default version of OpenSSL installed is version 1.1. This version of OpenSSL is not compatible with PowerPath for Linux 6.5 version. If only the 1.1 version of OpenSSL is installed, the remote management feature does not work.

To use the remote management feature, 1.0.x version of OpenSSL must be installed on the host. The rpm for 1.0.x version of OpenSSL is available in the operating system ISOs of SLES15 as legacy packages. This rpm must be installed on the SLES 15 host before enabling remote management feature using the `emcp_mond_edit` tool.

# APPENDIX C

## Disable native multipathing

This section provides procedure for disabling dm-multipath on an RHEL 6.x or OL 6.x host and RHEL 7.x or OL 7.x host.

- [Disabling dm-multipath on an RHEL 6.x or OL 6.x host](#)..... 126
- [Disabling dm-multipath on an RHEL 7.x or OL 7.x host](#)..... 126

## Disabling dm-multipath on an 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 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.

```
[root@hostname ~]# chkconfig multipathd off
[root@hostname ~]# chkconfig --list multipathd
```

Service should be off on all run levels.

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

```
[root@hostname ~]# 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 that are configured.

```
[root@hostname ~]# multipath -ll
```

Should not return any dm devices.

## Disabling dm-multipath on an RHEL 7.x or OL 7.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.

```
[root@hostname ~]# systemctl disable multipathd
```

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

```
[root@hostname ~]# 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. Update the new entry in the `grub.cfg` file and make it as default.

 **Note:** The system boots with the new ramdisk image automatically.

For RHEL 7.x/OL 7.x:

```
[root@hostname ~]# grub2-mkconfig -o /boot/grub2/grub.cfg
[root@hostname ~]# grub2-set-default 'PowerPath - Red Hat Enterprise Linux Server'
```

6. Restart the host with the new ramdisk image and ensure that the dm-multipath does not have any configured devices.

```
[root@hostname ~]# multipath -ll
```

Should not return any dm devices.

Disable native multipathing



# APPENDIX D

## Booting with duplicate PVs on SLES BFS hosts

This section provides the background information about how SLES hosts handle logical volumes when duplicate physical volumes (PVs) are found.

- [SLES hosts and LVM](#)..... 130
- [SLES BFS hosts and duplicate PVs](#)..... 130
- [Configure LVM to accept duplicate PVs](#)..... 131

## SLES hosts and LVM

On SLES12 SP4 and SLES15, if multiple physical volumes are found for a logical volume, such volumes are neither activated automatically nor can they be activated manually. This is to avoid data loss in scenarios where the physical volumes (PV) may belong to different underlying storage.

In such scenarios, a message similar to the following is displayed:

```
[root@hostname~] # vgchange -ay ppvg
WARNING: Not using lvmetad because duplicate PVs were found.
WARNING: Use multipath or vgimportclone to resolve duplicate PVs?
WARNING: After duplicates are resolved, run "pvscan --cache" to enable
lvmetad.
/dev/sdb: open failed: No medium found
/dev/sdc: open failed: No medium found
WARNING: Not using device /dev/sdd2 for PV USHeZ8-rhWS-AuOY-LYL1-AYHZ-Aew4-
xzPykl.
WARNING: Not using device /dev/sdg2 for PV USHeZ8-rhWS-AuOY-LYL1-AYHZ-Aew4-
xzPykl.
WARNING: Not using device /dev/sdf for PV 5MwDqG-WUAO-ZJT3-B0wB-CHnB-
e2w8-9UxVry.
WARNING: Not using device /dev/sdh for PV 5MwDqG-WUAO-ZJT3-B0wB-CHnB-
e2w8-9UxVry.
WARNING: PV USHeZ8-rhWS-AuOY-LYL1-AYHZ-Aew4-xzPykl prefers device /dev/sda2
because device is used by LV.
WARNING: PV USHeZ8-rhWS-AuOY-LYL1-AYHZ-Aew4-xzPykl prefers device /dev/sda2
because device is used by LV.
WARNING: PV 5MwDqG-WUAO-ZJT3-B0wB-CHnB-e2w8-9UxVry prefers device /dev/sde
because device is used by LV.
WARNING: PV 5MwDqG-WUAO-ZJT3-B0wB-CHnB-e2w8-9UxVry prefers device /dev/sde
because device is used by LV.
Cannot activate LVs in VG ppvg while PVs appear on duplicate devices.
o logical volume(s) in volume group "ppvg" now active
```

This behavior is caused by the following parameter in the `devices{}` section of the `/etc/lvm/lvm.conf` file on SLES15 hosts:

```
# Configuration option devices/allow_changes_with_duplicate_pvs.
# Allow VG modification while a PV appears on multiple devices.
# When a PV appears on multiple devices, LVM attempts to choose the
# best device to use for the PV. If the devices represent the same
# underlying storage, the choice has minimal consequence. If the
# devices represent different underlying storage, the wrong choice
# can result in data loss if the VG is modified. Disabling this
# setting is the safest option because it prevents modifying a VG
# or activating LVs in it while a PV appears on multiple devices.
# Enabling this setting allows the VG to be used as usual even with
# uncertain devices.
allow_changes_with_duplicate_pvs = 0
```

**Note:** Although this parameter is not present by default on SLES12 SP4, it is still active. This is an operating system/LVM feature.

## SLES BFS hosts and duplicate PVs

In case of SLES BFS hosts with root FS logical volume, it is essential that only one physical volume is presented at boot. Otherwise, the host may not boot as the root logical volume and the swap logical volume cannot be activated.

If PowerPath is configured properly, this is not an issue as the LVM filter in `/etc/lvm/lvm.conf` is modified to accept only emcpowerX pseudo devices and not the sd native devices.

In case PowerPath is going to be uninstalled or upgraded, as part of the process, the host needs to be rebooted and the LVM filter is changed to accept all devices. If there are multiple paths to the boot LUN, the host does not boot when the host comes up after reboot. Therefore, follow [Configure LVM to accept duplicate PVs](#) on page 131 before PowerPath is uninstalled or upgraded.

## Configure LVM to accept duplicate PVs

### About this task

In the `/etc/lvm/lvm.conf` file, in the `devices {}` section, change the value of the `allow_changes_with_duplicate_pvs` to 1 instead of 0. If the parameter is not present, as in the case of SLES12 SP4, create it and assign 1 as its value. This enables logical volumes to be activated even when there are multiple PVs present.

```
allow_changes_with_duplicate_pvs = 1
```

**Note:** It is important to configure LVM to accept duplicate PVs before performing steps to upgrade or uninstall PowerPath. As part of the upgrade or uninstall process, the ram disk image is updated and this parameter is automatically included in the image.

## Troubleshooting

### About this task

In case the `allow_changes_with_duplicate_pvs` parameter was not changed to 1 before upgrade or uninstall of PowerPath, the host may fail to boot if multiple paths are available to the boot LUN. In such scenarios, remove all the paths except one path and try booting again.



# INDEX

## A

Auto host registration 63

Auto Host Registration Settings 63

## D

Device in use 65

Device in Use 65

Device IO statistics 65

## I

install power path 118

Installation 122, 123

## L

Linux 123

## M

Mode of operation 62

## N

non retplaine 118

## O

OpenSSL 122, 123

## P

powerpath 62, 63

powerPath 123

Powerpath 62

PowerPath 122, 123

## R

retpline 117

Retpline 118

## S

SLES15 123

