Dell EMC PowerPath for Linux

Version 7.0

Installation and Administration Guide

302-005-912

04

March 2020



Copyright © 1998-2020 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

CONTENTS

Tables		7
	PREFACE	9
Chapter 1	Getting started	13
	Getting started	14
Chapter 2	Installing PowerPath	15
-	Before you install PowerPath	16
	Review compatibility with the environment	16
	Review host system compatibility	16
	Review PowerPath licensing	17
	Review third-party path management software	17
	Review the installation language	
	Install the RPM file	
	Start PowerPath after the installation	18
	Post-installation steps	
	Install PowerPath when Veritas Storage Foundation is already installed	. 20
Chapter 3	PowerPath and NVMe	21
	NVMe Overview	
	Operating Systems Supported with NVMe	. 22
	Install PowerPath in NVMe	. 22
	Configure PowerPath for NVMe	. 23
	Add NVMe namespaces dynamically	23
	Remove NVMe namespaces dynamically	23
	PowerPath limitations for NVMe	24
Chapter 4	Administering PowerPath in a Boot from SAN environment	25
-	Install and configure PowerPath in a Boot from SAN host	. 26
	Install Operating System on a Boot from SAN host	
	Configure PowerPath on a Boot from SAN host	. 28
	Non-Disruptive Upgrade of PowerPath in a Boot from SAN host	
	NDU upgrade if PowerPath 6.5 was configured using the	
	emcpbfs_linux script	31
	NDU upgrade if PowerPath 6.5 was configured using manual steps	31
	Upgrade PowerPath in a Boot from SAN host	. 36
	Upgrade PowerPath on the Boot from SAN host using the	
	emcpbfs_linux script	36
	Upgrade PowerPath on the Boot from SAN host using manual steps.	37
	Upgrade the Linux kernel in a Boot from SAN host	
	Uninstalling PowerPath from a Boot from SAN host	41
	Uninstall PowerPath from a Boot from SAN host using	
	emcpbfs_linux script	
	Uninstall PowerPath from a Boot from SAN host using manual steps	42

Chapter 5	Managing paths to a logical device	45
	Dynamic LUN addition and removal	
	Adding a LUN	
	Removing a LUN	
	Resizing a LUN	
	Before you add or remove paths	
	Adding new paths to a PowerPath logical device	
	Removing some of the paths to PowerPath logical devices	49
Chapter 6	Maintaining PowerPath	53
	/etc/fstab options	54
	Configure remote management service	
	Configure event monitoring with SNMP traps	
	Monitored events	
	Configure SNMP traps	
	Changing the language of installation messages	
	Saving new configurations	
	Configuring LVM2 support	
	Root file system not mounted on a logical volume	
	Root file system mounted on a logical volume	
	Root file system mounted on a logical volume residing on a local di	
	managed by PowerPath	
	Configuring PowerPath on cluster nodes	
	Hot swapping an HBA	
	Removing an HBA	
	Adding an HBA	
	Configuring for R1/R2 boot failover support	67
	R1/R2 supported configurations	67
	R1/R2 boot procedure	67
	Configure the network interface on the R1 host	68
	Enable scsi name on NetApp arrays	
	Removing devices from PowerPath control	
	Expand the file system on a pseudo-device	
	Supporting more than 256 storage devices	
Chapter 7	Upgrading PowerPath	71
Chapter /		
	Preparing for upgrade	
	Upgrading PowerPath	
	Upgrade PowerPath from a compressed archive	
	Upgrade PowerPath across Linux platforms	
	Upgrade the Linux kernel in a PowerPath installed host	
	Upgrade PowerPath after upgrading from one major Linux release	
	the next	
	Upgrade PowerPath before upgrading to the latest SLES Service 75	Pack
	Upgrade PowerPath before upgrading RHEL update on Linux	76
	Upgrade PowerPath in a SteelEye LifeKeeper cluster environment	
	Non-disruptive upgrade	
	Supported Operating Systems	
	Steps to perform PowerPath NDU	
		. -
Chapter 8	Removing PowerPath Before you remove PowerPath	81 82
		02

	Uninstall PowerPath	82
	Uninstall PowerPath	
	Saved configuration files	
	Remove the saved configuration files	
	_	
Chapter 9	PowerPath features for PowerMax array	85
	PowerPath Host Registration	
	Overview	
	Initial Registration and Storage Configuration	
	Modes of operation	
	Settings	
	Log messages	87
	FX Licensing	88
	PowerPath Device IO Statistics	
	Overview	89
	Settings	
	PowerPath Device in Use Reporting	
	Overview	
	Settings	
	Proactive Port Offline Autostandby	
	Proactive Port Offline Autostandby Overview	
	Flaky Path Reporting to PowerMax	
	Flaky Path Reporting to PowerMax Overview	
	SAN latency reportingSAN latency reporting overview	
	View PowerPath SAN Latency in PowerPath Management Applian 92	ice
Chapter 10	Coexistence with other path management software	93
	PowerPath and third-party path management software	
	PowerPath co-management behavior	94
Chapter 11	Troubleshooting PowerPath	97
51.up 151 11	Recover from a failed PowerPath installation	
	PowerPath custom settings not preserved	
	Recover from a failed PowerPath Non-Disruptive Upgrade due to a panic of	
	reboot	99
	Resolve missing or deleted files	100
	Problems booting the host	101
Chapter 12	Files changed by PowerPath	103
-	Linux files modified by PowerPath installation	
	Files created by PowerPath installation	
Appendix A	Retpoline Kernels	113
. ippolition A	PowerPath for Linux 7.0 and retpoline	
	PowerPath for Linux 7.0 and non-retpoline (older) kernels	
Appendix B		
	Dunamic Linking of OpenSSI Libraries	117
Appendix B	Dynamic Linking of OpenSSL Libraries	117
Appendix B	Dynamic Linking of OpenSSL Libraries Overview	118

	Expected behavior	118
	With a compatible version of OpenSSL library	118
	With no OpenSSL library or incompatible version	118
Appendix C	Disable native multipathing	121
	Disabling dm-multipath on a RHEL 6.x or OL 6.x host	122
	Disabling dm-multipath on a RHEL 7.x or OL 7.x host	
	Disabling dm-multipath on a RHEL8.x, SLES12 SP4, SLES15, or SLES15 host	
Appendix D	Booting with duplicate PVs on SLES and RHEL8 BFS hosts	125
	RHEL8, SLES hosts, and LVM	
	RHEL8, SLES BFS hosts, and duplicate PVs	
	Configure LVM to accept duplicate PVs	
	Troubleshooting	127
	Recreate new ram disk image	127
Appendix E	Disable NVMe multipath	129
l. l	Disabling NVMe multipath on SLES 15 and SLES 15 SP1 hosts	
Index		131

TABLES

1	Revision history	9
2	Getting started tasks	14
3	Monitored events	56
4	Supported languages	60
5	Linux Kernel Filter Fields	62
6	Kernel Filter	63
7	Operating System Support Steps	70
8	Supported Operating Systems	78
9	Software with disabled PowerPath support	94
10	Present or enabled software and multipathing behavior	94
11	Commands to disable support	

Tables

8

PREFACE

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

Contact your Dell EMC technical support professional if a product does not function properly or does not function as described in this document.

Note: This document was accurate at publication time. Go to Dell EMC Online Support to ensure that you are using the latest version of this document.

Purpose

This guide is part of the PowerPath documentation set, and is intended for use by a system or storage administrator during installation, configuration, and administration of the product.

Audience

Readers of this guide are expected to be familiar with:

- Storage systems in their environment
- · Linux operating system
- Applications (for example, clustering software) used with PowerPath

Related documentation

The following Dell EMC publications provide additional information:

- PowerPath Family- or Linux Release Notes
- PowerPath Family CLI and System Messages Reference Guide
- PowerPath Family Product Guide
- PowerPath Management Appliance Installation and Configuration Guide
- PowerPath Migration Enabler User Guide

Table 1 Revision history

Revision	Date	Description
04	March 2020	Updated the following sections:
		 Review compatibility with the environment
		Review host system compatibility
		Disabling dm-multipath on a RHEL8.x, SLES12 SP4, SLES15, or SLES15 SP1 host
03	January 2020	Updated the following sections:
		 Flaky Path Reporting to PowerMax Overview
		SAN latency reporting overview

Table 1 Revision history (continued)

Revision	Date	Description
02	September 2019	Added content for Foxtail and NVMe features PowerPath for Linux 7.0 Patch 01 release
01	July 2019	First release of the product

Special notice conventions used in this document

Dell EMC uses the following conventions for special notices:

- DANGER Indicates a hazardous situation which, if not avoided, will result in death or serious injury.
- WARNING Indicates a hazardous situation which, if not avoided, could result in death or serious injury.
- CAUTION Indicates a hazardous situation which, if not avoided, could result in minor or moderate injury.
- (i) NOTICE Addresses practices not related to personal injury.
- (i) Note: Presents information that is important, but not hazard-related.

Typographical conventions

Dell EMC uses the following type style conventions in this document:

Dell Livio does the following type style conventions in this document.			
Bold	Used for names of interface elements, such as names of windows, dialog boxes, buttons, fields, tab names, key names, and menu paths (what the user specifically selects or clicks)		
Italic	Used for full titles of publications referenced in text		
Monospace	Used for:		
	System code		
	System output, such as an error message or script		
	Pathnames, filenames, prompts, and syntax		
	Commands and options		
Monospace italic	Used for variables.		
Monospace bold	Used for user input.		
[]	Square brackets enclose optional values.		
1	Vertical bar indicates alternate selections - the bar means "or".		
{}	Braces enclose content that the user must specify, such as x or y or		

Ellipses indicate nonessential information that is omitted from the

example.

Where to get help

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

Product information

For documentation, release notes, software updates, or information about Dell EMC products, go to $\frac{1}{2}$ Dell EMC Online Support .

Technical support

Go to Dell EMC Online Support and click Services. You will see several options for contacting Dell EMC Technical Support. You must have a valid support agreement to open a service request. Contact your Dell EMC sales representative for details about obtaining a valid support agreement or with questions about your account.

PREFACE

CHAPTER 1

Getting started

Refer to related sections for procedures to install and administer PowerPath on a Linux serv	er fo
Dell EMC and third-party arrays.	

Getting started

This topic provides getting started information.

Table 2 Getting started tasks

Do you want to	Look at
install PowerPath?	Installing PowerPath on page 15 for quick installation with a compressed archive from Dell EMC Online Support.
administer PowerPath?	Administering PowerPath in a Boot from SAN environment on page 25 for administering PowerPath in a BFS environment, for configuring PowerPath logical device, managing PowerPath, or removing PowerPath from the host.
upgrade PowerPath?	Upgrading PowerPath on page 71 for prerequisites and procedures to upgrade PowerPath.
troubleshoot PowerPath?	Troubleshooting PowerPath on page 97 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 that is downloaded from Dell EMC Online Support.

After installation, license the PowerPath and configure it for operation. The $/var/log/powerpath_install.log$ file contains the details of all the operations that are performed during install, uninstall, and upgrade of PowerPath.

Note: The installation process is the same for major releases, service packs, patches, and hotfixes.

•	Before you install PowerPath	16
•	Install the RPM file	18
	Start PowerPath after the installation	
•	Post-installation steps	19
	Install PowerPath when Veritas Storage Foundation is already installed	

Before you install PowerPath

Before installing, understand the expected environment requirements as per the https://elabnavigator.emc.com/eln/elnhome 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.

Review compatibility with the environment

Review the following compatibility information with the environment:

Procedure

- 1. Review connectivity requirements as per the Dell EMC Host Connectivity Guide for Linux.
- 2. Ensure that environment meets the PowerPath requirements according to the *PowerPath Family for Linux Release Notes* and the E-lab Interoperability Navigator available at https://elabnavigator.emc.com/eln/elnhome.
- 3. Use EMC Grab utilities, Serviceability feature of PowerPath Management Appliance, and EMC Reporting tool and verify that your environment meets the requirements.
- Ensure that HBA BIOS and fibre drivers are configured as per the https:// elabnavigator.emc.com/eln/elnhome. Improper settings could cause unwanted I/O delays.

Review host system compatibility

Review the following host system compatibility for PowerPath:

Before you begin

(i) Note: PowerPath for Linux is not compatible to run in a virtual machine.

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 if there is inadequate space. Recover from a failed PowerPath installation on page 98 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 PowerPath in a Boot from SAN environment on page 25.
- 4. PowerPath for Linux 7.0 works only with retpoline enabled kernels. See Retpoline Kernels on page 113.
- 5. For hosts with NVMe devices, see PowerPath and NVMe on page 21 for pre-requisities.
- 6. Ensure any version of EMC Unisphere and Navisphere Application Transparent Failover (ATF) is removed from the host by EMC Professional Services.
- 7. Native Linux DM-multipathing (DM-MPIO) is not compatible with PowerPath, and installing both on the same host causes system instability.
 - See Disable native multipathing on page 121 for instructions on how to disable native multipathing.

- When installing on SLES 12 SP4, SLES 15, or SLES 15 SP1 ensure that native multipathing is not automatically activated.
- 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.
- If the remote management feature is enabled, PowerPath needs a compatible version of OpenSSL library to be installed on the host. See <u>Dynamic Linking of OpenSSL Libraries</u> on page 117.

Review PowerPath licensing

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 88.

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

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

About this task

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:

- 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 60.

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, 7.0.0.00.00.
- < 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.

Start PowerPath after the installation

PowerPath is not active immediately after the installation. It needs to be started manually or it automatically starts with every host reboot.

Procedure

1. Register the PowerPath license.

```
[root@hostname ~]# emcpreg -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.
 - (i) 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, RHEL8, OL7, SLES 12 SP4, SLES 15, or SLES 15 SP1

```
[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 LVM to use PowerPath pseudo devices only.

See Configuring LVM2 support on page 61.

2. 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 54.

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

For steps on configuring, see Configure SNMP traps on page 56.

4. 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
```

5. Configure only /dev/emcpowerX pseudo block devices that PowerPath created for the applications to use.

Note: Do not use /dev/sdX or native block devices for application IOs. Failover and load-balancing feature are not available for IOs sent directly to native block devices.

Install PowerPath when Veritas Storage Foundation 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, exclude Dell EMC arrays from DMP path management before and after the reconfiguration.

Procedure

- Log in as root.
- 2. Exclude Dell EMC arrays from DMP path management.
 - For SC devices, vxddladm excludearray libname=libvxSc.so
 - 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 vxdctl enable.

Verify that preceding DellEMC arrays are excluded from DMP path management. Run $vxdisk\ list$. The list of devices under VSF is displayed.

- 4. Install PowerPath, start it, and perform the post-installation steps as as explained in Post-installation steps on page 19.
- 5. Ensure that PowerPath claims the devices.

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

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

6. Run the command:

[root@hostname ~]# vxdisk scandisks

CHAPTER 3

PowerPath and NVMe

PowerPath for Linux 7.0 supports NVMe namespaces from PowerMax arrays running microcode 5978.444.444 and above through NVMeOF (NVMe Over Fabric) protocol.

•	NVMe Overview	22
	Operating Systems Supported with NVMe	
	Install PowerPath in NVMe	
	Configure PowerPath for NVMe	
	Add NVMe namespaces dynamically	
	Remove NVMe namespaces dynamically	
	PowerPath limitations for NVMe	

NVMe Overview

NVMe (Non-Volatile Memory Express) is an interface standard for accessing nonvolatile storage, commonly SSD disks.

NVMe supports higher speeds and has a lower latency than SATA. NVMe over Fabric is an architecture to access NVMe storage over networking fabrics such as NVMe over Fibre Channel (FC-NVMe).

Operating Systems Supported with NVMe

PowerPath for Linux 7.0 supports NVMe devices only on SLES 15 and SLES 15 SP1 operating systems. NVMe support on RHEL8 is available only with PowerPath for Linux 7.0 patch 1.

Note: Check the PowerPath support matrix for the exact kernel versions supported in PowerPath for Linux 7.0.

No extra steps are required for PowerPath installation and configuration on RHEL8. However, on SLES 15 and SLES 15 SP1, NVMe multipath is enabled by default and must be disabled manually before PowerPath is installed. If NVMe multipath is active, then PowerPath cannot manage the NVMe devices. See Disable NVMe multipath on page 129 for steps to disable NVMe multipath on SLES15 and SLES 15 SP1 hosts.

Besides as with SCSI devices, native multipathing (dm-multipath) should also be disabled before PowerPath is installed.

Install PowerPath in NVMe

Procedures to install and start PowerPath on hosts with NVMe devices are the same as on hosts with SCSI devices.

Before you begin

Refer Installing PowerPath on page 15 for steps to install and start PowerPath. Like SCSI devices, when PowerPath is installed and started, it automatically manages all the NVMe devices.

About this task

PowerPath for Linux 7.0 supports NVMe namespaces only from PowerMax arrays running microcode 5978.444.444 and above. These namespaces are listed under the PowerMax (powermax) class in powermt output.

```
[root@hostname ~]# powermt display class=powermax
PowerMax logical device count=10

----- Host Bus Adapters ------ I/O Paths ----- Stats -----
### HW Path Summary Total Dead IO/Sec Q-IOs Errors

0 lpfc optimal 40 0 - 0 0
3 lpfc optimal 40 0 - 0 0
```

(i) Note: SCSI devices from PowerMax arrays are still listed under the Symmetrix.

The default load-balancing policy with a valid license for PowerMax devices is Adaptive. The paths for NVMe devices are in the nvmeXnY format as seen below. Also, the PowerMax array ports are

listed with the prefix FN. Similar to SCSI, pseudo devices for NVMe devices are also created as emcpowerX block devices in /dev folder.

[root@hostname ~]# powermt display dev=all class=powermax Pseudo name=emcpowera PowerMax ID=000197600236 Logical device ID=0000073d Device WWN=600009700BC723EC3116005E000000D state=alive; policy=ADaptive; queued-IOs=0							
Host ### HW Path	I/O Paths	- Stor - Interf.	-, -		Stat Q-IOs E		
1 lpfc 3 lpfc 1 lpfc 3 lpfc	nvme0n1 nvme1n1 nvme2n1 nvme3n1	FN 2g:25 FN 4g:24	active active active active	alive alive alive alive	0 0 0 0	0 0 0 0	

Configure PowerPath for NVMe

This section lists some of the post installation steps that need to be performed for PowerPath to work properly with NVMe namespaces.

About this task

/etc/fstab options - There is no change in the options that need to be specified for filesystems on NVMe pseudo devices. See /etc/fstab options on page 54 for more details.

LVM filter - The LVM filter needs to exclude or reject all NVMe native devices when PowerPath is active. See Configuring LVM2 support on page 61 for the format of the LVM filter.

Add NVMe namespaces dynamically

NVMe namespaces can be added dynamically when PowerPath is already active.

About this task

After a new NVMe namespace is added to a host on the array, NVMe native block devices are automatically created by the kernel without any extra configuration steps. The new devices are also automatically configured in PowerPath. However, to complete the configuration and to create the pseudo device, the powermt config command needs to be run manually on the host.

Remove NVMe namespaces dynamically

NVMe namespaces can be removed dynamically when PowerPath is already active.

About this task

Before a NVMe namespace is removed from a host, the corresponding pseudo device should be removed from PowerPath control using the powermt remove dev=emcpowerX command. If the pseudo device is not in use, it is removed from PowerPath and then the namespace can be removed on the array.

PowerPath limitations for NVMe

PowerPath for Linux 7.0 does not support the following features with NVMe devices.

- Boot From SAN (BFS)
- · Cluster software
- PowerMax features (see PowerPath features for PowerMax array on page 85) except for FX licensing

CHAPTER 4

Administering PowerPath in a Boot from SAN environment

You can install and configure PowerPath in a Boot from SAN. It manages the boot LUN from an array like other LUNs attached to the host. As a boot LUN, it is always in use when the host is booted and hence more configuration steps need to be done on a Boot from SAN host when compared to a local boot host.

In PowerPath for Linux 7.0, enhancements have been made to support PowerPath in a Boot from SAN host. Procedures are simpler compared to previous versions. These updated procedures include not packaging PowerPath modules in the initial ramdisk image and not creating a separate PowerPath grub entry. Hence, carefully follow the steps that are listed in this chapter to install, upgrade, and uninstall PowerPath for Linux 7.0 in a Boot from SAN host.

Note: PowerPath is supported on Boot from SAN hosts only where the boot LUN is a SCSI device. It is not supported with a NVMe boot device.

•	Install and configure PowerPath in a Boot from SAN host	.26
	Non-Disruptive Upgrade of PowerPath in a Boot from SAN host	
	Upgrade PowerPath in a Boot from SAN host	
	Upgrade the Linux kernel in a Boot from SAN host	
	Uninstalling PowerPath from a Boot from SAN host	

Install and configure PowerPath in a Boot from SAN host

This section provides instructions about how to install and configure PowerPath in a Boot from SAN host. There are certain requirements for PowerPath to work properly on Boot from SAN (BFS) hosts. Install the operating system in the specified way and perform more steps to complete the configuration after installing the operating system and PowerPath. These steps are not required for hosts that boot from local disks.

Install Operating System on a Boot from SAN host

Procedure to install an operating system on a Boot from SAN host.

Procedure

- 1. Disable the disk controller that manages the local hard disks in the BIOS/UEFI configuration or remove all the local hard disks from the host.
 - Note: When host is powered up press F2 to enter the BIOS/UEFI setup to disable the disk controller.
- 2. Ensure that only one port of one of the FC HBAs, iSCSI HBAs, or FCoE adapters is connected to the storage network.
- 3. Ensure that only one LUN is presented to the host and only one path is available from the host to the array.

(i) 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 recommended 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.
- While installing on SLES 12 SP4, SLES 15, or SLES 15 SP1 ensure that native
 multipathing is not automatically activated. Because, there is no way to disable
 native multipathing after the operating system installation is complete.
- 4. Configure the boot LUN in the HBA BIOS by referring to the HBA vendor documentation.
- 5. Start the operating system installation.
- 6. Create a custom partition setup on the target storage device with the specified number of partitions depending on the configuration. Ensure that you select Basic Storage Devices as the type of device for installation where relevant.
 - For hosts with BIOS firmware and all supported operating systems:
 - Ensure that the partition layout type is MBR and not GPT.
 - A /boot partition that is at least 500 MB in size and formatted as an ext4 file system.
 - A second partition of type Linux LVM using the remaining space on the storage device.
 - For hosts with UEFI firmware and for SLES 12 SP4, SLES 15, and SLES 15 SP1 only:

- Ensure that the partition layout type is GPT and not MBR.
- A /boot/efi EFI system partition that is at least 500 MB in size and formatted as a vfat file system.
- A second partition of type Linux LVM using the remaining space on the storage device.
- For hosts with UEFI firmware and for OL 6, OL 7, RHEL 6, RHEL 7, and RHEL 8 only:
 - Ensure that the partition layout type is GPT and not MBR.
 - A /boot partition that is at least 500 MB in size and formatted as an ext4 file system.
 - A /boot/efi EFI system partition that is at least 500 MB in size and formatted as a vfat file system.
 - A third partition of type 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.
 - CAUTION The partitioning of the boot LUN and the creation of the logical volumes must be done precisely as specified in the preceding steps. Creation of any additional partitions or logical volumes on the boot LUN leads to the host not booting properly on reboot. It is also important that the type of file systems that are created on the boot partitions are exactly as specified.
- 9. Complete the operating system installation and reboot the host.
- 10. Ensure dm-multipath is not configured after the host is rebooted.

If multipath -ll returns configured devices, see Disable native multipathing on page 121.

11. If the operating system installation was completed with only a single path to the boot LUN, add additional paths to the boot LUN and any additional LUNs. Configure the boot LUN in the BIOS of the other HBA ports as well.

/|CAUTION

On SLES12 SP4, SLES15, SLES 15 SP1, and RHEL 8, LVM does not activate logical volumes if their physical volume (PV) appears on multiple block devices. In case of BFS hosts with the root file system on a logical volume, if LVM sees multiple PVs on boot, the host does not boot as the root logical volume and swap volume cannot be activated.

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. A new initial ramdisk image also needs to be created to include this setting so that the host can boot with the new image on the next reboot. Steps to change this parameter and update the initial ramdisk image are included in the steps to install, upgrade and uninstall PowerPath where relevant.

See Booting with duplicate PVs on SLES and RHEL8 BFS hosts on page 125 for more information.

Configure PowerPath on a Boot from SAN host

This section provides instructions to configure PowerPath on a Boot from SAN host.

Before you begin

Install and start PowerPath as explained in Installing PowerPath on page 15.

About this task

To complete the configuration of the boot LUN under PowerPath, use either one of the following methods:

- Configure PowerPath on the Boot from SAN host using the emcpbfs_linux script on page 28
- Configure PowerPath on the Boot from SAN host using the manual steps on page 28

Configure PowerPath on the Boot from SAN host using the emcpbfs linux script

Procedure to configure PowerPath on the Boot from SAN host using the <code>emcpbfs_linux</code> script.

Procedure

1. Run the following command to configure PowerPath on the Boot from SAN host:

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

Configure PowerPath on the Boot from SAN host using the manual steps

Procedure to configure PowerPath on the Boot from SAN host using the manual steps.

Procedure

1. Save a backup of the /etc/fstab file. Identify the emcpowerX pseudo device of the boot LUN.

In the following examples, it is assumed that emcpowera is the boot LUN. Then, modify /etc/fstab according to the configuration:

· For hosts with BIOS firmware and for OL 6 and RHEL 6 only:

Edit /etc/fstab to mount /boot from the first partition of the pseudo device (/dev/emcpowera1):

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

 For hosts with BIOS firmware and for OL 7, RHEL 7, RHEL 8, SLES 12 SP4, SLES 15, and SLES 15 SP1 only:

Edit /etc/fstab to mount /boot from the first partition of the pseudo device (/dev/emcpoweral):

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

For hosts with UEFI firmware and for SLES 12 SP4, SLES, 15 and SLES 15 SP1 only:

Edit /etc/fstab to mount /boot/efi from the first partition of the pseudo device (/dev/emcpoweral):

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

For hosts with UEFI firmware and for OL 6 and RHEL 6 only:

Edit /etc/fstab to mount /boot from the first partition of the pseudo device (/dev/emcpowera1) and /boot/efi from the second partition (/dev/emcpowera2):

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

#UUID=4AA3-84F3 /boot/efi vfat umask=0077,shortname=winnt 0 0
/dev/emcpowera2 /boot/efi vfat umask=0077,shortname=winnt 0 0
```

• For hosts with UEFI firmware and for OL 7, RHEL 7, and RHEL 8 only:

Edit /etc/fstab to mount /boot from the first partition of the pseudo device (/dev/emcpowera1) and /boot/efi from the second partition (/dev/emcpowera2):

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

#UUID=4AA3-84F3 /boot/efi vfat umask=0077,shortname=winnt 0 0
/dev/emcpowera2 /boot/efi vfat umask=0077,shortname=winnt 0 0
```

- 2. Remount the boot partitions so that it is mounted on pseudo device according to the configuration.
 - For hosts with BIOS firmware and all supported operating systems:

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

For hosts with UEFI firmware and for SLES 12 SP4, SLES 15, and SLES 15 SP1 only:

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

For hosts with UEFI firmware and for OL6, OL 7, RHEL 6, RHEL 7, and RHEL 8 only:

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

- 3. 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.
 - For OL6 and RHEL6:

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

· For all other operating systems:

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

- 4. Execute vgscan -v to rebuild the LVM2 cache.
- 5. Execute lvmdiskscan to ensure that the filter field is working correctly. Verify that the filtered device nodes are not listed in the output.
- 6. For RHEL 8, SLES 12 SP4, SLES 15, and SLES 15 SP1 hosts only, in /etc/lvm/lvm.conf file, enable the allow_changes_with_duplicate_pvs parameter if not already enabled as indicated as follows:

```
allow_changes_with_duplicate_pvs = 1
```

- (i) Note: If the above parameter is not already present, add it within the devices{} section.
- 7. Update initial ramdisk image only for RHEL 8, SLES 12 SP4, SLES 15, and SLES 15 SP1 hosts.
 - Note: This step is needed to ensure that the root and swap logical volumes are activated properly on host boot.
 - For RHEL 8: Save a backup of initial ramdisk image.

Recreate original initial ramdisk image.

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

 For SLES 12 SP4, SLES 15, and SLES 15 SP1: Save a backup of initial ramdisk image.

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

Recreate original initial ramdisk image.

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

Non-Disruptive Upgrade of PowerPath in a Boot from SAN host

This section provides instructions about Non-Disruptive Upgrade (NDU) of PowerPath in a Boot from SAN host.

About this task

PowerPath can be non-disruptively upgraded from PowerPath Linux 6.5 to 7.0 version without host reboot or application downtime. This NDU upgrade is supported only from PowerPath Linux 6.5.0.

Note: For all versions earlier than 6.5, see Upgrade PowerPath in a Boot from SAN host on page 36.

Depending on whether PowerPath for Linux 6.5.0 was configured with the <code>emcpbfs_linux</code> script or with manual steps, follow either of the following methods:

- NDU upgrade if PowerPath 6.5 was configured using the emcpbfs_linux script on page 31
- NDU upgrade if PowerPath 6.5 was configured using manual steps on page 31

NDU upgrade if PowerPath 6.5 was configured using the emcpbfs linux script

See Non-disruptive upgrade on page 77 for non-disruptively upgrading PowerPath from 6.5 version to 7.0 when ${\tt emcpbfs_linux}$ script was used to configure PowerPath after 6.5 installation. No additional steps are required.

About this task

Note: If the file /etc/emc/.emcpbfs_linux_done exists, then PowerPath 6.5 was configured using the emcpbfs linux script.

NDU upgrade if PowerPath 6.5 was configured using manual steps

Procedure for NDU upgrade if PowerPath 6.5 was configured using manual steps.

Procedure

- 1. Disable PowerPath from starting automatically on reboot.
 - For OL6 and RHEL6:
 Rename /etc/init.d/PowerPath to /etc/init.d/PowerPath.old.
 - For all other operating systems:

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

- 2. Restore /etc/fstab to mount boot partitions by UUID.
 - For hosts with BIOS firmware and for OL 6 and RHEL 6 only:

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

• For hosts with BIOS firmware and for OL 7, RHEL 7, SLES 12 SP4, and SLES 15 only:

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

• For hosts with UEFI firmware and for SLES 12 SP4 and SLES 15 only:

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

For hosts with UEFI firmware and for OL 6 and RHEL 6 only:

```
#/dev/emcpowera1 /boot ext4 _netdev 1 2
UUID=110223d8-f5f3-4058-812f-16d9e9ee7eab /boot ext4 defaults 1 2
#/dev/emcpowera2 /boot/efi vfat umask=0077,shortname=winnt 0 0
UUID=4AA3-84F3 /boot/efi vfat umask=0077,shortname=winnt 0 0
```

For hosts with UEFI firmware and for OL 7 and RHEL 7 only:

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

#/dev/emcpowera2 /boot/efi vfat umask=0077,shortname=winnt 0 0
UUID=4AA3-84F3 /boot/efi vfat umask=0077,shortname=winnt 0 0
```

- Remount the boot partitions so that it is mounted on the native device and not the pseudo device.
 - For hosts with BIOS firmware and all supported operating systems:

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

For hosts with UEFI firmware and for SLES 12 SP4 and SLES 15 only:

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

• For hosts with UEFI firmware and for OL6, OL 7, RHEL 6, and RHEL 7 only:

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

- 4. Edit the /etc/lvm/lvm.conf file to restore the filter parameter so that it accepts all devices.
 - For OL6 and RHEL6:

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

• For all other operating systems:

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

- 5. Execute vgscan -v to rebuild the LVM2 cache.
- 6. Execute lvmdiskscan to ensure that the filter field is working correctly. Verify all the block devices are listed in the output.
- 7. For SLES 12 SP4 and SLES 15 hosts only, in /etc/lvm/lvm.conf file, enable allow changes with duplicate pvs parameter, if not enabled earlier, as follows:

```
allow_changes_with_duplicate_pvs = 1
```

- Note: If the preceding parameter is not already present, add it within the devices{} section.
- 8. Recreate the ramdisk image and update the grub bootloader configuration.
 - For OL6 and RHEL6:
 - 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 a new ramdisk image file:

```
[root@hostname \sim]# dracut /boot/initramfs-\circ(uname -r).img \circ(uname -r)
```

- c. Add a new boot entry in the /boot/grub/menu.lst file for the newly created ramdisk image and make it the default.
- d. Remove previously added boot entry when PowerPath was installed and configured for BFS.
- For all other operating systems:
 - 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 a new ramdisk image file.
 - For OL 7 and RHEL 7:

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

■ For SLES 12 SP4 and SLES 15:

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

- c. Remove the entry that was previously added when PowerPath was installed and configured from the 40_custom file in the /etc/grub.d folder.
- d. Recreate the grub.cfg file to remove the PowerPath boot entry.
 - For hosts with BIOS firmware and for all supported operating systems:

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

• For hosts with UEFI firmware and for OL 7 and RHEL 7 only:

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

• For hosts with UEFI firmware and for SLES 12 SP4 and SLES 15 only:

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

- e. Get the relevant title by searching in grub.cfg file and set the default boot entry with grub2-set-default command.
 - For OL 7 and RHEL 7:

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

■ For SLES 12 SP4 and SLES 15:

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

- 9. Perform non-disruptive upgrade of PowerPath by following the steps in Non-disruptive upgrade on page 77.
- 10. Save a backup of the /etc/fstab file. Identify the emcpowerX pseudo device of the boot lun.

In the following examples, it is assumed that emcpowera is the boot lun. Modify /etc/ fstab according to the configuration.

• For hosts with BIOS firmware and for OL 6 and RHEL 6 only:

Edit /etc/fstab to mount /boot from the first partition of the pseudo device (/dev/emcpoweral):

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

For hosts with BIOS firmware and for OL 7, RHEL 7, SLES 12 SP4, and SLES 15 only:
 Edit /etc/fstab to mount /boot from the first partition of the pseudo device (/dev/emcpoweral):

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

For hosts with UEFI firmware and for SLES 12 SP4 and SLES 15 only:

Edit /etc/fstab to mount /boot/efi from the first partition of the pseudo device (/dev/emcpoweral):

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

For hosts with UEFI firmware and for OL 6 and RHEL 6 only:

Edit /etc/fstab to mount /boot from the first partition of the pseudo device (/dev/emcpowera1) and /boot/efi from the second partition (/dev/emcpowera2):

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

#UUID=4AA3-84F3 /boot/efi vfat umask=0077,shortname=winnt 0 0
/dev/emcpowera2 /boot/efi vfat umask=0077,shortname=winnt 0 0
```

• For hosts with UEFI firmware and for OL 7 and RHEL 7 only:

Edit /etc/fstab to mount /boot from the first partition of the pseudo device (/dev/emcpowera1) and /boot/efi from the second partition (/dev/emcpowera2):

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

#UUID=4AA3-84F3 /boot/efi vfat umask=0077,shortname=winnt 0 0
/dev/emcpowera2 /boot/efi vfat umask=0077,shortname=winnt 0 0
```

- 11. Remount the boot partitions so that it is mounted on pseudo device according to the configuration.
 - For hosts with BIOS firmware and all supported operating systems:

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

For hosts with UEFI firmware and for SLES 12 SP4 and SLES 15 only:

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

• For hosts with UEFI firmware and for OL6, OL 7, RHEL 6, and RHEL 7 only:

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

- 12. Save a backup of the /etc/lvm/lvm.conf file and change the LVM filter in /etc/lvm/lvm.conf to accept only pseudo-devices:
 - For OL6 and RHEL6:

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

· For all other operating systems:

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

- 13. Execute vgscan -v to rebuild the LVM2 cache.
- 14. Execute lvmdiskscan to ensure that the filter field is working correctly. Verify that the filtered device nodes are not listed in the output.

Upgrade PowerPath in a Boot from SAN host

This section provides instructions about how to upgrade PowerPath in a Boot from SAN host.

About this task

To upgrade PowerPath in a Boot from SAN host from versions 6.5 or earlier to 7.0, use one of the following methods:

- Upgrade PowerPath on the Boot from SAN host using the emcpbfs_linux script on page 36
- Upgrade PowerPath on the Boot from SAN host using manual steps on page 37
- (i) Note: These methods involve an application downtime as the host needs to be rebooted.

Upgrade PowerPath on the Boot from SAN host using the $emcpbfs_linux$ script

Procedure to upgrade PowerPath on the Boot from SAN host using the emcpbfs linux script.

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 <code>emcpbfs_linux</code> config has not been used during PowerPath configuration.

See Upgrade PowerPath on the Boot from SAN host using manual steps on page 37 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. For UEFI hosts, it also modifies /boot/efi entry.
- 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 72.
- 5. Execute /etc/opt/emcpower/emcpbfs linux config to configure PowerPath.

Upgrade PowerPath on the Boot from SAN host using manual steps

Procedure to upgrade PowerPath on the Boot from SAN host using the manual steps.

Procedure

- 1. Disable PowerPath from automatically starting on reboot.
 - For OL6 and RHEL6:

```
Rename /etc/init.d/PowerPath to /etc/init.d/PowerPath.old.
```

• For all other operating systems:

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

- 2. Restore /etc/fstab to mount boot partitions by UUID.
 - For hosts with BIOS firmware and for OL 6 and RHEL 6 only:

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

• For hosts with BIOS firmware and for OL 7, RHEL 7, SLES 12 SP4, and SLES 15 only:

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

For hosts with UEFI firmware and for SLES 12 SP4 and SLES 15 only:

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

For hosts with UEFI firmware and for OL 6 and RHEL 6 only:

```
#/dev/emcpowera1 /boot ext4 _netdev 1 2
UUID=110223d8-f5f3-4058-812f-16d9e9ee7eab /boot ext4 defaults 1 2
#/dev/emcpowera2 /boot/efi vfat umask=0077,shortname=winnt 0 0
UUID=4AA3-84F3 /boot/efi vfat umask=0077,shortname=winnt 0 0
```

• For hosts with UEFI firmware and for OL 7 and RHEL 7 only:

```
#/dev/emcpowera1 /boot ext4 nofail 1 2
UUID=110223d8-f5f3-4058-812f-16d9e9ee7eab /boot ext4 defaults 1 2
#/dev/emcpowera2 /boot/efi vfat umask=0077,shortname=winnt 0 0
UUID=4AA3-84F3 /boot/efi vfat umask=0077,shortname=winnt 0 0
```

3. Remount the boot partitions so that it is mounted on the native device and not the pseudo device.

• For hosts with BIOS firmware and all supported operating systems:

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

For hosts with UEFI firmware and for SLES 12 SP4 and SLES 15 only:

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

• For hosts with UEFI firmware and for OL6, OL7, RHEL 6, RHEL7, and RHEL 8 only:

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

- Edit the /etc/lvm/lvm.conf file to restore the filter parameter so that it accepts all devices.
 - For OL6 and RHEL6:

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

For all other operating systems:

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

- 5. Execute vgscan -v to rebuild the LVM2 cache.
- 6. Execute lvmdiskscan to ensure that the filter field is working correctly. Verify all the block devices are listed in the output.
- 7. For SLES 12 SP4 and SLES 15 hosts only, in /etc/lvm/lvm.conf file, enable allow_changes_with_duplicate_pvs parameter if not already enabled as indicated as follows:

```
allow_changes_with_duplicate_pvs = 1
```

- Note: If the preceding parameter is not already present, add it within the devices{} section.
- 8. Recreate the ramdisk image and update the grub bootloader configuration.
 - For OL6 and RHEL6:
 - 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 a new ramdisk image file:

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

- c. Add a new boot entry in the /boot/grub/menu.lst file for the newly created ramdisk image and make it the default.
- d. Remove previously added boot entry when PowerPath was installed and configured for BFS.
- · For all other operating systems:
 - 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 a new ramdisk image file.
 - For OL 7 and RHEL 7:

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

- c. Remove the entry that was previously added when PowerPath was installed and configured from the 40_custom file in the /etc/grub.d folder.
- d. Recreate the grub.cfg file to remove the PowerPath boot entry.
 - For hosts with BIOS firmware and for all supported operating systems:

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

For hosts with UEFI firmware and for OL 7 and RHEL 7 only:

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

• For hosts with UEFI firmware and for SLES 12 SP4 and SLES 15 only:

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

- e. Get the relevant title by searching in grub.cfg file and set the default boot entry with grub2-set-default command.
 - For OL 7 and RHEL 7:

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

■ For SLES 12 SP4 and SLES 15:

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

9. Reboot the host.

- 10. Upgrade and start PowerPath.
 - Note: To review the procedure for upgrading PowerPath, see Upgrading PowerPath on page 72.
- 11. Save a backup of the /etc/fstab file. Identify the emcpowerX pseudo device of the boot lun.

In the following examples, it is assumed that emcpowera is the boot lun. Modify /etc/ fstab according to the configuration.

• For hosts with BIOS firmware and for OL 6 and RHEL 6 only:

Edit /etc/fstab to mount /boot from the first partition of the pseudo device (/dev/emcpoweral):

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

For hosts with BIOS firmware and for OL 7, RHEL 7, SLES 12 SP4, and SLES 15 only:
 Edit /etc/fstab to mount /boot from the first partition of the pseudo device (/dev/emcpowera1):

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

· For hosts with UEFI firmware and for SLES 12 SP4 and SLES 15 only:

Edit /etc/fstab to mount /boot/efi from the first partition of the pseudo device (/dev/emcpoweral):

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

For hosts with UEFI firmware and for OL 6 and RHEL 6 only:

Edit /etc/fstab to mount /boot from the first partition of the pseudo device (/dev/emcpowera1) and /boot/efi from the second partition (/dev/emcpowera2):

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

#UUID=4AA3-84F3 /boot/efi vfat umask=0077,shortname=winnt 0 0
/dev/emcpowera2 /boot/efi vfat umask=0077,shortname=winnt 0 0
```

• For hosts with UEFI firmware and for OL 7 and RHEL 7 only:

Edit /etc/fstab to mount /boot from the first partition of the pseudo device (/dev/emcpowera1) and /boot/efi from the second partition (/dev/emcpowera2):

```
#UUID=110223d8-f5f3-4058-812f-16d9e9ee7eab /boot ext4 defaults 1 2 /dev/emcpoweral /boot ext4 nofail 1 2 #UUID=4AA3-84F3 /boot/efi vfat umask=0077,shortname=winnt 0 0 /dev/emcpowera2 /boot/efi vfat umask=0077,shortname=winnt 0 0
```

12. Remount the boot partitions so that it is mounted on pseudo device according to the configuration.

• For hosts with BIOS firmware and all supported operating systems:

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

For hosts with UEFI firmware and for SLES 12 SP4, SLES 15 and SLES 15 SP1 only:

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

• For hosts with UEFI firmware and for OL6, OL 7, RHEL 6, RHEL 7, and RHE only:

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

- 13. 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:
 - For OL6 and RHEL6:

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

For all other operating systems:

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

- 14. Execute vgscan -v to rebuild the LVM2 cache.
- 15. Execute lvmdiskscan to ensure that the filter field is working correctly. Verify that the filtered device nodes are not listed in the output.

Upgrade the Linux kernel in a Boot from SAN host

See the operating system vendor documentation to know how to upgrade the kernel. No extra steps are needed to reconfigure PowerPath.

Uninstalling PowerPath from a Boot from SAN host

This section provides instructions about uninstalling PowerPath from a Boot from SAN host.

About this task

To uninstall PowerPath from a Boot from SAN host, use one of the following methods:

- Uninstall PowerPath from a Boot from SAN host using emcpbfs_linux script on page 42
- Uninstall PowerPath from a Boot from SAN host using manual steps on page 42

Uninstall PowerPath from a Boot from SAN host using emcpbfs linux script

Procedure to uninstall PowerPath from a Boot from SAN host using emcpbfs linux script.

Before you begin

In a multiboot environment, ensure that the default value in the grub bootloader configuration is set properly with the currently booted 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 indicates that emcpbfs_linux config has not been used during PowerPath configuration. See Uninstall PowerPath from a Boot from SAN host using manual steps on page 42 to uninstall PowerPath.

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

Uninstall PowerPath from a Boot from SAN host using manual steps

Procedure to uninstall PowerPath from a Boot from SAN host using manual steps.

Procedure

- 1. Restore /etc/fstab to mount boot partitions by UUID.
 - For hosts with BIOS firmware and for OL 6 and RHEL 6 only:

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

• For hosts with BIOS firmware and for OL 7, RHEL 7, RHEL8, SLES 12 SP4, SLES 15, and SLES 15 SP1 only:

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

For hosts with UEFI firmware and for SLES 12 SP4, SLES 15 and SLES 15 SP1 only:

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

For hosts with UEFI firmware and for OL 6 and RHEL 6 only:

```
#/dev/emcpowera1 /boot ext4 _netdev 1 2
UUID=110223d8-f5f3-4058-812f-16d9e9ee7eab /boot ext4 defaults 1 2
#/dev/emcpowera2 /boot/efi vfat umask=0077,shortname=winnt 0 0
UUID=4AA3-84F3 /boot/efi vfat umask=0077,shortname=winnt 0 0
```

• For hosts with UEFI firmware and for OL 7, RHEL 7 and RHEL 8 only:

```
#/dev/emcpowera1 /boot ext4 nofail 1 2
UUID=110223d8-f5f3-4058-812f-16d9e9ee7eab /boot ext4 defaults 1 2
#/dev/emcpowera2 /boot/efi vfat umask=0077,shortname=winnt 0 0
UUID=4AA3-84F3 /boot/efi vfat umask=0077,shortname=winnt 0 0
```

- Remount the boot partitions so that it is mounted on the native device and not the pseudo device.
 - For hosts with BIOS firmware and all supported operating systems:

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

• For hosts with UEFI firmware and for SLES 12 SP4, SLES 15 and SLES 15 SP1 only:

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

For hosts with UEFI firmware and for OL6, OL 7, RHEL 6, RHEL 7, and RHEL 8 only:

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

- 3. Edit the /etc/lvm/lvm.conf file to restore the filter parameter so that it accepts all devices.
 - For OL6 and RHEL6:

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

• For all other operating systems:

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

- 4. Execute vgscan -v to rebuild the LVM2 cache.
- 5. Execute lvmdiskscan to ensure that the filter field is working correctly. Verify all the block devices are listed in the output.
- 6. Uninstall PowerPath.
 - Note: To review the procedure for uninstalling PowerPath, see Uninstall PowerPath on page 82.

Administering PowerPath in a Boot from SAN environment

CHAPTER 5

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.

Note: This section is only for managing SCSI device with PowerPath. See PowerPath and NVMe on page 21 for NVMe related procedures.

•	Dynamic LUN addition and removal	. 46
	Before you add or remove paths	
	Adding new paths to a PowerPath logical device	
•	Removing some of the paths to PowerPath logical devices	.49

Dynamic LUN addition and removal

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

(i) Note: Slow down I/O to the maximum extent possible during any dynamic addition or removal operations. Cluster or database integrity may be adversely affected if moderate or heavy I/O is running while LUN configuration is changed.

Adding a LUN

Follow this procedure to add a new LUN dynamically to PowerPath. If possible all I/O to the SCSI devices should be quiesced before this procedure as it rescans the SCSI buses.

Procedure

1. Add the new LUNs into the configuration.

(i) 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, isci, 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 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. Check the size of the pseudo and native devices using parted or fdisk -1 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

(i) 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. Scan and configure the additional paths.
 - Automatically scan and configure the additional paths:

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

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

- Manually scan and configure the additional paths:
- a. Scan each 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. 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.)
- 2. Confirm the configuration of the logical devices from which paths are removed.

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

3. Check the number of existing paths.

The path state should be alive for known good paths and dead for known bad paths.

- 4. If there is a problem, correct it before proceeding.
- 5. 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.

6. 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.

7. Obtain bus, TID, and LUN.

```
[root@hostname ~]# pp_inq -btl
```

Save the output of the command.

- 8. If using character I/O (sg) devices on Linux, run sg map -i -x and then save the output.
- 9. 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.

 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.

11. 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.

12. 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.
- 13. 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.

14. Remove the stale native devices from the kernel.

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

- 15. 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.

- 16. Correct any issues that are detected before saving the PowerPath configuration or using the new logical devices.
- 17. Save the new configuration.

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

Managing paths to a logical device

CHAPTER 6

Maintaining PowerPath

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

•	/etc/fstab options	. 54
•	Configure remote management service	. 54
	Configure event monitoring with SNMP traps	
•	Changing the language of installation messages	.60
•	Saving new configurations	. 61
•	Configuring LVM2 support	
•	Configuring PowerPath on cluster nodes	64
•	Hot swapping an HBA	.65
•	Configuring for R1/R2 boot failover support	.67
•	· · · · · · · · · · · · · · · · · · ·	
•	11 - 2	
•	Expand the file system on a pseudo-device	
•		

/etc/fstab options

/etc/fstab options that are used in FC, iSCSI, and Open-FCoE devices.

CAUTION Configure only /dev/emcpowerX pseudo block devices that PowerPath created for the use of applications. Do not use /dev/sdX native block devices for application IOs. Failover and load-balancing feature are not available for IOs sent directly to native block devices.

FC, iSCSI, and Open-FCoE 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/emcpoweral /mnt/emcpoweral ext3 defaults,_netdev,nofail 0 0
```

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

```
/dev/VG1/LV01 /mnt/VG1_LV01 ext3 defaults,_netdev,nofail 0 0
```

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.

Procedure

 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
   Change "Path Alive" event
                                                      Disabled, delay 60
sec
8 Change "All Paths Dead" event
                                                      Enabled, delay 300
sec
                                                      Disabled, delay 300
9
   Change "Not All Paths Dead" event
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
                                                   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
   Save configuration file, restart daemon and exit
   Save configuration file and exit
 3 Exit without saving configuration file
Enter number of desired choice [1, q, ?]: 1
Writing configuration file.
Changed parameters:
                                               Changed Value
Item
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 cannot 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 <code>emcp mond.conf</code> 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 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.

Tab	I_ 7			. _
ıab	ıe ə	· IV	lonitored	events

Monitored event	Event cause	Event ID	SNMP Trap
Path is Dead	Path is dead for 1 minute.	0x2,0x4,60	Yes
	Path is dead and then recovers within 1 minute.	0x2,0x4,60	No
All Paths Dead All paths are dead for 5 minutes.		0x20,0x40,300	Yes
	All paths are dead and then recover within 5 minutes.	0x20,0x40,300	No
Path Latency Max Exceeded	An I/O has taken more time than the set threshold on a path.	0x80000,0x80000, 0	Yes

The time threshold for each event can be modified as described in Editing the event on page 59.

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. These events, however, do 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 162
```

- 3. Configure snmpd persistence across restart.
 - a. Use the <code>chkconfig</code> or <code>systemctl</code> 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, RHEL8, OL7, SLES 12 SP4, SLES15, or SLES15SP1:

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

On all other kernels:

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

- (i) 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.

```
2 Change management service heartbeat interval 10 minutes
 3 Change management service port
                                                   9083
 4 Single path gatekeeper devices
5 Change SNMP trap generation status
                                                    Allow
                                                    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
                                                   Value
    It.em
 1 Change remote management service status
                                                  Disabled
 2 Change management service heartbeat interval 10 minutes
    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
                                                   Value
    Item
                                                    ____
                                                  Disabled
 1 Change remote management service status
    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
                                                   Enabled, delay 60
 6 Change "Path Dead" event
                                                  Disabled, delay 60
 7 Change "Path Alive" event
sec
 8 Change "All Paths Dead" event
                                                   Enabled, delay 300
sec
    Change "Not All Paths Dead" event
                                                    Disabled, delay 300
 sec
10 Change "Volume Dead" event
                                                    Disabled, delay 300
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
 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
    Change "Path Alive" event
                                                     Disabled, delay 60
 sec
    Change "All Paths Dead" event
                                                     Enabled, delay 300
 8
 sec
 9 Change "Not All Paths Dead" event
                                                     Disabled, delay 300
sec
10 Change "Volume Dead" event
                                                     Disabled, delay 300
sec
    Change "Max. Path Latency Exceeded" event
                                                     Disabled, delay 0
11
 sec
12 Write configuration file
13 Exit.
Enter number of desired choice [nn, q, ?]: 13
    It.em
 1 Save configuration file, restart daemon and exit
 2 Save configuration file and exit
    Exit without saving configuration file
 4 Back
Enter number of desired choice [1, q, ?]: 1
Writing configuration file.
Changed parameters:
                                                Changed Value
Item
 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

- 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, RHEL8, OL7, SLES 12 SP4, SLES15, or SLES15SP1:

```
[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 4 Supported languages

Language	Locale ID	Encoding
Chinese (simplified)	zh_CN.utf8	UTF8
English	En_US.utf8	
French	fr_FR.utf8	

Table 4 Supported languages (continued)

Language	Locale ID	Encoding
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 disable dm-multipath before continuing. To disable native multipathing, see Disabling dm-multipath on a RHEL 6.x or OL 6.x host on page 122 or Disabling dm-multipath on a RHEL 7.x or OL 7.x host on page 123.

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 as follows.

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 5 Linux Kernel Filter Fields

Kernel	Filter
SLES 15 SP1, RHEL 8	<pre>global filter = ["a/emcpower.*/", "r/sd.*/", "r/ nvme.*/", "r/disk.*/"]</pre>
SLES 12 SP4	<pre>global_filter = ["a/emcpower.*/", "r/sd.*/", "r/ disk.*/"]</pre>
SLES 15	<pre>global_filter = ["a/emcpower.*/", "r/sd.*/", "r/ nvme.*/", "r/disk.*/"]</pre>
RHEL 6.x	filter = ["a/emcpower.*/", "r/sd.*/", "r/disk.*/"]
RHEL 7.x	<pre>global filter = ["a/emcpower.*/", "r/sd.*/", "r/ disk.*/"]</pre>
Oracle Linux 6.x	filter = ["a/emcpower.*/", "r/sd.*/", "r/disk.*/"]
Oracle Linux 7.x	<pre>global filter = ["a/emcpower.*/", "r/sd.*/", "r/ disk.*/"]</pre>

- 2. Run vgscan v to rebuild the LVM2 cache.
- 3. Run lymdiskscan 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 Configure PowerPath on a Boot from SAN host on page 28:

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

Identify the underlying devices for the root/swap logical volumes and the /boot devices, if any. Modify the LVM filter field of the /etc/lvm/lvm.conf file to prevent LVM2 from scanning all SD device nodes, except for root/swap and /boot devices:

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.
 - (i) 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 6 Kernel Filter

Kernel Filter	Filter
SLES12 SP4, SLES 15, SLES 15 SP1, RHEL7.x, RHEL8, and OL7.x	global_filter = ["a /dev/.*/by-id/scsi-SSEAGATE_ ST373207LC_3KT4PMAN.* ", "a/emcpower.*/", "r/sd.*/", "r/disk.*/"]
RHEL 6.x and OL6.x	<pre>filter = ["a /dev/.*/by-id/scsi-SSEAGATE_ST373207LC_ 3KT4PMAN.* ", "a/emcpower.*/", "r/sd.*/", "r/ disk.*/"]</pre>

- (i) 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 lymdiskscan 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 emcpowerd
... Mapped emcpowere
... Mapped emcpowerf
... Mapped emcpowerf
```

- (i) Note: If devices are in use, import fails with device busy message.
- 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 Remapping Status
    Mapping
------
emcpowera no change
emcpowerc no change
emcpowerd no change
emcpowere no change
emcpowere no change
```

```
emcpowerf no change
emcpowerh remaps: emcpowerb emcpowerh (* Device name
already in use)
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:

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

Procedure to remove an HBA:

Procedure

- 1. Verify that the Fujitsu drivers shpchp 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. Also, 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.

(i) 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.

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 7 Operating System Support Steps

Operating System	Support steps:
RHEL 6.x and OL 6.x	 Create the lpfc.conf file in the /etc/modprobe.d/ directory. Edit the file and set the lpfc_max_luns parameter to 512. Recreate the initramfs by running the dracut -f command. Reboot the host. For more information, see the EMC Host Connectivity Guide for Linux or the vendor website.
RHEL 7.x, RHEL 8, OL 7.x, SLES12.x, SLES15, and SLES 15 SP1	 Edit the /etc/modprobe.d/lpfc.conf by adding the text: options scsi_mod max_luns=16383 max_report_luns=16384 options lpfc lpfc_max_luns=16384 Recreate the initrd. Point the initrd to the new initrd when booting.

CHAPTER 7

Upgrading PowerPath

Ensure that all the requirements are met, and then based on your requirement select a suitable upgrade procedure. The $/var/log/powerpath_install.log$ file contains the details of all the operations performed during install, uninstall and upgrade of PowerPath.

•	Preparing for upgrade	72
•	Upgrading PowerPath	72
•	Non-disruptive upgrade	77

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 6.0 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 60.
- If you encounter any problems while upgrading PowerPath, see Troubleshooting PowerPath on page 97.

Upgrading PowerPath

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

Upgrade PowerPath from a compressed archive on page 72

Upgrade PowerPath across Linux platforms on page 74

Upgrade PowerPath before upgrading RHEL update on Linux on page 76

Upgrade PowerPath from a compressed archive

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

Procedure

1. Log in as root.

2. Download the PowerPath archive from Dell EMC Online Support.

From Dell EMC Online Support, click **Support by Product** and then search for PowerPath Linux downloads.

3. Untar the PowerPath archive.

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

4. 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, 7.0.0.00.00.
- Build is the software build number.
- *Platform* is the Linux distribution, for example SLES12SP4 or RHEL6.
- 5. Start PowerPath.
 - On RHEL 6.x or OL 6.x

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

On RHEL 7.x, RHEL 8, OL 7.x, SLES 12 SP4, SLES 15, or SLES 15 SP1

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

Upgrade PowerPath across Linux platforms

Before you begin

(i) 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 PowerPath in a Boot from SAN environment on page 25.

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 74
- Upgrade PowerPath after upgrading from one major Linux release to the next on page 74
- Upgrade PowerPath before upgrading to the latest SLES Service Pack on page 75
- Upgrade PowerPath before upgrading RHEL update on Linux on page 76
- Upgrade PowerPath in a SteelEye LifeKeeper cluster environment on page 76

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, 7.0.0.00.00.
- 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, RHEL 8, OL 7.x, SLES 12 SP4, SLES 15, or SLES 15 SP1

```
[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:
 - a. Copy contents of /etc/emc/archive to another folder such as /tmp/ pp config backup.
 - b. Uninstall the new version of PowerPath.
 - c. Delete contents of /etc/emc/archive folder since that now contains configuration files that are updated by the new PowerPath version.
 - d. Install the old PowerPath version.
 - e. Copy the backed-up contents from step a to /etc/emc/archive folder.
 - f. Start PowerPath.
 While starting, PowerPath restores the config files from /etc/emc/archive and the old settings are restored.
- 5. 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:

Procedure

- 1. Upgrade PowerPath to the latest supported version for the current operating system version.
 - (i) Note: Do not start PowerPath service at this time.

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, 7.0.0.00.00.
- <build> is the software build number.
- <platform> is the Linux distribution, for example SLES.
- 2. Upgrade (migrate) 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, 7.0.0.00.00.
- <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 back up the LifeKeeper configuration, stop the existing PowerPath for each node, and then upgrade the node to the PowerPath of later version.

Before you begin

(i) Note: Ensure that the node being upgraded has no active resources or hierarchies.

Procedure

- 1. Log in as root.
- 2. Before upgrading PowerPath, type lkbackup -c --cluster to backup the LifeKeeper configuration.
 - Note: Ensure that you have a backup of the configuration that is stored in a safe location (for example, copy to a system not involved with the cluster).

- 3. Run 1kstop -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, 7.0.0.00.00.
- <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 1kstart 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.

Non-disruptive upgrade

Non-disruptive upgrade (NDU) enables PowerPath to upgrade from PowerPath for Linux 6.5 to any higher versions without any application downtime.

The applications need not be stopped, file systems on pseudo devices need not be unmounted, and volume groups need not be deactivated. The pseudo devices can be in use and IO going to them while the non disruptive upgrade is performed.

NDU is supported from all future releases including 6.5 patches, 6.5 service pack releases, 7.0 version, and above.

CAUTION PowerPath Non-Disruptive Upgrade from version 6.5 to 7.0 is not supported on cluster nodes. Perform only a normal upgrade on cluster nodes.

Supported Operating Systems

The following operating systems are supported for PowerPath NDU:

Table 8 Supported Operating Systems

Operating System		
RedHat	RHEL 7.2/7.3/7.4/7.5/7.6	RHEL 6.8/6.9/6.10
SuSE	SUSE 15	SUSE 12 SP4
Oracle	UEK 4 U6 , UEK4 U7	UEK 5.0, UEK 5.0 U1

Steps to perform PowerPath NDU

Procedure for performing PowerPath NDU.

Before you begin

Ensure that the following conditions are satisfied before a PowerPath NDU is attempted:

- A valid PowerPath license should be present on the host.
- No PowerPath Migration Enabler migrations are configured or in progress.
- · No dead, flaky, or iopf paths are present.
- No powermt, powermig, or emcpadm commands are in execution.
- PowerPath scripts such as emcplun linux or emcpbfs linux are not being executed.
- (i) Note: If any of these conditions are not met, then PowerPath NDU immediately aborts if attempted. Also, ensure that PowerPath NDU is attempted only during periods of low resource utilization on the host and low IO activity.

Procedure

1. Export environment variable *PP_NDU_ENABLE* in a terminal session.

```
#export PP NDU ENABLE=1
```

2. Execute the following command in the same terminal session in which the environment variable was set to begin PowerPath non-disruptive upgrade:

```
#rpm -Uvh DellEMCPower.LINUX-7.0.0.00.00-0XX.XXX.x86 64.rpm
```

(i) Note: If the preceding command is run from any other terminal session, a normal upgrade is performed. A normal upgrade requires that no pseudo devices are in use and cancels if any device is in use.

The following steps are automatically performed during a non-disruptive upgrade:

- Check if conditions are satisfactory to perform a non-disruptive upgrade.
- · Load thin module.
- Transition control of the devices and IOs from the core modules to the thin module.
 - Note: The thin module provides rudimentary failover and round robin load-balancing capabilities.

- Unload the modules of the currently installed version.
- · Install files of the new version.
- Load the core modules from the new version.
- Transition back the control of the devices and IOs from the thin module to the core modules.
- Unload the thin module.
- Uninstall files of the old version.
- Complete the upgrade.
- Confirm that the non-disruptive upgrade has completed successfully by checking for the following.
 - a. Check if rpm lists the upgraded PowerPath version properly.

```
[root@localhost ~]# rpm -qa | grep EMCpower.LINUX EMCpower.LINUX-7.0.0.00.00-064.el7.x86_64
```

b. Check if powermt display gives a valid output similar to the following:

```
[root@localhost ~]# powermt display
Symmetrix logical device count=13
VNX logical device count=10
Unity logical device count=15
SC logical device count=10
______
---- Host Bus Adapters ----- I/O Paths ----
Stats -----
### HW Path
                      Summary Total Dead IO/Sec Q-
IOs Errors
______
                    optimal
                           116
                                 0
0 qla2xxx
   0
                    optimal 116 0
3 qla2xxx
0
  Ω
```

Note: If there is a host panic or an accidental reboot while a NDU was in progress, see Recover from a failed PowerPath Non-Disruptive Upgrade due to a panic or reboot on page 99. If NDU fails for any other reasons, the reason for the failure is listed in the console and in /var/log/powerpath_install.log. Follow the recommended steps to fix the issue and try NDU again.

Upgrading PowerPath

CHAPTER 8

Removing PowerPath

Plan and prepare the host for removing or uninstalling PowerPath. You can either keep the existing PowerPath configuration and registration files for restoring in the next installation or completely remove them. The $/var/log/powerpath_install.log$ file contains the details of all the operations that are performed during install, uninstall, and upgrade of PowerPath.

•	Before you remove PowerPath	82
•	Uninstall PowerPath	82

Before you remove PowerPath

Before you remove PowerPath, ensure the following:

Procedure

- Ensure that no PowerPath devices are in use. Unmount all file systems from pseudo devices and ensure that none of the logical volumes on volumes groups from pseudo devices are in use.
- Run the lsof command and ensure that none of the storage devices that are managed by the PowerPath are listed. If the lsof command lists the processes along with opened PowerPath managed storage devices, then close such applications/processes and ensure that PowerPath managed storage devices are not in use.
- Manually remove references to PowerPath pseudo devices from system configuration files. For example, /etc/fstab.
- If the LVM filter had previously been updated to PowerPath recommended filter when PowerPath was initially installed and configured, restore the filter back to the default configuration. Pseudo devices are not available after PowerPath is uninstalled.

Uninstall PowerPath

Procedure

- To uninstall a host booting from a local disk:
 Uninstall PowerPath using the steps in Uninstall PowerPath on page 82.
- 2. To uninstall a host Boot from SAN setup:

Follow the relevant steps given in Chapter 3 Administering PowerPath in a Boot from SAN environment on page 25.

Uninstall PowerPath

Perform the following steps to uninstall PowerPath.

Procedure

- 1. Log in as root.
- 2. If the LVM filter had previously been updated to PowerPath recommended filter when PowerPath was initially installed and configured, restore the filter back to the default configuration. Pseudo devices will not be available after PowerPath is uninstalled.
- 3. Retrieve the full package name of the currently installed PowerPath version.

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

4. Uninstall the software.

```
[root@hostname ~]# rpm -ev EMCPower.LINUX-<version>_<build>
```

Note: In rare instances, 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 PowerPath is again installed and started on the host, it automatically restores the saved configuration from the archive directory.

Remove the saved configuration files

Procedure to remove the saved configuration files.

About this task

The PowerPath uninstallation process adds the extension . saved and saves the PowerPath configuration and registration files.

- /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 the saved configuration is not required, they can be removed with the following steps:

Procedure

1. Execute the following command to remove saved license and configuration 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.

Removing PowerPath

CHAPTER 9

PowerPath features for PowerMax array

PowerPath for Linux 7.0 supports the following features when the host is connected to a PowerMax array running microcode 5978 and above.

(i) Note:

- FX Licensing is supported with both SCSI and NVMe. All other features are supported only with SCSI devices.
- PowerPath for Linux 7.0 supports the following features when the host is connected to a PowerMax array running microcode 5978.444.444 and above:
 - Proactive Port Offline Autostandby
 - Flaky Path Reporting to PowerMax
 - SAN latency reporting

•	PowerPath Host Registration	86
	FX Licensing	
	PowerPath Device IO Statistics	
	PowerPath Device in Use Reporting	
	Proactive Port Offline Autostandby	
	Flaky Path Reporting to PowerMax	
	SAN latency reporting.	

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 feature 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.

(i) Note: All references to PowerMax refer 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
                     iopf limit
       Storage
       System Class Attributes
                    periodic autorestore = on
       Symmetrix
                      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 hostname 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 Ind198016 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 Ind198016 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 Ind198016 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.

```
: 000197801413
Symmetrix ID
   Host Name:
                          host123
                          SLES15SP1
   OS Version:
   OS Revision:
                           Linux - 4.12.14-197.10-default
   Hardware Vendor Name: Cisco Systems In
   PowerPath Version:
                          7.0.0
   PowerPath Patch Level: 00.00-064
   PowerPath License Info: FX
   Host Registration Time: 07/20/2019 01:58:15
   Host Connectivity type: FC
   Cluster Info:
      Cluster Name:
      Cluster Node Name:
       (1): 10000090FA53B94D
       (2): 10000090FA53B94C
```

(i) Note: For PowerMax arrays running microcode 5978.444.444 and above, PowerPath also sends one more detail in the registration command to the array, in addition to the preceding details. This is the count of the number of PowerPath Migration Enabler migrations that have taken place on the host, if either the source or target LUN is from a PowerMax array.

FX Licensing

Licensing PowerPath for Linux usually acquiring a license key and entering it with the **emcpreg** 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 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 ~]# emcpreg -I
 Host is licensed through FX array.

PowerPath Device IO Statistics

This section provides information about 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
                     iopf limit
       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 about 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.

(i) Note: If devices from PowerMax running microcode 5978.444.444 and above are used by Oracle database application, the database instance names are also sent to the array.

The storage administrators can use the device usage statistics to determine LUN usage on hosts and 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
                     iopf limit
       Storage
       System Class Attributes
                    periodic autorestore = on
       Symmetrix
                     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]
```

The details sent to the array can be viewed using the symclic command from the Solutions Enabler software:

Proactive Port Offline Autostandby

This section provides information about the PowerPath Proactive Port Offline Autostandby feature.

Proactive Port Offline Autostandby Overview

There is a new autostandby mode that is supported for devices from PowerMax running microcode 5978.444.444 and above.

This mode is known as *autostandby offline* (asb:ofln). PowerPath proactively sets this mode when it detects that a port maintenance operation is going to be conducted on a PowerMax Frontend adapter.

Before a planned maintenance activity, the array administrator runs a script for the specified front-end adapter. It triggers PowerPath to set paths to that adapter as *autostandby offline* (asb:ofln) and the paths are not used for IOs. Hence, when the administrator carries out the maintenance activity after a few minutes and the paths go dead, there are no unnecessary IO latencies as these paths are not picked for IOs. After the maintenance activity is completed and when the paths become alive, they are automatically set to active mode and then used for IOs.

The autostandby offline paths are listed like the following in powermt display output:

Log Messages

When PowerPath detects that a maintenance is going to be carried out on array ports and sets the paths as *autostandby offline*, messages such as the following are logged to the system log:

```
Aug 28 07:57:17 lnq244072 kernel: Info:Mpx:Path Bus 9 Tgt 37 Lun 4 to Symmetrix ID 000197600545 is set to asb:ofln.
Aug 28 07:57:17 lnq244072 kernel: Info:Mpx:Path Bus 9 Tgt 37 Lun 0 to Symmetrix ID 000197600545 is set to asb:ofln.
Aug 28 07:57:17 lnq244072 kernel: Info:Mpx:Path Bus 9 Tgt 37 Lun 8 to Symmetrix ID 000197600545 is set to asb:ofln.
Aug 28 07:57:17 lnq244072 kernel: Info:Mpx:Path Bus 9 Tgt 37 Lun 3 to Symmetrix ID 000197600545 is set to asb:ofln.
```

Flaky Path Reporting to PowerMax

This section provides information about PowerPath Flaky Path Reporting to PowerMax feature.

Flaky Path Reporting to PowerMax Overview

PowerPath already can identify paths with intermittent IO failures (also known as flaky paths) and mark them as *autostandby io-per-failure* (asb:iopf) automatically.

When paths are marked as autostandby, they are not used for IOs and hence performance issues can be avoided until the issues with those paths are fixed.

For devices from PowerMax running microcode 5978.444.444 and above, PowerPath also reports flaky path information to the array. This process helps array administrators to identify potential problems on the array if multiple hosts report flaky paths from the same port.

SAN latency reporting

This section provides information about SAN latency Reporting feature.

SAN latency reporting overview

PowerPath now reports SAN latency values for paths to PowerMax arrays running microcode 5978.444.444 and above.

As the vendor-specific commands used to calculate these values are not queued on the array port, their response is not dependent on how loaded the array is and hence accurately reflects the actual latency between the host initiator and array ports. SAN administrators can analyze these values to determine issues in the network that may be causing higher latencies and fix them.

Following latency values are determined automatically every 30 minutes for each host initiator to array target pair:

- Previous Latency Latency value reported from the previous iteration.
- Current Latency Most recently reported Latency value.
- Max Latency Max latency reported in the last 24 hours. This value gets reset every 24 hours.

View PowerPath SAN Latency in PowerPath Management Appliance

PowerPath Management Appliance (PPMA) displays SAN latency values in Bus View tab under PowerPath Monitor view section.

About this task

The appliance gets latency details periodically by polling hosts. These details can be refreshed manually by repolling the hosts.

Note: For more details about PowerPath Management Appliance 2.6, see *PowerPath Management Appliance Installation and Configuration Guide*.

CHAPTER 10

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	.94
•	PowerPath co-management behavior	. 94

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 121 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 9 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 10 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	powermt manage class=ess
HP StorageWorks XP AutoPath	powermt manage class=hpxp
MPxIO	powermt manage class= <class> Where class is any array that handles an MPxIO-configured array</class>

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 11 Commands to disable support

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

Coexistence with other path management software

CHAPTER 11

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	98
	PowerPath custom settings not preserved	
	Recover from a failed PowerPath Non-Disruptive Upgrade due to a panic or reboot	
	Resolve missing or deleted files	
	Problems booting the host	

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 18.

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, RHEL8, OL7, SLES 12 SP4, SLES 15, or SLES 15 SP1 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, RHEL8, OL7, SLES 12 SP4, SLES 15, or SLES 15 SP1:

```
[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
```

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.

Recover from a failed PowerPath Non-Disruptive Upgrade due to a panic or reboot

During a PowerPath Non-Disruptive Upgrade, if the host is accidentally restarted or if there is a host panic, then when the host comes up, PowerPath may not be active. This section describes the steps that are needed to recover from this scenario.

Procedure

1. After the host boots up, check if PowerPath is working properly. It returns an error like the following for powermt display command:

```
[root@localhost ~]# powermt display
Warning: PowerPath non disruptive upgrade is in progress.
Please retry after the upgrade is complete.
```

2. Check the PowerPath version installed:

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

```
EMCpower.LINUX-6.5.0.00.00-067.e17.x86_64

EMCpower.LINUX-7.0.0.00.00-064.e17.x86_64
```

- (i) Note: Two versions of PowerPath are listed. This is expected as the NDU has failed.
- 3. Uninstall all versions of PowerPath with the following command:

```
[root@localhost ~]# rpm --allmatches -ev EMCpower.LINUX
Preparing packages...
EMCpower.LINUX-6.5.0.00.00-067.el7.x86_64
EMCpower.LINUX-7.0.0.00.00-064.el7.x86_64
```

4. Install and start PowerPath again by following the procedure in Installing PowerPath on page 15.

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, RHEL8, OL7, SLES 12 SP4, SLES 15, or SLES 15 SP1:

```
[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 18.

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*.

Troubleshooting PowerPath

CHAPTER 12

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	104
•	Files created by PowerPath installation	104

Linux files modified by PowerPath installation

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

RHEL6 and OL6

- /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.

RHEL7 and OL7

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

RHEL8

• /lib/modules/*/modules.dep — For each /lib/modules subdirectory that gets a set of PowerPath drivers installed.

SLES12 SP4, SLES15, and SLES15 SP1

• /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
- ppme
- archive
- .pp_config_complete
- iotag.config
- ppmeusage

/etc/init.d/

RHEL6, OVM 3.x.x, and OL6

- PowerPath
- PowerPathPost

/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
- pp_reset_dm_deps.sh
- server.pem
- emcplun_linux
- 80-udev-pp-xio.rules
- 04-nvme-pp.rules
- 05-udev-pp.rules
- emcp_mond.conf
- emcp_mond.conf.bak
- emcp_serv.sh
- ppsetup

/etc/opt/emcpower/.tmp

SLES15, SLES15 SP1, SLES 12 SP4, RHEL 8, RHEL 7, RHEL 6, OL 6, OL 7, OVM 3.x.x

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

/etc/opt/emcpower/bin

- envsubst
- gettext
- gettext.sh

/etc/opt/emcpower/EMCpower.LINUX-7.0.0.00.00-064

- .pp_build_version
- .pp_vendor
- .pp_version
- EULA.pdf
- PowerPathPost
- PowerPath
- PowerPathPost-start.service (SLES15, SLES15 SP1, SLES 12 SP4, RHEL8, RHEL7 and OL7)
- PowerPathPost-stop.service (SLES15, SLES15 SP1, SLES 12 SP4, RHEL8, RHEL7 and OL7)
- PowerPath.service (SLES15, SLES15 SP1, SLES 12 SP4, RHEL8, RHEL7 and OL7)
- PowerPath.pp (RHEL8 only)
- boot-powerpath.sh
- boot.powerpath
- emcp_mond.sh
- emcpbfs_linux
- emcplun_linux
- EmcpMond.service
- emcp_serv.sh
- iotag.config
- modprobe.conf.pp
- pp_reset_dm_deps.sh
- pp_udev.sh
- pp_udev_xio.sh
- setup-powerpath.sh
- udev-pp.rules
- udev-pp-xio.rules
- 99powerpath
- nvme-pp.rules
- pp_nvme.sh
- ppsetup

/etc/opt/emcpower/EMCpower.LINUX-7.0.0.00.00-064/bin

- .cmds
- .drivers_base
- .drivers_ext
- .drivers_mgr

- .misc
- .sharedlibs
- .sharedlibs32
- .staticlibs

/etc/opt/emcpower/EMCpower.LINUX-7.0.0.00.00-064/bin/cmds

- emcp_cleanup
- emcp_mond
- emcp_mond_edit
- emcp_mond-sample.conf
- emcp_mond_no_ssl
- emcpadm
- emcpmgr
- emcpmigd
- emcpreg
- powercf
- powermig
- powermt
- pp_inq
- server.pem
- snmpd.conf
- powermt_display_dead
- powerndu

/etc/opt/emcpower/EMCpower.LINUX-7.0.0.00.00-064/bin/driver

SLES15 SP1

- sles15sp1default_x8664
- sles15sp1default_x8664/emcp.ko
- sles15sp1default_x8664/emcpdm.ko
- sles15sp1default_x8664/emcpgpx.ko
- sles15sp1default_x8664/emcpioc.ko
- sles15sp1default_x8664/emcpmpx.ko

SLES15

- sles15default_x8664
- sles15default_x8664/emcp.ko
- sles15default_x8664/emcpdm.ko
- sles15default_x8664/emcpgpx.ko
- sles15default_x8664/emcpioc.ko
- sles15default_x8664/emcpmpx.ko
- sles15default_x8664/emcpthin.ko

SLES 12 SP4

- sles12sp4default_x8664
- sles12sp4default_x8664/emcp.ko
- sles12sp4default_x8664/emcpdm.ko
- sles12sp4default_x8664/emcpgpx.ko
- sles12sp4default_x8664/emcpioc.ko
- sles12sp4default_x8664/emcpmpx.ko
- sles12sp4default_x8664/emcpthin.ko

RHEL7

- rhel7_x8664/emcp.ko
- rhel7_x8664/emcpdm.ko
- rhel7_x8664/emcpgpx.ko
- rhel7_x8664/emcpioc.ko
- rhel7_x8664/emcpmpx.ko
- rhel7_x8664/emcpthin.ko

RHEL6

- rhel6_x8664/emcp.ko
- rhel6_x8664/emcpdm.ko
- rhel6_x8664/emcpgpx.ko
- rhel6_x8664/emcpioc.ko
- rhel6_x8664/emcpmpx.ko
- rhel6_x8664/emcpthin.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/emcpmpx.ko
- ol6_uek4_r2_x8664/emcpthin.ko

OL7

- ol7_uek4_r2_x8664
- ol7_uek4_r2_x8664/emcp.ko
- ol7_uek4_r2_x8664/emcpdm.ko
- ol7_uek4_r2_x8664/emcpgpx.ko
- ol7_uek4_r2_x8664/emcpioc.ko
- ol7_uek4_r2_x8664/emcpmpx.ko
- ol7_uek4_r2_x8664/emcpthin.ko
- ol7_uek5_r2_x8664
- ol7_uek5_r2_x8664/emcp.ko
- ol7_uek5_r2_x8664/emcpdm.ko

- ol7_uek5_r2_x8664/emcpgpx.ko
- ol7_uek5_r2_x8664/emcpioc.ko
- ol7_uek5_r2_x8664/emcpmpx.ko
- ol7_uek5_r2_x8664/emcpthin.ko
- ol7_uek5_x8664
- ol7_uek5_x8664/emcp.ko
- ol7_uek5_x8664/emcpdm.ko
- ol7_uek5_x8664/emcpgpx.ko
- ol7_uek5_x8664/emcpioc.ko
- ol7_uek5_x8664/emcpmpx.ko
- ol7_uek5_x8664/emcpthin.ko

/etc/opt/emcpower/EMCpower.LINUX-7.0.0.00.00-064/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
- libemcpmp.so
- libpn.so

/etc/opt/emcpower/EMCpower.LINUX-7.0.0.00.00-064/bin/lib32

- libemcp.so
- libemcp_core.so
- libemcp_lam.so
- libemcp_lic_rtl.so
- libemcp_mp_rtl.so
- libemcp_mpapi_rtl.so
- libemcpmp.so
- libpn.so

/etc/opt/emcpower/EMCpower.LINUX-7.0.0.00.00-064/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.LINUX-7.0.0.00.00-064/man

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

/etc/opt/emcpower/EMCpower.LINUX-7.0.0.00.00-064/99powerpath/

- check
- install
- module-setup.sh

/etc/udev/rules.d

- 04-nvme-pp.rules
- 05-udev-pp.rules
- 80-udev-pp-xio.rules

/lib

- libemcp.so
- libemcp_core.so
- libemcp_lam.so
- libemcp_lic_rtl.so
- libemcp_mp_rtl.so
- libemcp_mpapi_rtl.so
- libemcpmp.so
- libpn.so

/lib/udev

- pp_udev.sh
- pp_udev_xio.sh
- pp_nvme.sh

/lib/modules/<kernel_version>/powerpath

- emcpdm.ko
- emcpgpx.ko
- emcpioc.ko
- emcp.ko
- · emcpmpx.ko
- · emcpthin.ko

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

/lib64

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

/sbin

- emcp_cleanup
- emcp_mond
- emcp_mond_edit
- emcp_mond_no_ssl
- emcpadm
- emcpmigd
- emcpreg
- powermig
- pp_inq
- powermt
- powermt_display_dead

/usr/lib/systemd/system

RHEL8, RHEL7, SLES15, SLES15 SP1, SLES 12 SP4 and OL7

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

/usr/lib/PowerPath

RHEL8, RHEL7, SLES15, SLES 15 SP1, SLES 12 SP4 and OL7

- PowerPath
- PowerPathPost
- emcp_mond.sh

/usr/share/man/man1

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

/usr/lib/dracut/modules.d/99powerpath

SLES15, SLES15 SP1, SLES 12 SP4, RHEL8, RHEL7, and OL7

· module-setup.sh

/usr/share/dracut/modules.d/99powerpath/

RHEL 6 and OL6

- check
- install

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 <code>-mindirect-branch=thunk-extern</code> 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 7.0 and retpoline

PowerPath for Linux 7.0 can be installed only on retpoline kernels. It does not work on older kernels that have been built with a pre-retpoline compiler.

CAUTION Check the PowerPath support matrix before installing or upgrading to PowerPath for Linux 7.0 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 functions are used to implement retpoline itself. Check if those functions are present in the running kernel. A snippet of the output is as follows:

```
# cat /proc/kallsyms | grep __x86_indirect_thunk_
ffffffffff81a04000 T __x86_indirect_thunk_rax
fffffffff81a04020 T __x86_indirect_thunk_rbx
fffffffff81a04040 T __x86_indirect_thunk_rcx
fffffffff81a04060 T __x86_indirect_thunk_rdx
fffffffff81a04080 T __x86_indirect_thunk_rsi
fffffffff81a040a0 T __x86_indirect_thunk_rdi
ffffffff81a040c0 T __x86_indirect_thunk_rbp
ffffffff81a040e0 T __x86_indirect_thunk_r8
```

(i) Note: For more information about retpoline kernels, see the relevant operating system documentation of vendors. All queries specific to retpoline kernels should be directed to the operating system vendor.

PowerPath for Linux 7.0 and non-retpoline (older) kernels

This section provides information about what is expected if PowerPath for Linux 7.0 is installed on a non-retpoline kernel.

Fresh install

Installing PowerPath 7.0 fails on non-retpoline kernels with the following error:

Upgrading PowerPath

Upgrading an older version of PowerPath to version 7.0 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-7.0 PowerPath modules are not built with a retpoline compiler. Even if PowerPath is upgraded to version 7.0, the IBRS mitigation is active until a host reboot. To avoid the second reboot, upgrading PowerPath to 7.0 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 7.0 on a non-retpoline kernel:

```
[root@hostname ~] # rpm -Uvh DellEMCPower.LINUX-7.0.0.00.00-064.RHEL6.x86 64.rpm
                 Preparing...
                  ############ [100%]
1:EMCpower.LINUX
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:
   - 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 7.0 starts normally on the new retpoline kernel.

Starting PowerPath after upgrade on a non-retpoline kernel

If PowerPath for Linux 7.0 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
```

Retpoline Kernels

APPENDIX B

Dynamic Linking of OpenSSL Libraries

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

•	Overview	. 118
	Changes	
	Expected behavior	

Overview

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

In versions of PowerPath for Linux 64 and earlier, 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.

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_ss*/ is used when OpenSSL library is not installed or an incompatible OpenSSL version is found to be installed.

Expected behavior

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

With a compatible version of OpenSSL library

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 about 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-7.0.0.00.00-064.RHEL7.x86_64.rpm Preparing... ########################## [100%]
Updating / installing...
1:EMCpower.LINUX-7.0.0.00.00-064.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-7.0.0.00.00-064.RHEL7.x86 64.rpm
                      ########### [100%]
Preparing...
Updating / installing...
   All trademarks used herein are the property of their respective owners.
NOTE: License registration is not required to manage the CLARiiON AX series
arrav.
*** 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 about 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

Dynamic Linking of OpenSSL Libraries

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 a RHEL 6.x or OL 6.x host	122
•	Disabling dm-multipath on a RHEL 7.x or OL 7.x host	123
•	Disabling dm-multipath on a RHEL8.x, SLES12 SP4, SLES15, or SLES15 SP1 host	124

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

Procedure for disabling dm-multipath on a RHEL 6.x or OL 6.x host.

Procedure

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

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

2. Restart multipathd service for the multipath config file to get into effect:

```
[root@hostname ~]# /etc/init.d/multipathd restart
```

OR

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

3. Stop dm-multipath service and ensure that dm-multipath does not start automatically on boot.

```
[root@hostname ~]# /etc/init.d/multipathd stop
[root@hostname ~]# chkconfig multipathd off
[root@hostname ~]# chkconfig --list multipathd
```

Service should be off on all run levels.

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

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

- 5. Add a new boot entry with the new ramdisk image file created in previous step in the / boot/grub/menu.lst file.
- 6. 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 a RHEL 7.x or OL 7.x host

Procedure for disabling dm-multipath on a RHEL 7.x, RHEL 8, 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. Restart multipathd service for the multipath config file to get into effect:

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

3. Stop dm-multipath service and ensure dm-multipath does not start automatically on boot.

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

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

```
[root@hostname \sim]# dracut /boot/initramfs-wo-DM-$(uname -r).img $(uname -r)
```

- 5. 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.
- 6. 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'
```

7. 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.

Disabling dm-multipath on a RHEL8.x, SLES12 SP4, SLES15, or SLES15 SP1 host

Procedure for disabling dm-multipath on a RHEL 8.x, SLES15, or SLES15SP1 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. Restart multipathd service for the multipath config file to get into effect:

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

3. Stop dm-multipath service and ensure dm-multipath does not start automatically on boot.

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

4. Run multipath -ll

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

Should not return any dm devices.

APPENDIX D

Booting with duplicate PVs on SLES and RHEL8 BFS hosts

This section provides the background information about how SLES and RHEL 8 hosts handle logical volumes when duplicate physical volumes (PVs) are found.

•	RHEL8, SLES hosts, and LVM	126
•	RHEL8, SLES BFS hosts, and duplicate PVs	127
	Configure LVM to accept duplicate PVs	
•	Recreate new ram disk image	.127

RHEL8, SLES hosts, and LVM

On RHEL8, SLES12 SP4, SLES15, and SLES15SP1, 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 lymetad 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-
xzPvkl.
 WARNING: Not using device /dev/sdg2 for PV USHeZ8-rhWS-AuOY-LYL1-AYHZ-Aew4-
xzPvkl.
 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 RHEL8, SLES15, and SLES15SP1 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.

RHEL8, SLES BFS hosts, and duplicate PVs

In case of RHEL8 and SLES BFS hosts with root FS logical volume, only one physical volume should present 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.

If 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.

(i) Note: Follow Configure LVM to accept duplicate PVs on page 127 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.

Recreate new ram disk image

On RHEL 8 host, the ramdisk image is named as initramfs. On SLES 12, SLES 15, and SLES 15 SP1 it is named as initrd.

About this task

On RHEL 8:

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

On SLES 12 SP4/SLES 15/SLES15P1:

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

APPENDIX E

Disable NVMe multipath

This section provides the procedure to disable	NVMe multipath on SLES	15 and SLES 15 SP1
hosts		

•	Disabling NVMe r	multipath on S	LES 15 and SLES	S 15 SP1 hosts	130

Disabling NVMe multipath on SLES 15 and SLES 15 SP1 hosts

Procedure for disabling NVMe multipath on SLES 15 and SLES 15 SP1 hosts.

Procedure

 Add an option to the default kernel command line to prevent NVMe multipath from starting automatically on host reboot. Edit the /etc/default/grub file and add nvmecore.multipath=//to the GRUB_CMDLINE_LINUX_DEFAULT system variable as follows:

```
GRUB_CMDLINE_LINUX_DEFAULT="resume=/dev/disk/by-id/scsi-3600508e00000000eclef0cfdc194b08-part7 splash=silent quiet showopts nvme-core.multipath=N crashkernel=2048M,high crashkernel=512M,low"
```

2. Update grub configuration and reboot host.

```
# grub2-mkconfig -o /boot/grub2/grub.cfg
Generating grub configuration file ...
Found theme: /boot/grub2/themes/SLE/theme.txt
Found linux image: /boot/vmlinuz-4.12.14-150.14-default
Found initrd image: /boot/initrd-4.12.14-150.14-default
Found linux image: /boot/vmlinuz-4.12.14-25.25-default
Found initrd image: /boot/initrd-4.12.14-25.25-default
done
```

3. Verify that the kernel started with NVMe multipath that is disabled by checking the value of /proc/cmdline.

```
# grep nvme-core.multipath /proc/cmdline
BOOT_IMAGE=/boot/vmlinuz-4.12.14-150.14-default
root=UUID=b5bc20fd-1fb3-4538-b906-b558fb914adf resume=/dev/disk/by-id/
scsi-3600508e000000000eclef0cfdc194b08-part7 splash=silent quiet
showopts nvme-core.multipath=N crashkernel=2048M,high
crashkernel=512M,low
```

NDU Operating System 78

Non-Disruptive Upgrade 31

new LUN 46 non retploine 114

NVMe 22

Α audience 9 OpenSSL 118, 119 Auto host registration 87 **Auto Host Registration Settings 87** automatic host registration 86 powerpath 86, 87 В Powerpath 86 **BFS 26** Boot from SAN 26 C comments 9 PPMA 54 conventions for publication 9 preface 9 D Device in use 89 Device in Use 90 R Device IO statistics 89 **Device IO Statistics 89** retpoline 113 Retpoline 114 Ε RPM file 18 environment compatibility 16 S F Flaky Path Reporting 92 SNMP traps 56 FX Licensing 88 Н Т Host Registration 86 host system compatibility 16 install power path 114 Installation 118, 119 license 18 М Mode of operation 86 Ν NDU 31, 77

PowerMax array 85 PowerPath 118, 119 PowerPath installation 15 PowerPath licensing 17 PowerPath Management Appliance 54 PowerPath NDU 78 Proactive Port Offline Autostandby 91 Proactive Port Offline Autostandy 91

ram disk image 127 related documentation 9

SAN latency reporting 92 support information 9

third-party path management 17

Veritas Storage Foundation 20

Index