



# EMC<sup>®</sup> PowerPath<sup>®</sup> for Solaris

Version 6.0 and minor releases

## Installation and Administration Guide

P/N 302-001-309  
REV 04

Copyright © 1997 - 2019 EMC Corporation. All rights reserved. Published in the USA.

Published January 2019

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

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

EMC<sup>2</sup>, EMC, EMC Centera, EMC ControlCenter, EMC LifeLine, EMC OnCourse, EMC Proven, EMC Snap, EMC SourceOne, EMC Storage Administrator, Acartus, Access Logix, AdvantEdge, AlphaStor, ApplicationXtender, ArchiveXtender, Atmos, Authentica, Authentic Problems, Automated Resource Manager, AutoStart, AutoSwap, AVALONidm, Avamar, Captiva, Catalog Solution, C-Clip, Celerra, Celerra Replicator, Centera, CenterStage, CentraStar, ClaimPack, ClaimsEditor, CLARiiON, ClientPak, Codebook Correlation Technology, Common Information Model, Configuration Intelligence, Connectrix, CopyCross, CopyPoint, CX, Dantz, Data Domain, DatabaseXtender, Direct Matrix Architecture, DiskXtender, DiskXtender 2000, Document Sciences, Documentum, elInput, E-Lab, EmailXaminer, EmailXtender, Engenuity, eRoom, Event Explorer, FarPoint, FirstPass, FLARE, FormWare, Geosynchrony, Global File Virtualization, Graphic Visualization, Greenplum, HighRoad, HomeBase, InfoMover, Infoscapes, InputAccel, InputAccel Express, Invista, Ionix, ISIS, Max Retriever, MediaStor, MirrorView, Navisphere, NetWorker, OnAlert, OpenScale, PixTools, Powerlink, PowerPath, PowerSnap, QuickScan, Rainfinity, RepliCare, RepliStor, ResourcePak, Retrospect, RSA, SafeLine, SAN Advisor, SAN Copy, SAN Manager, Smarts, SnapImage, SnapSure, SnapView, SRDF, StorageScope, SupportMate, SymmAPI, SymmEnabler, Symmetrix, Symmetrix DMX, Symmetrix VMAX, TimeFinder, UltraFlex, UltraPoint, UltraScale, Unisphere, Viewlets, Virtual Matrix, Virtual Matrix Architecture, Virtual Provisioning, VisualSAN, VisualSRM, VMAX, VNX, VNXe, Voyence, VPLEX, VSAM-Assist, WebXtender, xPression, xPresso, YottaYotta, the EMC logo, and the RSA logo, are registered trademarks or trademarks of EMC Corporation in the United States and other countries. Vblock is a trademark of EMC Corporation in the United States.

All other trademarks used herein are the property of their respective owners.

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

# CONTENTS

<b>Preface</b>	
<b>Chapter 1</b>	<b>Installing PowerPath</b>
	Before you install ..... 12
	Obtain current information ..... 12
	Choose a convenient time ..... 12
	Locate your license registration key(s)..... 12
	License ..... 13
	Prepare the host and storage system..... 13
	Installing PowerPath ..... 15
	Install PowerPath on Solaris 11.x ..... 15
	Install PowerPath using Oracle Solaris IPS repositories ..... 15
	Installing PowerPath using IPS repositories ..... 16
	Installing PowerPath on alternate Boot environment ..... 17
	Installing PowerPath offline..... 17
	bootadm update-archive error installation ..... 18
	PowerPath and MPxIO..... 18
	Arrays managed by MPxIO ..... 18
	Manage all array classes using PowerPath ..... 18
	Manage arrays using both PowerPath and MPxIO ..... 19
	Disable MPxIO on Solaris 11.x..... 19
	Remove or unmanage a device from PowerPath management ..... 20
	Best practices for PowerPath and MPxIO ..... 20
	Coexistence with third-party multipathing software..... 20
<b>Chapter 2</b>	<b>PowerPath in a Cluster Environment</b>
	PowerPath in an Oracle Solaris Cluster 4.2 ..... 24
	Installing PowerPath in a new Oracle Solaris Cluster 4.2..... 24
	Installing PowerPath into an existing Oracle Solaris Cluster 4.2 ..... 24
	PowerPath in a Veritas Cluster Server (VCS) cluster ..... 25
	Installing PowerPath in a new VCS cluster ..... 25
	Integrating/upgrading PowerPath into an existing VCS cluster..... 26
<b>Chapter 3</b>	<b>Configuring a PowerPath Boot Device on Solaris</b>
	Introduction ..... 30
	Boot from PowerPath pseudo device on Solaris 11.x ZFS (SPARC) ..... 30
	Removing PowerPath control over a boot device on Solaris 11.x..... 31
	Moving the boot device from a native device to an emcpower device (UFS) 32
	Changing the dump device path to the pseudo device in an R1/R2 boot failover setup ..... 33
	Recovery procedure..... 33
<b>Chapter 4</b>	<b>Upgrading PowerPath</b>
	Before you upgrade..... 36
	Upgrading PowerPath on Solaris 11 ..... 36
	Upgrading PowerPath on Solaris 11 ..... 36

	Upgrading PowerPath on Boot from SAN.....	36
<b>Chapter 5</b>	<b>Migrating to PowerPath</b>	
	Migrating from HP StorageWorks Secure Path .....	38
	Secure Path operation on Solaris .....	38
	Migrating from Secure Path to PowerPath .....	38
	Migrating from IBM Subsystem Device Driver (SDD).....	39
	SDD operation on Solaris .....	39
	Migrating from SDD to PowerPath.....	39
	Migrating from Hitachi Dynamic Link Manager (HDLM) .....	40
	HDLM operation on Solaris.....	40
	Migrating from HDLM to PowerPath .....	40
<b>Chapter 6</b>	<b>Removing PowerPath</b>	
	Before you remove PowerPath.....	42
	Removing PowerPath on Solaris 11.x .....	42
	Removing PowerPath on Boot from SAN (SPARC) .....	43
	Removing PowerPath on an alternate boot environment .....	44
	Removing PowerPath offline .....	44
	PowerPath interoperability with ZFS file systems.....	44
	Removing PowerPath with existing zpools .....	44
	After you remove PowerPath .....	44
	Removing PowerPath files .....	44
	Rebooting the host.....	45
<b>Chapter 7</b>	<b>PowerPath Administration on Solaris</b>	
	PowerPath Viewer and Management Component .....	48
	Event monitoring daemon using SNMP.....	48
	Monitored events.....	48
	Configuring the event monitoring daemon.....	49
	Trap packet format .....	50
	PowerPath and VNX OE and CLARiiON storage systems .....	51
	Ensuring a sufficient stack size .....	51
	Booting a host with built-in Fibre Channel ports .....	51
	Rebooting and custom settings.....	52
	R1/R2 boot failover support.....	52
	R1/R2 supported configurations .....	52
	Enabling R1/R2 boot disk failover .....	53
	Booting the R2 host after failover .....	53
	Booting the R1 host after fallback .....	53
	Device naming .....	54
	Selecting a device naming convention.....	55
	Reconfiguring PowerPath devices online .....	57
	Dynamic Reconfiguration .....	58
	Removing an HBA from a PowerPath configuration .....	58
	Cleaning up incorrect path or logical volume removal.....	59
	Adding an HBA to a PowerPath configuration.....	59
	Managing and unmanaging Celerra iSCSI devices .....	60
	Manage Celerra iSCSI devices under PowerPath .....	60
	Unmanage Celerra iSCSI devices under PowerPath.....	60
	Supporting Veritas Volume Manager.....	60
	Using native device names for EMC devices .....	61

	Exporting emcpower devices to Non-Global Zones .....	62
	powercf configuration utility.....	63
<b>Chapter 8</b>	<b>Troubleshooting</b>	
	Audit and error messages .....	66
	Logging messages to a common file .....	66
	Troubleshooting scenarios.....	66
<b>Appendix A</b>	<b>Files Changed by PowerPath</b>	
	Solaris files modified by PowerPath installation .....	72
	SPARC hosts .....	72
	x86 hosts .....	72
	/etc/system .....	72
	/etc/profile .....	73
	/etc/.login .....	73
	Files created by PowerPath installation .....	73
	Solaris 11.x x86 hosts.....	73
	Solaris11.x SPARC hosts .....	77
<b>Index</b>		



# PREFACE

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

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

---

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

---

## Audience

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

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

## Related documentation

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

## Revision history

The following table presents the revision history of this document:

Revision	Date	Description
04	January 2019	Removed Solaris 10 references.
03	May 2017	Included PowerPath in an Oracle Solaris Cluster 3.x in “PowerPath in a Cluster Environment”. Updated Forceload Statements in “Files Changed by PowerPath”.
02	March 2017	Corrected Solaris LDOM typo to I/O domain in “Installing PowerPath”.
01	May 2015	First release of the PowerPath for Solaris Version 6.0

## Conventions used in EMC documents

EMC uses the following conventions for special notices:



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

---

---

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

---

### **IMPORTANT**

---

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

---

## Typographical conventions

EMC uses the following type style conventions in this document:

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



## Where to get help

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

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

<http://support.EMC.com>

Technical support — For technical support, go to EMC WebSupport on Powerlink. To open a case on EMC WebSupport, you must be a WebSupport customer. Information about your site configuration and the circumstances under which the problem occurred is required.

## Your comments

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

[techpubcomments@emc.com](mailto:techpubcomments@emc.com)



# CHAPTER 1

## Installing PowerPath

This chapter describes how to install PowerPath on a Solaris host and how to upgrade from an earlier version of PowerPath.

- ◆ Before you install ..... 12
- ◆ Installing PowerPath ..... 15
- ◆ Install PowerPath on Solaris 11.x ..... 15
- ◆ PowerPath and MPxIO ..... 18
- ◆ Coexistence with third-party multipathing software ..... 20

## Before you install

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

### Obtain current information

Check the EMC Online Support site for current information:

- ❑ PowerPath documentation—EMC updates PowerPath documentation, including this installation guide, when it releases new features with a service pack or when documentation errors are reported.
- ❑ Service packs and upgrades—You can download PowerPath service pack software from EMC Online Support. Determine which service packs (if any) to install after PowerPath, and whether those service packs have any additional installation prerequisites.
- ❑ PowerPath Configuration Checker—PowerPath Configuration Checker (PPCC) is a software program that verifies that a host is configured to E-Lab Interoperability Navigator standards with the hardware and software required for PowerPath multipathing features (failover and load-balancing functions, licensing, and policies). Prior to installing or upgrading PowerPath, download the latest version of EMC Grab available on EMC Online Support and then run PPCC. This ensures that the system version used by PPCC includes the latest configuration information. PPCC is available for download on EMC Online Support.

EMC updates this database regularly between scheduled releases and service pack releases.

### Choose a convenient time

Installing PowerPath may require that you reboot the host. Plan to install or upgrade PowerPath when a reboot will cause minimal disruption to your site.

---

#### Note:

- A reboot may not be required after upgrading PowerPath.
  - PowerPath supports offline installation. [“Installing PowerPath offline” on page 17](#) provides more information.
- 

### Locate your license registration key(s)

Effective February 15, 2011, instead of the physical Right To Use (RTU), the default delivery method for PowerPath licenses is electronic. An electronic License Authorization Code (LAC) is sent by email in order to redeem the license key on the EMC Online Support Licensing Service Center. From EMC Online Support, go to **Service Center** > **Licensing**. This does not affect upgrades because PowerPath retains existing license information.

Physical RTU cards are still available as an option. EMC Global Support, at 1-800-svc4emc or 1-800-782-4362, can provide more information. The *PowerPath Family Electronic License Ordering Process Technical Notes*, available on EMC Online Support site, provides more information about the PowerPath license electronic ordering process.

## License

Take note of the following regarding licenses:

- ◆ In nonvirtual environments, installing PowerPath on a host connected exclusively to a CLARiiON AX-series array does not require a license. PowerPath provides full support, with or without a PowerPath license, when the host is connected exclusively to a CLARiiON AX-series array.
- ◆ In Oracle Solaris LDOM environments, the PowerPath license must be installed on the control domain only.
- ◆ If you are upgrading from an earlier version of PowerPath, you do not need to reregister. PowerPath will use your old key.
- ◆ Type the registration number exactly as it appears on the card or letter. To reduce common typographical errors, the **License Key** field accepts either uppercase or lowercase letters, and certain numbers and letters are interchangeable. Specifically, an entry of the alphabetic letters O, I, S, and B is equivalent to an entry of the numbers, 0, 1, 5, and 8.
- ◆ Apply one license per server. When different devices use different licenses, the **powermt set** command with the **dev=all** option may apply the most restrictive license to all devices.

## Prepare the host and storage system

To prepare the host and storage system:

- Before installing PowerPath, stop any I/O going to native devices.
- Go to the EMC Online Support site and verify that your environment meets the requirements found in the *PowerPath for Solaris Release Notes* and the E-Lab™ Interoperability Navigator by downloading the latest version of the EMC Grab utilities from Powerlink and then running the PowerPath Configuration Checker (PPCC). For PPCC installation information, the *PowerPath Configuration Checker User Guide* provides additional information.
- Root volume mirrors need to be broken before performing an installation or upgrade of PowerPath. Once installation is complete, the root volume mirror can be synchronized again.

---

**Note:** If you do not want MPxIO to manage any devices, disable it globally, otherwise consult official Oracle documentation regarding MPxIO.

---

- Ensure that the Symmetrix SCL director flag is set to OFF. This flag was required for Sun Cluster 2.x environments; it is no longer required for supported Sun Cluster versions (versions 3.x or later).

- ❑ For PowerPath Automatic Host-Array Registration to work properly, ensure that the **/etc/hosts** file or its equivalent is configured such that the local IP is mapped to only localhost and the valid network IP is mapped to the name of the host. For example:

```
<loopback IP address> localhost
<host IP address> <hostname.domainname> <hostname>
```

- ❑ When using a native Emulex HBA driver on a Solaris host, set the driver configuration file parameter to **no-device-delay=0**. This setting is not required when using Leadville drivers.

**⚠ CAUTION**

**Failure to set this parameter may result in very poor application performance when paths are down.**

- ❑ If your host is connected to either:
  - A Hitachi storage system
  - An VNX OE and CLARiiON storage system

disable Solaris native multipathing (MPxIO) globally before you install PowerPath. This prevents Solaris MPxIO from automatically managing VNX OE, CLARiiON or Hitachi devices. “PowerPath and MPxIO” on page 18 provides more information.

- ❑ Set up the Fibre Channel port and LUN addresses. The host connectivity guides (Symmetrix®, VNX OE, and CLARiiON arrays), VNX OE Storage-System Support website, CLARiiON Storage-System Support website, or the appropriate documentation from your vendor (third-party arrays) provides more information.
- ❑ Use the Solaris **format** utility to format, partition, and label the unused storage system devices such that each unused device has a partition 2 that occupies the entire device. Do not, however, use or mount these devices before installing PowerPath.
- ❑ Run **powermt config** after changing from EFI to SMI labels or from SMI to EFI labels.
- ❑ Configure HBA drivers.

**⚠ CAUTION**

**Be sure to follow HBA driver configuration guidelines outlined by EMC in the interoperability information and product documentation on EMC Online Support. Using improper settings can cause erratic failover behavior, such as greatly increased I/O delays.**

- ❑ For hosts connected to storage arrays through a Fibre Channel switch (that is, a fabric), configure the HBAs using persistent binding for SCSI target IDs. This is required for native HBA drivers, but not for Leadville drivers.

**⚠ CAUTION**

**Failure to configure persistent binding could result in the loss or corruption of data.**

The *Host Connectivity Guide for Sun Solaris*, the HBA driver documentation, and the .conf file for your HBA type provides more information about persistent binding.

- ❑ If any ControlCenter® agents are running on the host, stop the agents. Follow the instructions in the ControlCenter documentation.

- ❑ If the Navisphere® or Unisphere™ agent is installed on the host, stop the agent. Follow the instructions in the Navisphere and Unisphere documentation.

---

**Note:** The PowerPath installation stops the Navisphere and Unisphere agent before installing PowerPath and attempts to restart the agent after installation of PowerPath completes. However, in some situations, such as SP reboot during the upgrade, PowerPath may have problems restarting the agent. Therefore, EMC recommends stopping the Navisphere and Unisphere agent before upgrading and restarting the agent after the upgrade. Alternatively, once the upgrade is complete, ensure that the agent is running, and restart it if necessary.

---

- ❑ Stop all applications that use PowerPath; for example, unmount all file systems mounted on PowerPath pseudo devices.
- ❑ Ensure that you are logged in as root or that you have administrative privileges.

## Installing PowerPath

This section describes how to install PowerPath when no version of PowerPath has ever been installed on the host. The procedure for installing PowerPath is different depending on the host operating system.

If you are installing a service pack onto a host on which PowerPath is already installed, refer to [“PowerPath and MPxIO” on page 18](#).

Note the following when installing PowerPath:

- ◆ PowerPath Migration Enabler is installed by default when you install PowerPath.
- ◆ If you plan to install both PowerPath and VxVM, install PowerPath first.
- ◆ Note the following considerations if you are installing in an Oracle Solaris LDOM environment:
  - The PowerPath license key must be installed on the control domain.
  - PowerPath for Solaris multipathing software can be installed either in the control domain or in I/O domain.

EMC Knowledgebase Solution emc294681 provides supported Oracle Solaris LDOM configuration information.

- ◆ PowerPath supports offline installation.
- ◆ You can install PowerPath in single-user mode.

## Install PowerPath on Solaris 11.x

Install PowerPath on Solaris 11.x.

### Install PowerPath using Oracle Solaris IPS repositories

Solaris11.x uses Oracle Solaris Image Packaging System (IPS) packaging rather than SVR4 packaging. PowerPath for Solaris 11.x supports the Oracle Install, Upgrade, and Uninstall management service on Solaris 11.x using the Oracle Solaris Image Packaging System (IPS) and repositories.

## IPS repositories and commands

A repository can contain more than one package and can be a local, file-based repository or a remote web-based repository. Each PowerPath IPS package is a repository containing one package.

- ◆ Use the **pkg publisher** command to display a list of repositories.

The Solaris 11.x OS maintains a list of known repositories. The following is an example PowerPath package listing. For PowerPath packages, the publisher value is **emc.com**:

```
> pkg publisher
PUBLISHER TYPE STATUS URL
solaris origin online http://pkg.oracle.com/solaris/release/
emc.com origin online file://PowerPath/EMCpower/
```

To reduce installation time, create a local Solaris 11.x repository for installation. The Solaris 11.x repository is accessed by IPS scripts during installation. A remote repository and slow network connections impact the install time. Oracle Solaris repository at <http://pkg.oracle.com/solaris/release> provides more information. This is pertaining to OPT 382316.

The complete Oracle manual Copying and Creating Oracle Solaris 11.x Package Repositories is available at <http://docs.oracle.com>.

- ◆ Use the **pkg unset-publisher** command to remove all PowerPath repositories.

```
pkg unset-publisher emc.com
```

- ◆ Use the **pkg set-publisher** command to add a package to the list of known repositories.

```
pkg set-publisher -p <absolute_path_to_package>
```

When PowerPath for Solaris 11.x is installed, configuration scripts are run by the PowerPath configuration service, **system/emcpower/config**, after the package files are installed.

For more information on Solaris IPS installation, the Solaris 11.x Express Image Packaging System documents provide more information.

## Installing PowerPath using IPS repositories

This procedure is applicable to fresh installations and to full package installations of service packs and patches.

This installation procedure is not interactive. This procedure recovers saved configuration files and always installs into the **/opt** directory.

1. Untar the package into the **/var/spool/pkg** directory.
2. Use the **pkg install** command to install the package.

```
pkg install -g /var/spool/pkg/EMCpower system/EMCpower
```

---

**Note:** To install PowerPath without updating repository information which causes install delay, use **pkg install --no-refresh** command option. This is pertaining to OPT 382316.

---

When the **pkg** command finishes the PowerPath configuration may still be running.



If you receive an error after running the **pkg install** command, see [“bootadm update-archive error installation” on page 18](#).

3. Check the `/var/svc/log/system-emcpower-config:default.log` log file to find out if the installation was successful and if a reboot is required.

The log file is cumulative. The result of the latest installation is at the end of the file.

4. For PowerPath to manage the storage arrays over MPxIO, set the values of **mpxio-disable="yes"** in the **fp.conf** file and the **iscsi.conf** file.
5. Register the PowerPath license on the host.

## Installing PowerPath on alternate Boot environment

Use the **pkg install** command to install the package on an alternate Boot environment.

```
pkg install -g /var/spool/pkg/EMCpower --be-name <alternate boot environment name> system/EMCpower
```

## Installing PowerPath offline

This procedure is applicable to fresh installations and to full package installations of service packs and patches.

1. Untar the package into the `/var/spool/pkg` directory.
2. Use the **zpool import** command to import the root zpool of the image where PowerPath is to be installed.

```
zpool import -R /a <root zpool>
```

3. Set the zfs mount point.

```
zfs set mountpoint=<mountpoint location> <boot-path of the root zpool>
```

Where `<mountpoint location>` is any mount directory that is created and `<boot-path of the root pool>` is the path of the Solaris root zpool from where PowerPath is to be installed. For example:

```
zfs set mountpoint=/fixup rpool1/ROOT/solaris
```

4. Mount the root zpool.

```
zfs mount <boot-path of the root zpool>
```

5. Install the package.

```
pkg -R /fixup -g /var/spool/pkg/EMCpower system/EMCpower
```

If you receive an error after entering the **pkg -R** command, see [“bootadm update-archive error installation” on page 18](#).

6. Use the **zfs unmount /fixup** command to unmount the root zpool.
7. For PowerPath to manage the storage arrays over MPxIO, set the values of **mpxio-disable="yes"** in the `/etc/driver/drv/fp.conf` and `/etc/driver/drv/iscsi.conf` files.

## bootadm update-archive error installation

If you receive the following message:

```
PHASE ITEMS
Reading Existing Index 8/8
Indexing Packages 1/1
pkg: '/sbin/bootadm update-archive -R /' failed.
with a return code of 1.
```

It does not affect the PowerPath installation. It means the bootadm archive did not get updated during the PowerPath installation.

Run the following command after the **pkg** command finishes:

```
/sbin/bootadm update-archive
```

This is pertaining to issue 389977.

## PowerPath and MPxIO

This section discusses PowerPath and the Solaris native I/O multipathing software, MPxIO, interoperability, and array management.

MPxIO can be enabled or disabled on a host. When MPxIO is enabled, it can manage storage arrays implicitly or explicitly. Disabling MPxIO and rebooting moves the management of the arrays under PowerPath.

There are several Solaris configuration files used to control MPxIO behavior:

- ◆ **fp.conf**: value of `mpxio-disabled` (yes/no) controls MPxIO behavior for Fibre Channel devices. For Solaris 11.x: `/etc/driver/drv/fp.conf`
- ◆ **iscsi.conf**: value of `mpxio-disabled` (yes/no) controls MPxIO behavior for iSCSI devices. For Solaris 11.x use: `/etc/driver/drv/iscsi.conf`
- ◆ **scsi\_vhci.conf**: used to specify array classes explicitly managed by MPxIO.

MPxIO is enabled on the system when there is a `mpxio-disable=no` entry in the `fp.conf` and/or `iscsi.conf` files.

On x86 hosts, MPxIO also manages some array classes by default even if there are no entries in the `fp.conf` and `iscsi.conf` files.

## Arrays managed by MPxIO

MPxIO manages array classes defined in the `scsi_vhci.conf` file. Across versions of Solaris, MPxIO has varied behavior related to managing array classes by default. The Oracle documentation provides more details.

## Manage all array classes using PowerPath

- ◆ To use PowerPath to manage all array classes attached to the host, ensure that the flag- "mpxio-disable" is set to "yes" in both files `iscsi.conf` and `fp.conf`. For more information on the coexistence of PowerPath and MPxIO, refer to [“Coexistence with third-party multipathing software” on page 20](#).

- ◆ If the values of the flag "mpxio-disable" is "yes" in one file and "no" in the other file, change both values to "yes". Note that for this operation to take effect a reboot is required.

## Manage arrays using both PowerPath and MPxIO

On a host having multiple storage arrays connected, you can configure PowerPath to manage some array classes and configure MPxIO to manage other storage system by modifying `scsi_vhci.conf` file appropriately. For example, if the host has VNX, CLARiiON and XYZ arrays connected to it, PowerPath can manage the VNX and CLARiiON arrays; whereas, MPxIO can manage the XYZ array.

### NOTICE

For enabling or disabling multipath support by MPxIO on particular storage please refer to Oracle Solaris Documentation for MPxIO.

Continuing with the VNX, CLARiiON, and XYZ array example, perform the following procedure.

1. Modify the `scsi_vhci.conf` file to disable MPxIO for VNX and CLARiiON arrays and then enable MPxIO for the XYZ array. The Oracle documentation provides more information on Disabling Multipath Support for third party arrays.
2. Enable MPxIO in the `fp.conf` file. The Oracle Documentation provides more information.
3. Reboot the host as part of enabling MPxIO.
4. (optional) If PowerPath is not yet installed, install PowerPath. If PowerPath is installed, proceed to the next step.
5. Type **powermt display** to see if PowerPath is managing the VNX and CLARiiON arrays.
6. Check whether MPxIO is managing the XYZ array.

## Disable MPxIO on Solaris 11.x

For Solaris 11.x, the relevant System configuration files are in multiple system directories.

The default read-only conf files are:

- ◆ `/kernel/drv/fp.conf`
- ◆ `/kernel/drv/iscsi.conf`
- ◆ `/kernel/drv/scsi_vhci.conf`

The editable conf files are:

- ◆ `/etc/driver/drv/fp.conf`
- ◆ `/etc/driver/drv/iscsi.conf`
- ◆ `/etc/driver/drv/scsi_vhci.conf`

A fresh install of Solaris 11.x is by default set up to use MPxIO. The defaults can be seen in the **fp.conf** and **iscsi.conf** files located in **/kernel/drv**.

To override the defaults, copy the **fp.conf** and **iscsi.conf** files to **/etc/driver/drv** and edit them accordingly. Comment out all the lines in the editable conf file except for **mpxio-disable** and set it to **yes** or **no**, as detailed in [step 1 on page 20](#) and [step 2 on page 20](#).

The configuration files in **/etc/driver/drv** take precedence over the configuration files in **/kernel/drv**. The conf file in **/kernel/drv** is only used when the corresponding file is missing from **/etc/driver/drv**. Although the Solaris OS can distinguish between FC and iSCSI devices, PowerPath does not make this distinction for manage and unmanage. The **mpxio-disable** value must be set to **yes** in both the **fp.conf** and **iscsi.conf** files for PowerPath to manage storage arrays.

1. Disable MPxIO in the FC environment:
  - a. Copy **/kernel/drv/fp.conf** to **/etc/driver/drv/fp.conf**
  - b. Set the flag as follows:

```
mpxio-disable="yes"
```

Alternatively, disable MPxIO by HBA port on a Solaris host by typing the following command in the **/etc/driver/drv/fp.conf** file:

```
name="fp" parent="<parent_name>" port="<port-number>" mpxio-disable="yes"
```

2. Disable MPxIO in iSCSI environment:
  - a. Copy **/kernel/drv/iscsi.conf** to **/etc/driver/drv/iscsi.conf**
  - b. Set the **mpxio-disable** flag to **yes**.
 

```
mpxio-disable="yes"
```
3. Use the Solaris **stmsboot -d** command to disable Solaris I/O multipathing and then reboot as prompted.

## Remove or unmanage a device from PowerPath management

1. Type **powermt remove dev=<pseudo name>**.
2. Type **powermt unmanage class=<class>**.
3. Type **powermt config**.

## Best practices for PowerPath and MPxIO

Before installing PowerPath on a host, check the contents of **fp.conf**, **iscsi.conf**, and **scsi\_vhci.conf** files and adjust the entries for the desired manage/unmanage behavior. [“Manage all array classes using PowerPath” on page 18](#) provides more information.

## Coexistence with third-party multipathing software

PowerPath can coexist with the following third-party path management software:

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

- ◆ HP StorageWorks Auto Path XP Virtual Array
- ◆ Oracle Solaris MPxIO
- ◆ Veritas Volume Manager (VxVM) Dynamic MultiPathing (DMP)

However, PowerPath cannot co-manage devices with some third-party path management software. During installation of PowerPath, the installation script tests for the presence of third-party path management software on the system.

For all third-party path management software except MPxIO and DMP:

- ◆ When third-party path management software is installed on the system, the PowerPath installation script disables support for the corresponding array type in PowerPath.

On Solaris 11.x, the message is placed in the `system-emcpower-config:default.log` file. The message prompts you to ensure that devices are not managed by both MPxIO and PowerPath. PowerPath unmanages devices that are explicitly managed by MPxIO.

If this software is present	Support for this class is disabled in PowerPath
Hitachi HDLM	hitachi
IBM SDD	ess
Solaris MPxIO	Classes that would ordinarily handle the MPxIO-configured arrays

PowerPath co-management is not supported by third-party array path management products; therefore, when you enable PowerPath (**powermt manage class=<class>**) for a class that is being managed by a third-party product, multipathing behavior is undefined.

If this software is present or enabled	Multipathing behavior is undefined if you execute this command
Hitachi HDLM	<code>powermt manage class=hitachi</code>
IBM SDD	<code>powermt manage class=ess</code>
HP StorageWorksXP Auto Path	<code>powermt manage class=hpxp</code>
Solaris MPxIO	<code>powermt manage class=class</code> Where <i>class</i> is any class that handles an MPxIO-configured array

Since PowerPath and the third-party software cannot co-manage devices, do not initiate co-management by executing this command on third-party array class machines. As long as this command is not executed, support for the relevant third-party arrays will remain disabled across reboots.

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

Before you install or enable this software	Execute this command
Hitachi HDLM	<code>powermt unmanage class=hitachi</code>
IBM SDD	<code>powermt unmanage class=ess</code>
HP StorageWorks XP Auto Path	<code>powermt unmanage class=hpxp</code>

# CHAPTER 2

## PowerPath in a Cluster Environment

This chapter describes how to install and configure PowerPath in Solaris cluster environments.

- ◆ [PowerPath in an Oracle Solaris Cluster 4.2 .....](#) 24
- ◆ [PowerPath in a Veritas Cluster Server \(VCS\) cluster .....](#) 25

## PowerPath in an Oracle Solaris Cluster 4.2

This section describes how to:

- ◆ Install PowerPath and Oracle Solaris Cluster 4.2 on Solaris 11.x in a new cluster, that is, where neither the PowerPath nor the Oracle Solaris Cluster 4.2 software is installed on any host to be included in the cluster.
- ◆ Integrate PowerPath into an existing Oracle Solaris Cluster 4.2 cluster.

### Installing PowerPath in a new Oracle Solaris Cluster 4.2

1. Prepare the cluster environment. Refer to the relevant Sun Cluster documentation. In a Symmetrix, VNX OE, CLARiiON, XtremIO, or VPLEX environment, refer also to the *Host Connectivity Guide for Sun Solaris*, VNX OE Storage-System Support website, or the CLARiiON Storage-System Support website.
2. Install PowerPath on all nodes.
3. Install Oracle Solaris Cluster 4.2 on all nodes. Refer to the relevant Oracle Solaris Cluster documentation.

#### **IMPORTANT**

On x86 hosts connected to a Symmetrix storage system, ensure that the hosts in the cluster have unique host IDs (type **hostid** from each node in the cluster). Failure to use unique host IDs breaks the cluster disk-fencing scheme. Contact Oracle Solaris customer support for instructions on changing host IDs.

4. Configure the quorum device on all nodes.

**Note:** You cannot use gatekeepers and VCMDB devices as quorum devices.

5. Initialize PowerPath devices on all nodes, such as registering PowerPath and setting array policies.
6. Start cluster services on the master node.
7. Designate/create shared disk groups on the master node.
8. Create logical volumes from the designated shared disks.
9. Register the disk group.

### Installing PowerPath into an existing Oracle Solaris Cluster 4.2

Upgrades PowerPath into an existing Oracle Solaris Cluster 4.2 on Solaris 11.x are supported.

To install PowerPath into an existing Oracle Solaris Cluster 4.x, follow these steps on each cluster node, one node at a time.

1. Stop cluster services on the node.

```
reboot -- -x
```



2. Install or upgrade PowerPath on the node as described in [“Install PowerPath on Solaris 11.x” on page 15](#) Register PowerPath on the host before adding the node into the cluster.

#### **IMPORTANT**

On x86 hosts connected to a Symmetrix storage system, ensure that the hosts in the cluster have unique host IDs (type **hostid** from each node in the cluster). Failure to use unique host IDs breaks the cluster disk-fencing scheme. Contact Oracle Solaris customer support for instructions on changing host IDs.

3. Type **reboot** to start cluster services on the node.  
Wait for the node to be fully integrated into the cluster before continuing.
4. If your environment includes Hitachi Lightning, HP xp, HP StorageWorks EVA, or IBM ESS systems:
  - a. Wait for the node to join the cluster.
  - b. Run the following commands to create pseudo devices for the HDS, HP xp, HP StorageWorks EVA, and ESS devices:

```
powermt config
```

```
powermt save
```

- c. Verify that the pseudo devices have been created.

```
powermt display dev=all
```

## PowerPath in a Veritas Cluster Server (VCS) cluster

This section describes how to:

- ◆ Install PowerPath and VCS in a new cluster, that is, where neither the PowerPath nor the VCS software is installed on any host to be included in the cluster.
- ◆ Integrate PowerPath into an existing VCS cluster.

If emcpower devices exist in a VxVM disk group that is being added as a resource group to VCS, those devices should remain in the disk group as emcpower devices for as long as they comprise active VxVM logical volumes.

#### **Note:**

—On x86 hosts in a Symmetrix environment, before installing PowerPath in a VCS or Oracle Solaris Cluster, verify that each node has a unique host ID. Failure to use unique host IDs breaks the cluster disk-fencing scheme.

—When configuring the VCS cluster, set the vxio number in the name\_to\_major file to be the same across all the nodes of the VCS cluster. You can check the vxio number in the name\_to\_major file by running **cat etc/name\_to\_major | grep vxio**.

## Installing PowerPath in a new VCS cluster

To install PowerPath and VCS when neither PowerPath nor VCS is installed on any host:

1. On each host to be included in the cluster:

- a. Prepare the cluster hardware, making the necessary networking and disk connections among the hosts and the storage system. Refer to the relevant VCS documentation. In Symmetrix, VNX OE, or CLARiiON environments, refer also to the *Host Connectivity Guide for Oracle Solaris*, VNX OE Storage-System Support website, or the CLARiiON Storage-System Support website.
  - b. Use the Solaris **format** utility to verify that all storage system devices are seen by each host.
  - c. Install PowerPath and then verify that PowerPath can see all the devices.
  - d. Install any applications.
  - e. Install the VCS software, following the installation procedure described in the relevant Veritas Cluster Server documentation. Initially configure VCS to run without a service group.
2. On each node in the cluster, define the resources (for example, VxVM volumes) that make up the service group. (You will configure the service group in step 3.)

#### **IMPORTANT**

On x86 hosts connected to a Symmetrix storage system, ensure that the hosts in the cluster have unique host IDs (type **hostid** from each node in the cluster). Failure to use unique host IDs breaks the cluster disk-fencing scheme. Contact Oracle Solaris customer support for instructions on changing host IDs.

3. On one host in the cluster:
  - a. Configure the service group by adding the resources you defined in step 2 to the `/etc/VRTSvcs/conf/config/main.cf` file. The disk or logical device resources can be pseudo (for EMC devices only) or native devices.

---

**Note:** Using a disk for service group heartbeat instead of a network is subject to restrictions. Not all disks can be used. Consult the Veritas documentation.

---

  - b. Start cluster services on the host.
4. Start cluster services on each remaining node in the cluster.
5. On each node in the cluster:
  - a. Verify that the service group is up and running, and use either the VCS GUI or the **hagr** command to verify that the service group can successfully fail over to all hosts in the cluster.
  - b. Add other service groups as needed.

## Integrating/upgrading PowerPath into an existing VCS cluster

1. On each node in the cluster, stop cluster services on the node using the **hastop -local -evacuate** command.
2. On each node in the cluster, install or upgrade PowerPath on the node.

**IMPORTANT**

---

On x86 hosts connected to a Symmetrix storage system, ensure that the hosts in the cluster have unique host IDs (type **hostid** from each node in the cluster). Failure to use unique host IDs breaks the cluster disk-fencing scheme. Contact Oracle Solaris customer support for instructions on changing host IDs.

---

3. Run the **hastart** command to start cluster services on the node where you changed the main.cf file, and wait for the node to be fully reintegrated into the cluster.
4. Run **hastart** to start the remaining nodes in the cluster, waiting for each node to be fully integrated into the cluster before running **hastart** on the next node.



# CHAPTER 3

## Configuring a PowerPath Boot Device on Solaris

This chapter describes how to configure a PowerPath device as the boot device for a Solaris host and how to remove PowerPath control over a storage system boot device.

- ◆ Introduction..... 30
- ◆ Boot from PowerPath pseudo device on Solaris 11.x ZFS (SPARC) ..... 30
- ◆ Removing PowerPath control over a boot device on Solaris 11.x..... 31
- ◆ Moving the boot device from a native device to an emcpower device (UFS) ..... 32

## Introduction

On some storage systems, you can use either a PowerPath native device or an emcpower device as a boot device—the device that contains the startup image. Refer to the PowerPath release notes to see whether your storage system supports PowerPath boot devices.

Once the root is mounted, using a PowerPath device as the boot device provides load balancing and path failover for the boot device. Native devices, however, do not provide boot time boot path failover.

The following sections describe how to configure a PowerPath native device as the boot device and then move the boot device to an emcpower device.

### **CAUTION**

**When booting off an external storage system using an emcpower device, ensure that all device paths from the host to the storage system are connected and available for I/O at the time of boot. Also ensure that all physical connections (for example, hardware paths and switches) are in working order. Otherwise, the host may not boot.**

## Boot from PowerPath pseudo device on Solaris 11.x ZFS (SPARC)

Use the following procedure to boot from PowerPath pseudo device on ZFS (SPARC).

1. Choose a LUN the same size or larger than the internal disk in the rpool.
2. Attach the boot LUN (pseudo) to the existing rpool.

```
# zpool attach rpool c0t5000CCA025765610d0s0 emcpower1a
```

Make sure you wait until resilver is done before rebooting.

```
-bash-4.1# zpool status rpool
pool: rpool
state: ONLINE
scan: resilvered 26.8G in 0h6m with 0 errors on Wed May 6 08:28:32
2015
config:
```

NAME	STATE	READ	WRITE	CKSUM
rpool	ONLINE	0	0	0
mirror-0	ONLINE	0	0	0
c0t5000CCA025765610d0s0	ONLINE	0	0	0
emcpower1a	ONLINE	0	0	0

```
errors: No known data errors
```

3. After resilvering is done, specify the pseudo device as the boot device.
  - a. Get the ls -l output of all four native devices of the boot LUN:

```
# powermt display dev=1
Pseudo name=emcpower1a
Symmetrix ID=000194901016
Logical device ID=0EBE
```

```

state=alive; policy=SymmOpt; queued-IOS=0
=====
----- Host ----- - Stor - -- I/O Path -- -- Stats ---
### HW Path          I/O Paths   Interf.  Mode   State  Q-IOS Errors
=====
3078 pci@400/pci@2/pci@0/pci@8/SUNW,qlc@0,1/fp@0,0 c11t50000972C00FE15Dd10s0 FA 8fB active
alive      0      0
3078 pci@400/pci@2/pci@0/pci@8/SUNW,qlc@0,1/fp@0,0 c11t50000972C00FE159d10s0 FA 7fB active
alive      0      0
3076 pci@400/pci@2/pci@0/pci@8/SUNW,qlc@0/fp@0,0 c10t50000972C00FE15Dd10s0 FA 8fB active
alive      0      0
3076 pci@400/pci@2/pci@0/pci@8/SUNW,qlc@0/fp@0,0 c10t50000972C00FE159d10s0 FA 7fB active
alive      0      0

# ls -l /dev/dsk/c11t50000972C00FE15Dd10s0
lrwxrwxrwx 1 root root      85 Sep  8 2014 /dev/dsk/c11t50000972C00FE15Dd10s0 ->
../../../../devices/pci@400/pci@2/pci@0/pci@8/SUNW,qlc@0,1/fp@0,0/ssd@w50000972c00fe15d,a:a
# ls -l /dev/dsk/c11t50000972C00FE159d10s0
lrwxrwxrwx 1 root root      85 Sep  8 2014 /dev/dsk/c11t50000972C00FE159d10s0 ->
../../../../devices/pci@400/pci@2/pci@0/pci@8/SUNW,qlc@0,1/fp@0,0/ssd@w50000972c00fe159,a:a
# ls -l /dev/dsk/c10t50000972C00FE15Dd10s0
lrwxrwxrwx 1 root root      83 Sep  8 2014 /dev/dsk/c10t50000972C00FE15Dd10s0 ->
../../../../devices/pci@400/pci@2/pci@0/pci@8/SUNW,qlc@0/fp@0,0/ssd@w50000972c00fe15d,a:a
# ls -l /dev/dsk/c10t50000972C00FE159d10s0
lrwxrwxrwx 1 root root      83 Sep  8 2014 /dev/dsk/c10t50000972C00FE159d10s0 ->
../../../../devices/pci@400/pci@2/pci@0/pci@8/SUNW,qlc@0/fp@0,0/ssd@w50000972c00fe159,a:a

```

b. set boot path using eeprom:

```

eeprom boot-device=
/pci@400/pci@2/pci@0/pci@8/SUNW,qlc@0,1/fp@0,0/disk@w50000972c00
fe15d,a:a,raw (replace 'ssd' with 'disk')

```

c. Setting multiple paths as boot-device for failover:

```
eeprom boot-device="disk2 disk3 disk4 disk5"
```

```

(Setting nvram parameter:
eeprom "use-nvramrc?=true"
eeprom "nvramrc=devalias disk5
/pci@400/pci@2/pci@0/pci@8/SUNW,qlc@0/fp@0,0/disk@w50060162086035e9,7:a devalias disk4
/pci@400/pci@2/pci@0/pci@8/SUNW,qlc@0/fp@0,0/disk@w5006016a086035e9,7:a devalias disk3
/pci@400/pci@2/pci@0/pci@8/SUNW,qlc@0,1/fp@0,0/disk@w50060163086035e9,7:a devalias disk2
/pci@400/pci@2/pci@0/pci@8/SUNW,qlc@0,1/fp@0,0/disk@w5006016b086035e9,7:a"
eeprom boot-device="disk2 disk3 disk4 disk5"
)

```

4. Reboot.

5. Verify that the system booted with pseudo device:

```

# prtconf -vp | grep bootpath
bootpath:  '/pci@400/pci@2/pci@0/pci@8/SUN

```

**Note:** Refer to the Oracle documentation for BFS on Solaris 11.x x86 or Opteron (ZFS) for additional information.

## Removing PowerPath control over a boot device on Solaris 11.x

1. Detach pseudo device from rpool, for example:

```
# zpool detach rpool emcpower1a
```

```
# zpool status rpool
```

2. Set local device as a boot device, for example:

```
# eeprom
boot-device="boot-device=/pci@400/pci@1/pci@0/pci@4/scsi@0/disk@w50
00cca025765611,0:a
```

## Moving the boot device from a native device to an emcpower device (UFS)

The procedures in this section applies to SPARC platform.

**Note:** Do not boot from an emcpower device on an unsupported HBA. For information on HBAs that are supported with Symmetrix, VNX OE, and CLARiiON storage systems, refer to the E-Lab Interoperability Navigator. For information on HBAs supported with other storage systems, refer to the appropriate documentation from your vendor.

1. Boot from the native device.
2. Install PowerPath as described in [Chapter 1, “Installing PowerPath.”](#)
3. Locate the native device from which you are booting, and correlate this device to an emcpower device.

```
/etc/powermt display dev=all
```

Output similar to the following appears:

```
Pseudo name=emcpower0a
Symmetrix ID=000000005543
Logical device ID=0001
state=alive; policy=SymmOpt; priority=0; queued-I/Os=0
=====
----- Host -----          - Stor - - I/O Path - - Stats ---
### HW Path                I/O Paths   Interf.  Mode   State Q-I/Os Errors
=====
1281 pci@4/QLGC,qla@1      c3t0d0s0   FA 3aA   active alive  0 0
1283 pci@6/QLGC,qla@1      c4t0d0s0   FA 3aA   active alive  0 0
```

Scroll through the output until you locate the native device used as the boot device; for example, **c3t0d0**. In this example, this native device corresponds to **emcpower0a**.

4. Identify the device node that corresponds to the emcpower device.

```
ls -l /dev/dsk/emcpower0a
```

Output similar to the following appears:

```
lrwxrwxrwx 1 root other 33 May 30 17:42 /dev/dsk/emcpower0a ->
../../devices/pseudo/emcp@0:a,blk
```

Looking at the output, you can see that **/pseudo/emcp@0:a,blk** corresponds to **emcpower0a**. You will use this value in step 6.

5. Make backup copies of the **/etc/system** and **/etc/vfstab** files so that, if necessary, you can restore the host to its pre-PowerPath settings.

```
cp /etc/system /etc/system.pre_EMCPower
```

```
cp /etc/vfstab /etc/vfstab.pre_EMCPower
```



- Using a text editor such as **vi**, add the following line above the **forceload: drv/emcp** statement:

```
rootdev: /pseudo/emcp@0:a,blk
```

The `/etc/system` file now includes the following lines:

```
forceload: drv/sd
rootdev: /pseudo/emcp@0:a,blk
forceload: drv/emcp
```

- Using a text editor such as **vi**, edit the `/etc/vfstab` file, replacing each native partition (`c#t#d#s#`) for the boot device with an `emcpower` partition name. In this example, you would replace `c3t0d0s0` with `emcpower0a`. You must change both the `/dev/dsk` and `/dev/rdisk` entries.
- Type **reboot** to restart the host.

## Changing the dump device path to the pseudo device in an R1/R2 boot failover setup

- While booted off the pseudo path, run **dumpadm** to check the dump device path.
- If the dump device is being configured with the native path, change the path to pseudo by typing **dumpadm -d /dev/dsk/emcpower<N>b**.

## Recovery procedure

If you cannot boot the host after configuring the PowerPath device as the boot device, you may have made a typing error when editing the `/etc/system` and `/etc/vfstab` files. To recover:

- Insert the Solaris Operating System CD into the host's CD-ROM drive.
- At the **OK** prompt, type **boot cdrom -s**
- Mount the storage system boot device that is experiencing the problem. For example, type **mount /dev/dsk/c3t0d0s0 /a**
- Type **TERM=sun-cmd export TERM**
- Check the `/etc/system` and `/etc/vfstab` files against the changes you made to these files when you set up multipathing to the storage system boot device. Use a text editor such as **vi** to correct any problems you find.
- Type **reboot -- -r** to restart the host.



# CHAPTER 4

## Upgrading PowerPath

This chapter describes how to upgrade PowerPath on a Solaris host.

- ◆ [Before you upgrade.....](#) 36
- ◆ [Upgrading PowerPath on Solaris 11 .....](#) 36

## Before you upgrade

This section describes how to upgrade PowerPath from an earlier version of PowerPath.

- ◆ Download the latest version of the EMC Grab utilities, available on EMC Online Support, and then run the PowerPath Configuration Checker.
- ◆ When you upgrade PowerPath, customized PowerPath settings are preserved (except write throttle queue settings).
- ◆ You must remove the earlier version of PowerPath before upgrading to the latest PowerPath version. When you remove PowerPath, follow the instructions in the documentation for that version. Note, however, that you do not need to reboot the host after you remove the earlier version.

---

**Note:** In a cluster, you must follow the cluster upgrade procedure, which includes booting the host out of the cluster before you upgrade.

---

- ◆ The upgrade preserves your existing PowerPath license. You do not need to retype the PowerPath license information.
- ◆ You must shut down all applications before you upgrade PowerPath. If the host boots off a PowerPath pseudo device, you must reboot after the upgrade.
- ◆ You cannot upgrade from PowerPath on a SPARC host to PowerPath on an x86 host or the reverse. You can only upgrade from SPARC to SPARC or from x86 to x86.
- ◆ Root volume mirrors need to be broken before performing an installation or upgrade of PowerPath. Once installation is complete, the root volume mirror can be synchronized again.

## Upgrading PowerPath on Solaris 11

This section describes how to upgrade PowerPath on Solaris 11 hosts.

### Upgrading PowerPath on Solaris 11

1. Untar the PowerPath bundle into the `/var/spool/pkg` directory.
2. Use the **pkg update** command to upgrade the PowerPath package to the latest.

```
# pkg update -g /var/spool/pkg/EMCpower system/EMCpower
```

### Upgrading PowerPath on Boot from SAN

1. Run **pkg update** command with options to create new boot environment for PowerPath upgrade on a new boot environment. For example:
 

```
# pkg update -v -g ./EMCpower --be-name BE_Name
```
2. Type **beadm list** to check the updated boot environment list. The `BE_Name` should display R for Active.
3. Type **init 6** to reboot the host and boot to the alternate boot environment.

# CHAPTER 5

## Migrating to PowerPath

This chapter describes how to migrate to PowerPath from other multipathing software without loss of data.

- ◆ [Migrating from HP StorageWorks Secure Path .....](#) 38
- ◆ [Migrating from IBM Subsystem Device Driver \(SDD\).....](#) 39
- ◆ [Migrating from Hitachi Dynamic Link Manager \(HDLM\) .....](#) 40

## Migrating from HP StorageWorks Secure Path

This section describes how to migrate to PowerPath from Secure Path.

### Secure Path operation on Solaris

Secure Path replaces all native `ctd` devices belonging to a particular EVA or HSG80 LUN (there is one such native device for each path) with one unique `ctd` Secure Path pseudo device. The original native devices are not usable as long as Secure Path is installed. The pseudo devices can be initialized and added to Veritas disk groups, and volumes can be built on them.

### Migrating from Secure Path to PowerPath

1. Stop I/O activity from the host to the Secure Path pseudo devices:
  - a. Unmount any layered file systems.
  - b. In `/etc/vfstab`, comment out any file system entries that directly use Secure Path pseudo names. (You need not comment out file system entries that use Veritas volume names.)
2. Install PowerPath.
 

PowerPath will detect an existing Secure Path installation, causing HP HSx devices to come up unmanaged.
3. Uninstall Secure Path.
 

---

**Note:** At this point, the Secure Path driver, `hsx`, is still running. `modinfo` will show this driver. You cannot unload the driver using `modunload`.

---
4. Ensure that PowerPath will manage the HP HSx devices upon the next reboot.
 

```
powermt manage class=hphsx
```
5. Ensure that HBA driver configuration files and `sd.conf` are in good condition in `/kernel/drv`.
 

Secure Path edits these files during installation, and may not return them to their original condition during deinstallation. We suggest that you:

  - Use the PowerPath default settings for the HBA driver configuration file (`lpfc.conf`, `fcaw.conf`, `fca-pci.conf`, `qla2200.conf`, or `qla2300.conf`).
  - Ensure the `sd.conf` file has LUN number entries for EVA or HSG80 targets.
6. Type `reboot -- -r` to reboot the host with device reconfiguration.
 

During reboot, Veritas recovery manager will rediscover volumes on native `ctd` devices (instead of Secure Path pseudo `ctd` devices).
7. Ensure that PowerPath will now handle multipathing to the StorageWorks devices and the volumes build on them:
  - a. Run the `vxdiskadm` utility, suppressing all but one path to every EVA or HSG80 device from the VxVM point of view.

- b. If, for performance, LUNs were originally distributed *preferred* across both EVA or HSG80 controllers, run **powermt restore**.
  - c. Run **powermt display dev=all class=hpsx** and verify that PowerPath can access every path to every StorageWorks device.
8. Edit `/etc/vfstab` for those file systems that formerly used Secure Path pseudo devices directly (that is, with no volume manager): Add the corresponding `emcpower` device names, then **mount -a**.

---

**Note:** This procedure works for all types of Veritas volumes, whether created directly in `rootdg` or in some other user-created disk group.

---

## Migrating from IBM Subsystem Device Driver (SDD)

This section describes how to migrate to PowerPath from SDD.

### SDD operation on Solaris

For each ESS LUN on Solaris, there are typically several native `cXt YdZ` devices in `/dev/dsk` and `/dev/rdisk`—one native device per path to the LUN. SDD adds one extra pseudo device, `vdiskN`, for each unique ESS LUN.

Note that *all* these device entries are usable for any given LUN:

- ◆ For I/O issued directly to the native `cXt YdZ` device, SDD does no multipathing.
- ◆ For I/O issued to the `vdiskN` device, SDD does multipathing over all available native paths to the LUN.

Veritas can recognize only the native `cXt YdZ` devices. SDD does not provide any install-time utilities to make its pseudo devices visible to Veritas. Thus, an ESS LUN is typically configured within Veritas using one native `cXt YdZ` path. Since Veritas/DMP does support ESS arrays, DMP can provide multipathing for ESS LUNs.

### Migrating from SDD to PowerPath

1. Stop I/O activity from the host to the **sdd** pseudo devices:
  - Unmount any layered file systems.
  - In `/etc/vfstab`, comment out any file system entries that directly use the pseudo names.
2. Install PowerPath.
 

PowerPath will detect an existing **sdd** installation, causing only the IBM ESS devices to come up unmanaged.
3. Uninstall **sdd**, which requires a shutdown and reboot.
4. Ensure that PowerPath will now handle multipathing:
  - a. Run the **vxdiskadm** utility, suppressing all but one path to every ESS device from the VxVM point of view.
  - b. Run **powermt manage class=ess**.

- c. Run **powermt display dev=all class=ess** to verify that PowerPath can see all paths to every ESS LUN.
5. Edit `/etc/vfstab` for those file systems that formerly used **sdd** pseudo devices directly: Add the equivalent emcpower device names, then **mount -a**.

## Migrating from Hitachi Dynamic Link Manager (HDLM)

This section describes how to migrate to PowerPath from HDLM.

### HDLM operation on Solaris

HDLM replaces all native `cXtYdZ` devices belonging to a particular Hitachi Lightning LUN (there is one such native device for each path) with one unique `cRQdZ` device. The native devices are not usable as long as HDLM is installed.

In addition, HDLM provides a script that can be run to create `ChRQdZ` entries in `/dev/vx/dmp` and `/dev/vx/rdmp`. The HDLM documentation describes how to add these `Ch` pseudo devices to a Veritas disk group. These devices can be added as *simple* disks only (not *sliced*), within Veritas. If you use this feature, we recommend that you back up the HDLM-based volumes before you uninstall HDLM, then restore them to PowerPath-based volumes later. Veritas recovery manager does not automatically recover the volumes from the underlying `cXtYdZ` device after HDLM multipathing is uninstalled.

### Migrating from HDLM to PowerPath

1. Stop I/O activity from the host to the HDLM pseudo devices:
  - a. Unmount any layered file systems.
  - b. In `/etc/vfstab`, comment out any file system entries that directly use the pseudo names.
2. Install PowerPath.
 

PowerPath will detect an existing HDLM installation, causing Hitachi Lightning devices to come up unmanaged.
3. Uninstall HDLM, which requires a reboot with device reconfiguration. Type **reboot -- -r** to reboot.
4. Ensure that PowerPath will now handle multipathing for all Hitachi Lightning devices:
  - a. Run the **vxdiskadm** utility, suppressing all but one path to every Hitachi Lightning device from the VxVM point of view.
  - b. Run **powermt manage class=hitachi**.
  - c. Run **powermt display dev=all class=hitachi** to verify that PowerPath can now see all paths to every Hitachi Lightning LUN.
5. Edit `/etc/vfstab` for those file systems that formerly used HDLM pseudo devices directly: Add the equivalent native `cXtYdZ` device names, then **mount -a**.



# CHAPTER 6

## Removing PowerPath

This chapter describes how to remove PowerPath from a Solaris host.

- ◆ Before you remove PowerPath..... 42
- ◆ Removing PowerPath on Solaris 11.x ..... 42
- ◆ PowerPath interoperability with ZFS file systems..... 44
- ◆ After you remove PowerPath ..... 44

## Before you remove PowerPath

Before you remove PowerPath from the host:

- ❑ Check the EMC Online Support site for the most current information. EMC updates the PowerPath release notes periodically and posts them on the EMC Online Support site.
- ❑ Stop all applications and databases using emcpower devices.
- ❑ Cleanup migrations using **powermig cleanup -handle <handle>**.
- ❑ Make sure no pseudo devices are in use. Unmount any mounted file systems contained on pseudo devices. Disable VxVM or Sun Volume Manage volumes on pseudo devices.
- ❑ If you have a database partition, discontinue use of the PowerPath devices as follows:
  1. Stop the database manager.
  2. Unmount PowerPath devices.
  3. Edit the appropriate database configuration files so they no longer refer to emcpower devices.
- ❑ If your PowerPath installation uses a storage system device as the boot device, remove PowerPath control over the boot device.
- ❑ If you are upgrading to a later version of PowerPath, save your configuration by running **powermt save**.
- ❑ Ensure that you are logged in as root or that you have administrative privileges.

## Removing PowerPath on Solaris 11.x

Use the following procedure to remove PowerPath on Solaris 11.x.

**Note:** **pkgremove** is the recommended procedure on Solaris 11.x when uninstalling on current Boot Environment

Because of an Oracle IPS limitation, using the **pkg uninstall** command might not exit if Migration Enabler migration handles are still present and devices are in use.

1. Determine if PowerPath is installed on host.

```
# pkg info EMCpower

root@lclg082:~# pkg info EMCpower
Name: system/EMCpower
Summary: EMC Powerpath Multipathing, version <>
Description: EMC Powerpath Multipathing, version <>
State: Installed
Publisher: emc.com
Version: <version>
Build Release: 5.11
Branch: 449
Packaging Date: November  8, 2011 05:28:11 PM
Size: 124.50 MB
FMRI:
pkg://emc.com/system/EMCpower@<version>,5.11-449:20111108T172811Z
```

2. Remove PowerPath on Solaris 11.x.

```
/opt/EMCpower/install/pkgremove
```

The command result will be displayed on the screen, and not stored in a log file.

For more information on Solaris IPS removal, go to [docs.oracle.com](http://docs.oracle.com) and search for Solaris 11.x Express Image Packaging System documents.

## Removing PowerPath on Boot from SAN (SPARC)

1. Remove the pseudo device from rpool and then set the internal disk as a boot-device.
  - a. Detach pseudo device from rpool.

```
-bash-4.1# zpool status rpool
pool: rpool
state: ONLINE
scan: resilvered 125M in 0h0m with 0 errors on Wed May  6 08:40:05
2015
config:

          NAME                                STATE      READ  WRITE CKSUM
          rpool                                ONLINE      0     0     0
          c0t5000CCA025765610d0s0            ONLINE      0     0     0

errors: No known data errors
```

- b. Set local device as a boot device.

2. Uninstall PowerPath.

```
# pkg uninstall EMCpower
Packages to remove: 1
Create boot environment: No
Create backup boot environment: No
Services to change: 1

PHASE                                ITEMS
Removing old actions                  368/368
Updating package state database        Done
Updating package cache                 1/1
Updating image state                   Done
Creating fast lookup database          Done
```

The following unexpected or editable files and directories were salvaged while executing the requested package operation; they have been moved to the displayed location in the image:

```
var/svc/manifest/system/emcpower ->
/var/pkg/lost+found/var/svc/manifest/system/emcpower-20150506T09
1850Z
opt/EMCpower/install ->
/var/pkg/lost+found/opt/EMCpower/install-20150506T091850Z
opt/EMCpower/bin ->
/var/pkg/lost+found/opt/EMCpower/bin-20150506T091850Z
opt/EMCpower ->
/var/pkg/lost+found/opt/EMCpower-20150506T091850Z
etc/emc/ppme ->
/var/pkg/lost+found/etc/emc/ppme-20150506T091850Z
etc/emc/bin -> /var/pkg/lost+found/etc/emc/bin-20150506T091850Z
etc/emc -> /var/pkg/lost+found/etc/emc-20150506T091850Z
```

3. Reboot the host.

4. Verify that the host booted from internal disk:

```
# prtconf -vp | grep bootpath
```

## Removing PowerPath on an alternate boot environment

1. Use the **pkg uninstall** command to uninstall an existing package.

```
pkg uninstall -be-name <alternate be-name> system/EMCpower
```

Where <alternate be-name> is the alternate boot environment.

2. Type **beadm list** to check the updated Boot Environment list. The *BE\_name* should display **R** for **Active**.
3. Type **init 6** to reboot the host and boot to the alternate boot environment.

## Removing PowerPath offline

1. Mount the boot environment, from which you need to remove PowerPath.

```
beadm mount solaris-be /mount
```

2. Uninstall PowerPath from mounted boot environment.

```
pkg -R /mount uninstall system/EMCpower
```

3. Clean up the configuration on the given boot environment.

```
/mount/etc/emcp_cleanup -R /mount
```

## PowerPath interoperability with ZFS file systems

For all versions of PowerPath, export zpools before uninstalling or upgrading PowerPath. Additional steps may be required for some PowerPath versions, as described in the procedures that follow.

### Removing PowerPath with existing zpools

1. Use the **zpool export** command to export the zpools.
2. Remove PowerPath.
3. Reboot the host (required only if VxVM is installed).
4. Use the **zpool import** command to import the zpools.

## After you remove PowerPath

After you remove PowerPath, you may need to:

- Remove PowerPath configuration and registration files.
- Reboot the host.

### Removing PowerPath files

The removal process saves the following files, adding the extension **.55.saved**:

- ◆ /kernel/drv/emcp.conf
- ◆ /kernel/drv/emcpsf.conf
- ◆ /etc/emc/mpaa.lams
- ◆ /etc/emc/mpaa.excluded
- ◆ /etc/emcp\_registration
- ◆ /etc/emcp\_devicesDB.dat
- ◆ /etc/emcp\_devicesDB.idx
- ◆ /etc/powermt\_custom.xml

If the removal program detects existing files with the **.55.saved** extension on the host, it overwrites these files. It does not overwrite saved files that have a different format, for example, **powermt.custom.saved**.

*If you are removing PowerPath from the host entirely* (that is, you are not planning to re-install PowerPath), use `emcp_cleanup` to remove all these saved files from the host. `Emc_cleanup` is invoked with one argument, either a PowerPath revision number, without periods, to remove the files saved from a specific PowerPath revision or a quoted asterisk to remove files saved from any PowerPath revision.

1. Type `/etc/emcp_cleanup "*"` to remove the PowerPath license and all supporting files. The argument `"*"` is used to signify saved files for all PowerPath revisions.
2. Type **y** to continue and then press **Enter**.
3. Type **y** to confirm uninstall and then press **Enter**.
4. After running `emcp_cleanup` and verifying no files remain, manually delete `emcp_cleanup`.

If saved files remain on the host

`emcp_cleanup "*"` removes all saved files. However, if after running `emcp_cleanup "*"` you discover saved files remaining on the host, you can run `emcp_cleanup` again, specifying a revision number.

For example, if after removing PowerPath and running `emcp_cleanup "*"`, you see saved files ending in **53.saved**, type `emcp_cleanup 53`

After removing all files with `emcp_cleanup`, manually delete `/etc/emcp_cleanup`.

## Rebooting the host

You need not reboot the host after you remove PowerPath if:

- ◆ The message at the end of the removal process indicates that reboot is not necessary.
- ◆ You plan to reinstall PowerPath.
- ◆ You plan to upgrade to a later version of PowerPath.

Otherwise, reboot the host by typing `reboot -- -r`



# CHAPTER 7

## PowerPath Administration on Solaris

This chapter discusses PowerPath issues and administrative tasks specific to Solaris.

◆ PowerPath Viewer and Management Component .....	48
◆ PowerPath and VNX OE and CLARiiON storage systems .....	51
◆ Ensuring a sufficient stack size .....	51
◆ Booting a host with built-in Fibre Channel ports .....	51
◆ Rebooting and custom settings .....	52
◆ R1/R2 boot failover support .....	52
◆ Device naming .....	54
◆ Reconfiguring PowerPath devices online .....	57
◆ Dynamic Reconfiguration .....	58
◆ Managing and unmanaging Celerra iSCSI devices .....	60
◆ Supporting Veritas Volume Manager .....	60
◆ Exporting emcpower devices to Non-Global Zones .....	62
◆ powercfc configuration utility .....	63

## PowerPath Viewer and Management Component

PowerPath Viewer is a utility that provides centralized, remote monitoring of your PowerPath-managed storage environment. PowerPath Viewer comprises two main components: PowerPath Viewer Console (or, the *Console*) and PowerPath Management Component (or, the *Management Component*). On Solaris, the Management component is part of PowerPath Monitor Daemon (emcp\_mond).

On PowerPath for Solaris, the Management Component is included in the PowerPath software download available on EMC Online Support. The *PowerPath Viewer Installation and Administration Guide*, available on EMC Online Support, provides information on configuring the Management Component.

PowerPath Viewer Console is available as a separate download in the PowerPath software downloads section of EMC Online Support. The *PowerPath Viewer Release Notes*, available on EMC Online Support, provides additional information.

## Event monitoring daemon using SNMP

The PowerPath event monitoring daemon monitors specific PowerPath events and sends an SNMP (Simple Network Management Protocol) trap when access to devices is disrupted. The configuration file for the event monitoring daemon specifies the events to monitor and a filter parameter for each event. This way, only those events needing intervention from an administrator are displayed.

An SNMP trap is issued for events listed in the configuration file, `/etc/emc/emcp_mond.conf`. The traps are sent to a master SNMP agent running locally. The master agent forwards the trap to a remote network manager where an administrator can view them and take corrective action as necessary.

## Monitored events

The PowerPath event monitoring daemon monitors the following events:

---

**Note:** All Paths Dead Management daemon traps are generated when there is I/O running and as a result of periodic path tests. Email in log files contains more information.

---

- ◆ Path Dead — One path to a device goes down.
- ◆ Path Alive — A dead path becomes operational.
- ◆ All Paths Dead — All paths to a device go down.
- ◆ Not all Paths Dead — At least one path to a device with all paths dead becomes operational.
- ◆ Path Latency Threshold Exceeded — The time to complete an I/O operation on a path exceeds a preset time.



When a monitored event occurs on a path, a timer is started. After a given delay time elapses, a trap is generated to report the event. If a 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. [Table 1 on page 49](#) shows each event, its canceling event and the default delay time for a monitored event.

**Table 1** Event monitoring daemon monitored events

Event name	Event ID	Event description in trap packet	Cancelling event	Default delay time
Path Dead	0x2	Path is dead	Path alive	60 seconds
Path Alive	0x4	Path is alive	Path dead	60 seconds
All Paths Dead	0x20	All paths dead	Not all paths dead	5 minutes
Not all Paths Dead	0x40	All paths not dead	All paths dead	5 minutes
Latency Threshold Exceeded	0x80000	Latency threshold exceeded	N/A	N/A

The configuration file for the event monitoring daemon contains comment lines, beginning with a hash mark, and event lines, containing the ID of each event to monitor, the ID of its cancelling event, and the delay time in seconds for the event. These three event parameters are delimited by commas. For example, the line for the path dead event is 0x2,0x4,60.

## Configuring the event monitoring daemon

After PowerPath is installed on the host being managed, configure the event monitoring daemon to send traps to the SNMP manager.

Configuration is performed in two steps:

- ◆ Configure the event monitoring daemon and `emcp_mond.conf` to select the events to be monitored.
- ◆ Configure SNMP files to enable trap generation.

Configure each host on which path monitoring is desired.

### Step 1. Configuring the event monitoring daemon

1. Go to `emcp_mond_edit` in `/etc/emc/bin` to start the event monitoring daemon editor.
2. To make configuration changes:
  - a. Type the number options to make the desired configuration changes.
  - b. Type the number option to confirm `emcp_mond.conf` and save.
3. Type the number option to exit, save, and restart the event monitoring daemon. If the Latency Threshold Exceeded event has been configured, enable the kernel to report path latency events as follows:
  - a. Run `powermt set path_latency_monitor=on` to enable generation of latency events.

- b. Run **powermt set path\_latency\_threshold=<seconds>** to set the time in seconds that must be exceeded to generate a latency event.
4. Configure and restart the master SNMP daemon as described in [“Step 2. Configuring the master SNMP files” on page 50.](#)

## Step 2. Configuring the master SNMP files

1. Copy the /opt/EMCpower/scripts/snmpd.local.conf file to /etc/sma/snmp.
2. Uncomment the **trap2sink** line and fill in the IP address or host name of the machine to receive traps generated by emcp\_mond.

```
trap2sink <Trap_destination> <Trap_destination_port#>
```

where:

- <Trap\_destination> is the IP address or host name of the host where the SNMP manager is installed.
- <Trap\_destination\_port#> is 162. This is the default port number of the host where the SNMP manager is listening.

For example:

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

3. On the Solaris 11.x system, determine if the master SNMP daemon is running.

```
svcs <master_snmp_service>,
```

where *master\_snmp\_service* is **/application/management/net-snmp**

If the service is running, restart it by typing:

```
svcadm restart -s <master_snmp_service>
```

If the service is not running, start it by typing:

```
svcadm enable -s <master_snmp_service>
```

## Trap packet format

The trap generated from a path event has SNMP three variable bindings shown in the example below.

OID	Parameter	Example
PowerPath.1.1.2	event description	Path is dead
PowerPath.1.1.3	event time	Fri Jan 8 11:05:19 2010
PowerPath.1.1.9	path	BUS=3073 TGT=5006048ACB1B9D93 LUN=31 ARRAYID=000187461238

In the preceding example, PowerPath stands for 1.3.6.1.4.1.1139.12. The event description for each event is described in [“Monitored events” on page 48.](#)

## PowerPath and VNX OE and CLARiiON storage systems

PowerPath events or actions that entail LUN trespasses (for example, SP failovers or **powermt restore** commands) can cause the Solaris disk driver to log warning and/or error messages. You can ignore these messages, as PowerPath intercepts them and hides them from the application sending the I/O.

For more information on VNX OE and CLARiiON configuration requirements, refer to the *Host Connectivity Guide for Sun Solaris*, available on the EMC Online Support site and the VNX OE and CLARiiON Storage-System Support websites. The *Host Connectivity Guide for Sun Solaris* describes how to edit the `/kernel/drv/sd.conf` file to suppress ODS device overlap error messages.

## Ensuring a sufficient stack size

PowerPath requires a stack size of at least 0x6000 for SPARC hosts. For x86 hosts, the required stack is 0x5000. To ensure an adequate stack size, the value of several kernel parameters in `/etc/system` are checked when you install PowerPath. If a parameter is missing, or has a value below the minimum required by PowerPath, a line that sets an acceptable value is added to the end of `/etc/system`. If `/etc/system` already contains a line that sets the parameter to a value greater than PowerPath's required minimum, then no change is made.

If you subsequently install another application that resets the stack size (for example, VxVM or VxFS), ensure that the appropriate kernel parameters are not set to values below PowerPath's required minimum. Other applications may reset the parameters to values that are insufficient for PowerPath.

If any kernel parameter in `/etc/system` contains an invalid value, Solaris reverts to its default value. Examples of invalid values include one that is larger than the system-specific maximum, or one that is not aligned to a page boundary.

A stack size that is too small for PowerPath can cause a stack overflow and kernel panic. Thus, it is important to ensure that any modifications to kernel parameters in `/etc/system` both are valid and meet the minimum size requirement for PowerPath.

## Booting a host with built-in Fibre Channel ports

If you boot a Solaris host with all host adapters to storage system volumes disconnected or dysfunctional, PowerPath will not configure any host adapter paths. After physically restoring the connections, run the following commands to restore the paths in PowerPath:

```
devfsadm
powercf -q
powermt config
```

---

**Note:** Running `powercf -q` is no longer necessary; the command is included for backward compatibility only.

---

## Rebooting and custom settings

On every reboot, all saved custom files (`powermt_custom.xml.[0-2]`) are incremented by one version. Thus, at any time, the custom configurations from the last three reboots are available.

You can ascertain from the custom file timestamps which version of the custom file contains the last valid settings saved prior to a reboot. To restore a custom configuration, type the following commands:

```
/etc/powermt load file=/etc/powermt_custom.xml.version
/etc/powermt save
```

**Note:** Although you can restore an earlier `powermt_custom.xml` file, it is not currently possible to restore an earlier `emcp_devicesDB` file.

---

## R1/R2 boot failover support

Mirroring a bootable `emcpower` device using Symmetrix Remote Data Facility (SRDF) makes it possible, should a server failure at the local storage system fail, to fail over the boot disk to the remote mirror disk and then boot the server on an identical remote host.

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

## R1/R2 supported configurations

R1/R2 boot support is provided with both SPARC and x86 hosts and with both Leadville and non-Leadville drivers. Refer to the E-Lab Interoperability Navigator for more information on supported hosts and drivers.

EMC supports the following specific R1/R2 configuration:

- ◆ The R1 and R2 hosts must have identical hardware. The host should have the same number and types of HBAs.
- ◆ The R1 and R2 hosts must have the same number of devices, the devices must all be paired, and added to the disk group for SRDF failover and failback commands.
- ◆ The R1 and R2 hosts cannot have non-mirrored devices, BCVs, or gatekeepers.
- ◆ Each R1 device has only one mirror. Concurrent SRDF is not supported.
- ◆ All R1 devices reside on one Symmetrix (for example, *Symmetrix A*) and are visible only to a single host (for example, *Host A*).
- ◆ All R2 devices reside on a separate Symmetrix (for example, *Symmetrix B*) and are visible only to the identical host in reserve (for example, *Host B*).
- ◆ SRDF is managed from either of the following two facilities:
  - ControlCenter Management Server
  - Symmetrix Service Processor

## Enabling R1/R2 boot disk failover

For R1/R2 boot support, the hosts must be configured to boot from an emcpower device. Refer to [Chapter 3, “Configuring a PowerPath Boot Device on Solaris,”](#) for more information.

To enable R1/R2 boot disk failover follow these steps:

1. Execute `powercfg -z` from the R1 host.
2. Execute `bootadm update-archive` to update the archive files.

The `powercfg -Z` command ensures that the `emcp.conf` file is updated with an entry that contains R1 and R2 Symmetrix volume IDs for the pseudo (emcpower) device. When the host on R1 or R2 side boots, PowerPath intelligently determines the R1 and R2 site based on the information in the configuration file and acts accordingly.

## Booting the R2 host after failover

1. Boot off the external pseudo device on the R2 host:
  - On a SPARC host, type the `boot <hardware_path> -r` command at the **OK** prompt.
  - On an x86 host, power on the host, stop the boot process, and update the kernel boot path with the `-r` flag to perform a reconfiguration boot.
2. Run the following commands to remove any dead paths:

```
powermt check [force]
devfsadm -C
powermt config
powermt save
```

3. Run `powermt display` to ensure that no dead paths remain.

---

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

---

## Booting the R1 host after failback

1. Boot off the external pseudo device on the R1 host:
  - On a SPARC host, type the `boot <hardware_path> -r` command at the **OK** prompt.
  - On an x86 host, power on the host, stop the boot process, and update the kernel boot path with the `-r` flag to perform a reconfiguration boot.
2. Run the following commands to remove any dead paths:

```
powermt check [force]
devfsadm -C
powermt config
powermt save
```

3. Run `powermt display` to ensure that no dead paths remain.

## Device naming

PowerPath for Solaris presents PowerPath-enabled storage system logical devices to the operating system by all their native devices plus a single PowerPath-specific pseudo device. Applications and operating system services can use any of these devices—native or pseudo—to access a PowerPath-enabled storage system logical device.

**Note:** Run `powermt config` after changing the label from EFI to SMI or from SMI to EFI, to update the PowerPath device name/links.

VxVM 4.2 supports PowerPath pseudo device names with third-party arrays.

- ◆ VxVM 4.2 support PowerPath pseudo device names with third-party arrays.
- ◆ Oracle cluster version 4.2 does not support the use of an emcpower pseudo device as a quorum device on all currently supported versions of PowerPath for Solaris.

### Device-name format

SPARC platforms use a VTOC with 8 slices; x86 platforms use a VTOC with 16 slices. Consequently, device names for these architectures differ. x86 hosts also use fdisk partitioning, which allows you to partition the disk from one to four bootable fdisk partitions.

[Table 2 on page 54](#) summarizes the device names used for SPARC and x86 hosts.

**Note:** The maximum number of target pseudo device instances per x86 is 4095 and must be in the range of 0 - 4094.

**Table 2** Solaris device names

Architecture	Disk label	PowerPath pseudo-device names	Native-device names
SPARC	VTOC	emcpower#[a-h]	c#t#d#s[0-7]
	EFI	<ul style="list-style-type: none"> <li>• emcpower#[a-g]</li> <li>• emcpower# (represents the entire disk)</li> </ul>	<ul style="list-style-type: none"> <li>• c#t#d#[0-6]</li> <li>• c#t#d# (represents the entire disk)</li> </ul>
x86	VTOC	emcpower#[a-p]	<ul style="list-style-type: none"> <li>• c#t#d#s[0-15] (slices)</li> <li>• c#t#d#p[0-4] (fdisk partition)</li> </ul>

### Pseudo device

A *pseudo device* describes a device special file of the format emcpower#[a-h] or emcpower#[a-p]:

Where:

- ◆ # = disk number
- ◆ [a-h] = slice

---

**Note:** EMC does not provide pseudo-name support for the fdisk partitions (native device names `c#t#d#p#[p0-p4]`). To access blocks from the fdisk partition or to use the fdisk command, you must use native device names.

---

Slices in Sys V identifiers are designated `s0`, `s1`, `s2`, and so on. They correspond exactly to emcpower slices designated `a`, `b`, `c`, and so on. Therefore, if device `c0t0d0` corresponds to device **emcpower0**, slice `c0t0d0s2` corresponds to slice **emcpower0c**.

### Examples

- ◆ Block device — `/dev/dsk/emcpower3c`
- ◆ Raw device — `/dev/rdisk/emcpower3c`

## Native device

A native device describes a device special file of the format `c#t#d#s#` or `c#t#d#p#`.

Where:

- ◆ `c#` = instance number for the interface card.
- ◆ `t#` = target address of the storage system logical device, the bus.
- ◆ `d#` = storage system logical device at the target.
- ◆ `s#` = slice, ranging from 0 to 7 or 0 to 15, depending on the architecture.
- ◆ `p#` = fdisk partition, ranging from 0 to 4 (x86 only).

### Examples

- ◆ Block device (native) — `/dev/dsk/c3t4d0s0`
- ◆ Raw device (native) — `/dev/rdisk/c3t4d0s0`

## Selecting a device naming convention

After PowerPath is installed, a host has both native devices and emcpower devices enabled and available for use. Both native devices and emcpower devices can be active simultaneously on a host.

Pseudo (emcpower) devices offer the following advantages:

- ◆ Pseudo devices allow you to perform non-disruptive migrations with PowerPath Migration Enabler.
- ◆ When you upgrade or change HBAs, pseudo device names do not need to change.
- ◆ There is only one pseudo device name for each multipathed logical device; however, note that there are multiple native device names for each multipathed logical device, which is harder to manage.

Pseudo device names are easier to manage because there is a one-to-one relationship between pseudo device names and logical volumes. Native device names are based on HBA, target, and device assignments as recognized at system startup.

- ◆ Implementing PowerPath's boot-time, boot-path failover feature requires pseudo devices.

- ◆ Operations of Sun Microsystem’s Dynamic Reconfiguration (DR) feature are transparent to applications using pseudo devices. (With native devices, these operations are not transparent: c#t#d#s# paths are removed, which can be disruptive to applications using those paths.) Refer to [“Dynamic Reconfiguration” on page 58](#) for information on using DR to add and remove HBAs in a PowerPath environment.

Native devices offer the following advantages:

- ◆ VxVM 5.1 does not support PowerPath pseudo device names with third-party arrays. If both PowerPath and VxVM 5.1 or later are installed, VxVM automatically uses pseudo devices for EMC arrays.
- ◆ Existing applications, like volume managers and DBMSs, need not be modified to provide PowerPath multipathing and path failover functionality, because they can directly access PowerPath logical devices through native devices. (With pseudo devices, existing applications need to be modified to use this functionality.)

[Table 3 on page 56](#) summarizes the functional differences between native devices and emcpower devices in the Solaris environment.

**Table 3** Native devices compared with emcpower devices

Function	Native device	Pseudo device
I/O failover	✓	✓
I/O load balancing	✓	✓
Booting: boot-path failover	No	✓
Reboot (reconfiguration)	✓ (Partial support) If a path is missing, PowerPath does not create a “replacement” c#t#d# device.	✓ (Full support)
Support for VxVM sliced disks	✓	✓ Supported with VxVM 5.1 or later with EMC storage devices only.
Support for VxVM simple disks	✓	✓
Support for Solaris disk partitions (slices)	✓	✓
Support for interaction with VxVM DMP (Dynamic MultiPathing) (When using PowerPath 5.0 or later, you can manage a VNX OE and CLARiiON storage system with either PowerPath or DMP, but not with both.)	✓	✓ (VxVM 5.1 or later support pseudo names with EMC storage devices only)
DR transparency	Limitations — c#t#d#s# paths are removed, which can be disruptive to applications using those paths.	✓
IOCTL deterministic path selection	✓ (PowerPath selects the specific path.)	No (PowerPath selects an arbitrary path.)
Nondisruptive data migrations with PowerPath Migration Enabler	No Minimal disruption is necessary to reconfigure applications after a migration.	✓



Table 4 on page 57 indicates when native and pseudo devices are supported, and which device naming conventions we prefer in environments with specified software requirements. If no preference is specified, both supported options are equally good.

**Table 4 Support for native and pseudo devices**

Software features	Are native devices supported?	Are pseudo devices supported?
<b>Boot requirements</b>		
Boot-time failover—Symmetrix, VNX OE, CLARiiON, XtremIO and VPLEX	No	✓  <b>Note:</b> Failover is not supported during the initial phase of boot; after boot, failover is supported.
<b>Volume Managers</b>		
VxVM 5.1 <ul style="list-style-type: none"> <li>Pseudo device names are used by default for EMC arrays</li> <li>Pseudo device names are not supported for third-party arrays</li> </ul>	✓ (EMC and third-party devices)	✓ (EMC devices only) (Preferred)
<ul style="list-style-type: none"> <li>VxVM 6.0 + latest patches + latest ASL (Supports PowerPath on i386 only)</li> <li>VxVM 6.1 + latest patches + latest ASL (Supports PowerPath on SPARC only)</li> </ul>	✓ (EMC and third-party devices)	✓ (EMC devices only) (Preferred)
Other volume managers (including Sun's DiskSuite, raw devices, file systems on raw devices, and raw table spaces) <ul style="list-style-type: none"> <li>If DR is used</li> <li>If DR is not used</li> </ul>	✓ ✓	✓ (Preferred) ✓

## Reconfiguring PowerPath devices online

Whenever the physical configuration of the storage system or the host changes, you must reconfigure the PowerPath devices to avoid data loss. Configuration changes that require you to reconfigure PowerPath devices include:

- ◆ Removing or adding HBAs (see also [“Dynamic Reconfiguration” on page 58](#))
  - ◆ Adding, removing, or changing storage system logical devices
  - ◆ Changing the cabling routes between HBAs and storage system ports
  - ◆ Adding or removing storage system interfaces
1. To reconfigure PowerPath devices, run the **devfsadm -C** command to create the device nodes.
  2. Type **powermt config** to configure the paths to the storage system device.  
Run **powermt config** before sending any I/O to the new devices.
  3. Type **powermt save** to save the new PowerPath configuration.
  4. Run the **format** command and look for emcpower devices.
  5. Type **powermt display dev=all** to display the new device.

## Dynamic Reconfiguration

The Solaris Dynamic Reconfiguration (DR) feature allows you to add or remove an HBA from a Solaris system while the system continues running. You can logically attach and detach system boards from the operating system without halting and rebooting. For example, with DR you can detach a board from the operating system, physically remove and service the board, and then re-insert the board and re-attach it to the operating system—without halting the operating system or terminating any user application.

PowerPath supports DR. The following procedures describe how to use DR to remove or add HBAs in a PowerPath environment.

As you perform these procedures, ensure that the Sun Dynamic Reconfiguration documentation is available for your platform.

---

**Note:** If you have a custom PowerPath configuration that you have not yet saved, run **powermt save** before completing the procedures in this section, to save your configuration changes. Run **powermt load** after completing these procedures, to restore your configuration.

---

### Removing an HBA from a PowerPath configuration

#### **IMPORTANT**

Failure to follow the steps exactly as documented could result in data loss.

To use DR to remove an HBA from a Sun Fire server in a PowerPath configuration:

1. Ensure that all applications have closed the paths to be removed.
2. Type **cfgadm -vl** to list the status of the dynamically configurable resources.

Output similar to the following appears:

Ap_Id	Type	Busy	Receptacle Phys_Id	Occupant	Condition	Information
c0	scsi-bus	n	/devices/pci@1f,4000/scsi@3:scsi	connected configured	unknown	
c1	scsi-bus	n	/devices/pci@1f,4000/scsi@2:scsi	connected configured	unknown	
<b>c2</b>	fc-fabric	n	/devices/ <b>pci@4,4000/fibre-channel@2</b> /fp@0,0:fc	connected configured	unknown	
c3	fc-fabric	n	/devices/pci@4,4000/fibre-channel@2,1/fp@0,0:	connected configured	unknown	
c4	fc-fabric	n	/devices/pci@6,4000/lpfc@2/fp@0,0:fc	connected configured	unknown	
c5	fc-fabric	n	/devices/pci@6,2000/lpfc@1/fp@0,0:fc	connected configured	unknown	

Locate the information for the PCI adapter you want to remove. For this example, assume you want to remove **pci@4,4000/fibre-channel@2/fp@0,0**, which is shown in bold text in the output above.

3. Type **powermt display** to locate the PCI adapter information in the **powermt display** output.

Output similar to the following appears:

```

Symmetrix logical device count=18
CLARiiON logical device count=0
Hitachi logical device count=7
Invista logical device count=0
HP xp logical device count=10
Ess logical device count=10
HP HSx logical device count=3
=====
----- Host Bus Adapters ----- I/O Paths ----- Stats -----
### HW Path Summary Total Dead IO/Sec Q-IOs Errors
=====
3072 pci@4,4000/fibre-channel@2/fp@0,0 optimal 6 0 - 0 0
3073 pci@4,4000/fibre-channel@2,1/fp@0,0 optimal 6 0 - 0 0
3074 pci@6,4000/lpfc@2/fp@0,0 optimal 79 0 - 0 0
3075 pci@6,2000/lpfc@1/fp@0,0 optimal 79 0 - 0 0

```

Locate the adapter number of the HBA you want to remove. In this example, **pci@4,4000/fibre-channel@2/fp@0,0** corresponds to PowerPath adapter **3072**.

4. Type **vxdmpadm disable ctr=cX** in a cluster environment, to remove the HBA from the Veritas DMP environment. Where *X* is the adapter number shown in the **cfgadm -vl** output in step 2, excerpted here:

```

c2                connected    configured    unknown
unavailable fc-fabric    n
/devices/pci@4,4000/fibre-channel@2/fp@0,0:fc

```

In this example, the controller number is 2 (shown as **c2** in the output):

```
vxdmpadm disable ctr=c2
```

5. Type **powermt disable hba=#** to remove the HBA from the PowerPath configuration, where: # corresponds with the PowerPath adapter number identified in step 3. For example

```
powermt disable hba=3072
```

6. Type **cfgadm -vc unconfigure cX** to unconfigure the HBA, where *X* is the controller number typed in step 4. For example:

```
cfgadm -vc disconnect C2
```

7. Type **cfgadm -vc disconnect cX** to disconnect the HBA, where *X* is the controller number typed in step 4. For example:

```
cfgadm -vc disconnect C2
```

## Cleaning up incorrect path or logical volume removal

1. Run **powermt display** to identify the removed volume and its paths. In the output these should all display as dead.
2. Run **powermt check dev=<volume identifier>** for volume in question.

## Adding an HBA to a PowerPath configuration

To use DR to add an HBA to a Solaris system in a PowerPath configuration:

1. Add the new HBA to the system, following the instructions in the Sun Dynamic Reconfiguration documentation.

2. If you are replacing an HBA and had previously disabled the HBA, enable it. Type **powermt enable hba=*hba***.

Where *hba* is the HBA number, as shown in **powermt display**. For example, **powermt enable hba=3072**.

3. Type **powermt config** to configure the new HBA.
4. Type **powermt restore** to restore dead paths.
5. Type **powermt save** to save the new configuration.

## Managing and unmanaging Celerra iSCSI devices

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

MPxIO must be globally disabled before configuring PowerPath's generic LAM for Celerra support. To disable MPxIO globally, replace occurrences of `mpxio-disable=no` with `mpxio-disable=yes` in `/kernel/drv/fp.conf` and `/kernel/drv/iscsi.conf`.

### Manage Celerra iSCSI devices under PowerPath

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

```
managed:generic:EMC:Celerra
```
2. Run **/etc/powermt init**.
3. Run **/etc/powercf -u**.
4. Reboot the host.

### Unmanage Celerra iSCSI devices under PowerPath

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

```
managed:generic:EMC:Celerra
```
2. Run **/etc/powermt init**.
3. Run **/etc/powercf -u**.
4. Reboot the host.

## Supporting Veritas Volume Manager

Veritas Volume Manager (VxVM) 5.1 or later includes a new array support library (ASL) for EMC PowerPath pseudo devices. When you install VxVM 5.1 or later in a PowerPath environment, VxVM automatically does the following:

- ◆ Converts any Symmetrix, VNX OE, and CLARiiON devices with native names to pseudo-named devices (emcpower names).
- ◆ Puts the EMC pseudo-named devices under VxVM control.

EMC recommends using the default VxVM 5.1 or later naming convention (pseudo names) for EMC storage devices. Veritas Volume Manager 5.1 and later does not support PowerPath pseudo device names with third-party arrays. To configure native names (c#t#d# names), you can change the VxVM third-party driver mode (tpdmode) to native, as described in [“Using native device names for EMC devices.”](#) This change also applies to any EMC devices in a mixed array environment.

The *PowerPath Family for Solaris Release Notes*, available on EMC Online Support, provides information on supported VxVM versions on PowerPath for Solaris. The *PowerPath for Solaris Interoperability with VERITAS Volume Manager Technical Notes*, available on EMC Online Support, provides more information on PowerPath for Solaris with VxVM.

---

**Note:**

- The **powervxvm** command (which was required with earlier releases of VxVM) is not needed with VxVM 5.1 or later.
  - When installing PowerPath and VxVM 5.1 or later, install PowerPath first.
- 

## Using native device names for EMC devices

To change EMC storage devices from the default pseudo names to native names:

---

**Note:** Use the procedure in this section only if the default disk-naming scheme is set to operating system–based naming, and the TPD-controlled enclosure does not contain fabric disks.

---

1. View the pseudo device names that appear after installing VxVM 5.1 or later.

```
# vxdisk list

DEVICE      TYPE          DISK  GROUP  STATUS
emcpower10s2 auto:sliced  disk1  mydg  online
emcpower11s2 auto:sliced  disk2  mydg  online
emcpower12s2 auto:sliced  disk3  mydg  online
emcpower13s2 auto:sliced  disk4  mydg  online
emcpower14s2 auto:sliced  disk5  mydg  online
emcpower15s2 auto:sliced  disk6  mydg  online
emcpower16s2 auto:sliced  disk7  mydg  online
emcpower17s2 auto:sliced  disk8  mydg  online
emcpower18s2 auto:sliced  disk9  mydg  online
emcpower19s2 auto:sliced  disk10 mydg  online
```

2. Change the pseudo names to native names.

```
# vxddm padm setattr enclosure <enclosure> tpdmode=native
# vxddm padm listenclosure
```

where *<enclosure>* is the enclosure name of the EMC device.

3. Repeat [step 2 on page 61](#) for all the enclosures.
4. List the device names and note the different naming convention in the **Device** column.

```
# vxdisk list

DEVICE      TYPE          DISK  GROUP  STATUS
c6t0d10s2  auto:sliced  disk1  mydg  online
c6t0d11s2  auto:sliced  disk2  mydg  online
```

```

c6t0d12s2 auto:sliced disk3 mydg online
c6t0d13s2 auto:sliced disk4 mydg online
c6t0d14s2 auto:sliced disk5 mydg online
c6t0d15s2 auto:sliced disk6 mydg online
c6t0d16s2 auto:sliced disk7 mydg online
c6t0d17s2 auto:sliced disk8 mydg online
c6t0d18s2 auto:sliced disk9 mydg online
c6t0d19s2 auto:sliced disk10 mydg online

```

If **tpdmode** is set to **native**, the subpath with the smallest device number is displayed.

For information about troubleshooting a PowerPath with VxVM, environment, refer to [“Troubleshooting scenarios” on page 66](#).

## Exporting emcpower devices to Non-Global Zones

**Note:** This section applies to PowerPath Solaris 11.x.

This section describes how to export PowerPath pseudo devices (emcpower devices) to Non-Global Zones (also called NG Zones). After you export the pseudo devices, you can create file systems from the NG Zones, or you can use them through volume managers.

To export emcpower devices to NG Zones:

1. Manually type the pseudo device names in the `/etc/devlink.tab` file.

**Note:** For each pseudo device to be exported, create two entries: one for the raw device and one for the block device.

—You must type a single tab between the two fields in each line. For example, the tab appears between **blk** and **dsk** in the first sample entry. Adding spaces instead of tabs causes the entries to be ignored.

For example, type these lines in the `/etc/devlink.tab`:

```

type=ddi_pseudo;name=emcp;addr=0;minor=a,blk    dsk/emcpower\A0\M1
type=ddi_pseudo;name=emcp;addr=0;minor=a,raw    rdsk/emcpower\A0\M1
type=ddi_pseudo;name=emcp;addr=22;minor=c,blk    dsk/emcpower\A0\M1
type=ddi_pseudo;name=emcp;addr=22;minor=c,raw    rdsk/emcpower\A0\M1

```

2. After you modify the `/etc/devlink.tab` file, reboot the host.
3. After rebooting the host, ensure that the NG Zones are not in a booted condition and then make configuration changes by typing the following commands:

**Note:** In this example, the NG Zone is **zone1**.

```

# zonecfg -z zone1
zonecfg:zone1> add device
zonecfg:zone1:device> set match=/dev/rdsk/emcpower0a
zonecfg:zone1:device> end
zonecfg:zone1> add device
zonecfg:zone1:device> set match=/dev/dsk/emcpower0a
zonecfg:zone1:device> end
zonecfg:zone1> exit
zonecfg:zone1> add device
zonecfg:zone1:device> set match=/dev/rdsk/emcpower22c
zonecfg:zone1:device> end
zonecfg:zone1> add device

```

```
zonecfg:zone1:device> set match=/dev/dsk/emcpower22c
zonecfg:zone1:device> end
zonecfg:zone1> exit
# zoneadm -z zone1 boot
```

- For Solaris 11 and later releases, add the native devices corresponding to the pseudo device to the non-global zone.

For example,

```
# powermt display dev=0
Pseudo name=emcpower0a
Symmetrix ID=000195700642
Logical device ID=046A
state=alive; policy=SymmOpt; queued-I/Os=0
=====
----- Host ----- - Stor - -- I/O Path -- -- Stats ---
### HW Path          I/O Paths   Interf.  Mode    State  Q-I/Os Errors
=====
3074 pci@0,0/pci8086,3410@9/pci1077,15d@0,1/fp@0,0 c9t50000973000A0995d0s0 FA 6gB active
alive      0      0
3074 pci@0,0/pci8086,3410@9/pci1077,15d@0,1/fp@0,0 c9t50000973000A0991d0s0 FA 5gB active
alive      0      0
3072 pci@0,0/pci8086,3410@9/pci1077,15d@0/fp@0,0 c8t50000973000A0995d0s0 FA 6gB active
alive      0      0
3072 pci@0,0/pci8086,3410@9/pci1077,15d@0/fp@0,0 c8t50000973000A0991d0s0 FA 5gB active
alive      0      0
zonecfg:zone1> add device
zonecfg:zone1:device> set match=/dev/*dsk/c9t50000973000A0995d0*
zonecfg:zone1:device> end
zonecfg:zone1:device> end
zonecfg:zone1> exit

# zoneadm -z zone1 reboot
```

---

**Note:** Repeat these steps for each native device listed under **powermt display dev=0**

---

- Boot the NG Zone (**zone1** in this example) and then run the EMC **inq** command or the Solaris **format** command to verify that the emcpower device (**emcpower0a** in the example) is visible on the NG Zone you created.

## powercf configuration utility

During system boot on Solaris hosts, the **powercf** utility configures PowerPath devices by scanning HBAs for both single-ported and multiported storage system logical devices. (A multiported logical device shows up on two or more HBAs with the same storage system subsystem/device identity. The identity comes from the serial number for the logical device). For each storage system logical device found in the scan of the HBAs, **powercf** creates a corresponding emcpower device entry in the emcp.conf file and it saves a primary path and an alternate primary path to that device.

After PowerPath is installed, you need to run **powercf** only when the physical configuration of the storage system or the host changes. Configuration changes that require you to reconfigure PowerPath devices include the following:

- ◆ Adding or removing HBAs
- ◆ Adding, removing, or changing storage system logical devices
- ◆ Changing the cabling routes between HBAs and storage system ports

- ◆ Adding or removing storage system interfaces

[“Reconfiguring PowerPath devices online” on page 57](#) for instructions on reconfiguring PowerPath devices on Solaris and the *PowerPath Family CLI and System Messages Reference* for information on the command syntax.



# CHAPTER 8

## Troubleshooting

This chapter describes problems you might encounter and how to resolve them.

- ◆ [Audit and error messages](#) ..... 66
- ◆ [Troubleshooting scenarios](#)..... 66

## Audit and error messages

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

### Logging messages to a common file

EMC recommends configuring the `syslog.conf(4)` file to write **local0.info** messages to a common file. It is possible to specify more than one location for the messages.

**Note:** Refer to the **syslog.conf(4)** man page for specific information on configuring the **syslog.conf(4)** file.

To configure the `syslog.conf(4)` file to receive informational, error, warning, and audit (**emcpAudit**) messages:

1. Append the following line to the `/etc/syslog.conf` file unconditionally:

```
local0.info /var/adm/messages
```

#### **IMPORTANT**

You must place a tab after `local0.info`; a space will not work.

2. Restart the daemon for the configuration change to take effect. Refer to the **syslog.conf** man page for more information.

## Troubleshooting scenarios

This section describes problems you might encounter and suggests how to resolve them.

### Problem

You attempt to remove PowerPath in a Sun Cluster 3.1 with VxVM 4.1 environment and **pkgrm** fails.

#### Cause

Pseudo (**emcpower**) devices appear to be in use.

#### Solution

Type the following commands and retry **pkgrm**.

```
# cd /dev/rdisk
# for path in `ls emcpower*c`
do
/usr/sbin/vxdmpadm -f disable path=${path}
done
#
```

Uninstall PowerPath following the directions in [Chapter 6](#).

## Problem

**powermt display** shows an HBA as **UNKNOWN** after removal of a Fibre Channel cable and a host reboot.

### Cause

**powermt display** output does not reflect the correct HBA status.

### Solution

1. Run **powermt remove dev=all** to remove all devices.
2. Run **powermt config**.

## Problem

Pseudo device mappings in Veritas Volume Manager are inconsistent with PowerPath pseudo device mappings.

### Cause

This issue has been seen in the following circumstances:

- ◆ VxVM 4.1 is uninstalled and reinstalled after a LUN configuration change.  
VxVM maps pseudo devices based on device information saved before the uninstall and these mappings may differ from the current PowerPath mappings and configuration.
- ◆ A pseudo device instance is renamed using **emcpadm** which causes VxVM 4.1 to lose the pseudo device mapping.

The issue occurs because PowerPath pseudo device links are unavailable during boot, and Veritas VxVM 4.1 depends on these links before they can be created. This usually occurs when PowerPath is uninstalled and reinstalled without clearing the knowledge VxVM 4.1 has of these devices. It can also happen if a new LUN is mapped or zoned to a host and VxVM had prior knowledge of this device.

### Solution

1. Type **cp /etc/vx/disk.info /etc/vx/disk.info.old**
2. Type **echo > /etc/vx/disk.info**
3. Type **vxconfigd -k**

## Problem

With versions of VxVM earlier than VxVM 4.1 MP2 and VxVM 5.0 MP1 RP1, **vxdisk list** displays multiple entries for the same device.

Additionally, after unmanaging Symmetrix, VNX OE, or CLARiiON devices with the **powermt unmanage** command, VxVM continues to report that the removed devices belong to the Symmetrix, VNX OE, or CLARiiON array.

### Cause

This may occur following a reboot where PowerPath has just been installed or new devices have been mapped or zoned to a host.

The issue occurs because PowerPath pseudo device links are unavailable during boot, and Veritas VxVM 4.1 depends on these links before they can be created. This usually occurs when PowerPath is uninstalled and reinstalled without clearing the knowledge VxVM 4.1 has of these devices. It can also happen if a new LUN is mapped or zoned to a host and VxVM had prior knowledge of this device.

### Solution

Install VxVM 4.1 MP2 or VxVM 5.0 MP1 RP1 or later.

Alternatively, follow the appropriate action described in [Table 5 on page 68](#).

**Table 5 Duplicate device entries in VxVM**

If	Then
Upgrading or reinstalling PowerPath in a VxVM environment	Reboot after the initial package remove, if PowerPath requests the reboot.
Duplicate device entries are seen in <b>vxdisk list</b>	A second reboot will clear the duplicates. The PowerPath pseudo device links will be available during the second reboot.
Mapping/zoning new LUNS to the host	To avoid duplicate device entries: <ol style="list-style-type: none"> <li>1. Run <b>vxddladm stop eventsource</b>.</li> <li>2. Map/zone the new LUNS to the host.</li> <li>3. Dynamically recreate the LUN device (either using <b>cfgadm</b> routines or HBA-specific routines to dynamically create the device nodes).</li> <li>4. Run <b>powermt config</b> to discover and create the pseudo device links.</li> <li>5. Run <b>vxddladm start eventsource</b>.</li> </ol>
A host is rebooted with a path down	To avoid duplicate device entries: <ol style="list-style-type: none"> <li>1. Run <b>vxddladm stop eventsource</b>.</li> <li>2. Reconnect the paths.</li> <li>3. Run <b>powermt config</b>.</li> <li>4. Run <b>vxddladm start eventsource</b>.</li> </ol>
If <b>vxddladm</b> commands were not run	Reboot the system to remove the duplicate device entries.

### Problem

When installing PowerPath on a Solaris x86 host, the NFS client service goes offline, and NFS file systems are inaccessible.

### Cause

The package name is not included as part of the installation command (**pkgadd -d <package\_name>**).

### Solution

If you do not include the package name and encounter this problem:

1. Type **svcs | grep nfs** to verify the NFS client service is offline.
2. Type **svcadm disable nfs/client** to disable the NFS client service.
3. Kill the NFS mount/mountall commands started by the NFS start service.
4. Type **df -k** to verify the NFS mount is accessible.

5. Type `svcadm enable -r nfs/client` to enable the NFS client service.



# APPENDIX A

## Files Changed by PowerPath

This appendix lists files that are created or modified by PowerPath installation and upgrade.

- ◆ [Solaris files modified by PowerPath installation .....](#) 72
- ◆ [Files created by PowerPath installation .....](#) 73

## Solaris files modified by PowerPath installation

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

### SPARC hosts

- ◆ /etc/.login
- ◆ /etc/name\_to\_major
- ◆ /etc/profile
- ◆ /etc/rem\_name\_to\_major
- ◆ /etc/system
- ◆ /usr/share/man/windex

### x86 hosts

- ◆ /etc/.login
- ◆ /etc/name\_to\_major
- ◆ /etc/profile
- ◆ /etc/rem\_name\_to\_major
- ◆ /etc/system
- ◆ /platform/i86pc/boot\_archive

### /etc/system

---

**Note:** Before modifying /etc/system, the PowerPath installation saves a copy in /etc/system.pre-EMCpower.

---

### Forceload statements

On all hosts, PowerPath adds forceload statements for the PowerPath driver and miscellaneous kernel modules. For example, on a host running Solaris 5.7, PowerPath adds the following forceload statements to /etc/system:

```
forceload: drv/sd
forceload: drv/ssd
forceload: drv/emcp
forceload: misc/emcpx
forceload: misc/emcpsapi
forceload: misc/emcpcg
forceload: misc/emcpgpx
forceload: misc/emcpdm
forceload: misc/emcpioc
```

### Kernel stack size settings

On all hosts, PowerPath adds set statements for kernel **stksize** variables to increase default kernel stack sizes and avoid stack overflow panics. For example, on a host running Solaris 5.7, PowerPath adds the following stack size set statements to /etc/system:

```
set lwp_default_stksize=0x6000
set rpcmod:svc_run_stksize=0x6000
```

For more information, refer to [“Ensuring a sufficient stack size” on page 51](#).



## /etc/profile

On all hosts, PowerPath adds a line to /etc/profile that causes the PowerPath setup script to run in the current Bourne or Korn shell environment at system boot:

```
. /basedir/EMCpower/scripts/emcp_setup.sh
```

## /etc/.login

On all hosts, PowerPath adds a line to /etc/.login that causes the PowerPath setup script to run in the current C shell environment at system boot:

```
source /basedir/EMCpower/scripts/emcp_setup.csh
```

## Files created by PowerPath installation

**Note:** The directory *basedir* is the directory you specified as the PowerPath installation directory (**opt**, if you selected the default).

### Solaris 11.x x86 hosts

#### **/etc/emc/bin**

- ◆ emcp\_discover
- ◆ emcp\_ini
- ◆ emcp\_mond
- ◆ emcp\_mond\_edit
- ◆ emcp\_purge
- ◆ emcp\_setup.csh
- ◆ emcp\_setup.sh
- ◆ emcpcfg
- ◆ emcpdiscover
- ◆ emcpkvm
- ◆ emcpmgr
- ◆ emcpmigd
- ◆ emcppurge
- ◆ inquiry.pp
- ◆ installUtils.sh
- ◆ powerformat
- ◆ powermig

#### **/etc/emc**

- ◆ catalog/de\_DE.UTF-8/LC\_MESSAGES/EMCpower.mo
- ◆ catalog/es\_VE.UTF-8/LC\_MESSAGES/EMCpower.mo
- ◆ catalog/fr\_FR.UTF-8/LC\_MESSAGES/EMCpower.mo
- ◆ catalog/it\_IT.UTF-8/LC\_MESSAGES/EMCpower.mo
- ◆ catalog/ja\_JP.UTF-8/LC\_MESSAGES/EMCpower.mo
- ◆ catalog/ko\_KR.UTF-8/LC\_MESSAGES/EMCpower.mo
- ◆ catalog/pt\_BR.UTF-8/LC\_MESSAGES/EMCpower.mo
- ◆ catalog/zh\_CN.UTF-8/LC\_MESSAGES/EMCpower.mo
- ◆ emcpxarchive/

- ◆ mpa.excluded
- ◆ mpa.lams
- ◆ ppme/
- ◆ server.pem

#### **/etc/emc/install**

- ◆ installUtils.sh
- ◆ pkginfo
- ◆ pkgParms
- ◆ postinstall
- ◆ preremove

#### **/etc**

- ◆ rc2.d/S99powermig
- ◆ emcp\_devicesDB.dat
- ◆ emcp\_devicesDB.idx
- ◆ emcpcvt
- ◆ emcpmgr
- ◆ emcpreg
- ◆ emcpupgrade
- ◆ powercf
- ◆ powermt
- ◆ powermt\_custom.xml
- ◆ system.pre-EMCpower

#### **/lib/svc/method**

- ◆ emcp\_mond
- ◆ powerconfig
- ◆ powershift
- ◆ powershutdown
- ◆ powerstartup

#### **/kernel/drv**

- ◆ amd64/emcp
- ◆ amd64/emcpcf
- ◆ emcp
- ◆ emcp.conf
- ◆ emcpcf
- ◆ emcpcf.conf

#### **/kernel/misc**

- ◆ amd64/emcpdm
- ◆ amd64/emcpgpx
- ◆ amd64/emcpioc
- ◆ amd64/emcpx
- ◆ emcpdm
- ◆ emcpgpx
- ◆ emcpioc
- ◆ emcpx

#### **/opt/EMCpower/bin**

- ◆ amd64/emcpminor
- ◆ i386/emcpminor
- ◆ emcpadm

- ◆ emcpminor

#### **/opt/EMCpower/kernel/drv**

- ◆ amd64/emcp
- ◆ amd64/emcpsf
- ◆ emcp
- ◆ emcpsf

#### **/opt/EMCpower/kernel/misc**

- ◆ amd64/emcpdm
- ◆ amd64/emcpgpx
- ◆ amd64/emcpioc
- ◆ amd64/emcpx
- ◆ emcpdm
- ◆ emcpgpx
- ◆ emcpioc
- ◆ emcpx

#### **/opt/EMCpower/lib**

- ◆ amd64/libemcp.so
- ◆ amd64/libemcp\_64.so
- ◆ amd64/libemcp\_core.so
- ◆ amd64/libemcp\_core\_64.so
- ◆ amd64/libemcp\_lam.so
- ◆ amd64/libemcp\_lam\_64.so
- ◆ amd64/libemcp\_lic\_rtl.so
- ◆ amd64/libemcp\_lic\_rtl\_64.so
- ◆ amd64/libemcp\_mp\_rtl.so
- ◆ amd64/libemcp\_mp\_rtl\_64.so
- ◆ amd64/libemcpcg\_64.so
- ◆ amd64/libemcpmp.so
- ◆ amd64/libemcpmp\_64.so
- ◆ amd64/libmp\_64.so
- ◆ amd64/libpn.so
- ◆ amd64/libpn\_64.so
- ◆ libemcp.so
- ◆ libemcp\_32.so
- ◆ libemcp\_core.so
- ◆ libemcp\_core\_32.so
- ◆ libemcp\_lam.so
- ◆ libemcp\_lam\_32.so
- ◆ libemcp\_lic\_rtl.so
- ◆ libemcp\_lic\_rtl\_32.so
- ◆ libemcp\_mp\_rtl.so
- ◆ libemcp\_mp\_rtl\_32.so
- ◆ libemcpcg\_32.so
- ◆ libemcpmp.so
- ◆ libemcpmp\_32.so
- ◆ libmp\_32.so
- ◆ libpn.so
- ◆ libpn\_32.so

**/opt/EMCpower/tools**

- ◆ amd64/emcpserialfix
- ◆ amd64/emcpserialfix\_64
- ◆ emcpserialfix
- ◆ emcpserialfix\_32

**/opt/EMCpower/scripts**

- ◆ emcp\_cleanup
- ◆ emcp\_mond.conf
- ◆ emcp\_mond\_configure
- ◆ emcp\_setup.csh
- ◆ emcp\_setup.sh
- ◆ snmpd.local.conf
- ◆ sol\_liveup.sh
- ◆ update\_sma\_method

**/opt/EMCpower**

- ◆ smf/powerpath\_services.xml
- ◆ Postinstall.log

**/usr/lib**

- ◆ amd64/libcg.so
- ◆ amd64/libemcp.so
- ◆ amd64/libemcp\_lic\_rtl.so
- ◆ amd64/libemcp\_mp\_rtl.so
- ◆ amd64/libemcp\_mpapi\_rtl.so
- ◆ amd64/libemcpcg.so
- ◆ amd64/libemcpmp.so
- ◆ amd64/libemcppn.so
- ◆ amd64/libbpn.so
- ◆ libcg.so
- ◆ libemcp.so
- ◆ libemcp\_32.so
- ◆ libemcp\_lam.so
- ◆ libemcp\_lam\_32.so
- ◆ libemcp\_lic\_rtl.so
- ◆ libemcp\_lic\_rtl\_32.so
- ◆ libemcp\_mp\_rtl.so
- ◆ libemcp\_mp\_rtl\_32.so
- ◆ libemcp\_mpapi\_rtl.so
- ◆ libemcpcg.so
- ◆ libemcpmp.so
- ◆ libemcppn.so
- ◆ libmp\_32.so
- ◆ libbpn.so
- ◆ libpn\_32.so

**/usr/share/man**

- ◆ man1/emcpadm.1
- ◆ man1/emcpminor.1
- ◆ man1/emcpreg.1
- ◆ man1/emcupgrade.1
- ◆ man1/powercf.1

- ◆ man1/powerformat.1
- ◆ man1/powermig.1
- ◆ man1/powermt.1

## Solaris11.x SPARC hosts

### **/etc/emc/bin**

- ◆ emcp\_discover
- ◆ emcp\_ini
- ◆ emcp\_mond
- ◆ emcp\_mond\_edit
- ◆ emcp\_purge
- ◆ emcp\_setup.csh
- ◆ emcp\_setup.sh
- ◆ emcpcfg
- ◆ emcpdiscover
- ◆ emcpkvm
- ◆ emcpmgr
- ◆ emcpmigd
- ◆ emcppurge
- ◆ inquiry.pp
- ◆ installUtils.sh
- ◆ powerformat
- ◆ powermig

### **/etc/emc/lib**

- ◆ emcp\_mond\_netsnmp.so

### **/etc/emc/rsa/csf/xml**

- ◆ CSP-Admin.xml
- ◆ CSP-Authn.xml
- ◆ CSP-CatalogPosix.xml
- ◆ CSP-CatalogService.xml
- ◆ CSP-CatalogWindows.xml
- ◆ CSP-CertificateManagement.xml
- ◆ CSP-Config.xml
- ◆ CSP-DestinationAlert.xml
- ◆ CSP-DestinationEventLog.xml
- ◆ CSP-DestinationFile.xml
- ◆ CSP-DestinationODBC.xml
- ◆ CSP-DestinationSyslog.xml
- ◆ CSP-LDAP.xml
- ◆ CSP-LayoutFlat.xml
- ◆ CSP-LayoutRaw.xml
- ◆ CSP-LayoutXML.xml
- ◆ CSP-LocalDirectory.xml
- ◆ CSP-LocalizationService.xml
- ◆ CSP-Logging.xml
- ◆ CSP-OSLogin.xml
- ◆ CSP-PAM.xml
- ◆ CSP-PositionalFormatter.xml

- ◆ CSP-RoleManagement.xml
- ◆ CSP-SM.xml
- ◆ CSP-SecurID.xml
- ◆ CSP-SprintfFormatter.xml
- ◆ CSP-TrustedUser.xml
- ◆ CSP-UserManagement.xml
- ◆ CSP-saml-schema-assertion-1.1.xsd
- ◆ CSP-saml-schema-assertion-2.0.xsd
- ◆ CSP-xenc-schema.xsd
- ◆ CSP.xml
- ◆ Config.xml
- ◆ csp-app.xml
- ◆ xmldsig-core-schema.xsd

**/etc/emc/rsa/rkm\_client**

- ◆ config/rkm\_init.conf.tmpl
- ◆ config/rkm\_keyclass.conf.tmpl
- ◆ config/rkm\_registration.conf.tmpl
- ◆ config/rkm\_svc.conf.tmpl
- ◆ lib/libccme\_base.so
- ◆ lib/libccme\_ecc.so
- ◆ lib/libccme\_eccaccel.so
- ◆ lib/libccme\_eccnistaccel.so
- ◆ lib/libcryptocme2.sig
- ◆ lib/libcryptocme2.so
- ◆ lib/libkmclient\_shared.so
- ◆ lib/libkmcryptolib.so
- ◆ lib/libkmsvcshlib.so

**/etc/emc**

- ◆ catalog/de\_DE.UTF-8/LC\_MESSAGES/EMCpower.mo
- ◆ catalog/es\_VE.UTF-8/LC\_MESSAGES/EMCpower.mo
- ◆ catalog/fr\_FR.UTF-8/LC\_MESSAGES/EMCpower.mo
- ◆ catalog/it\_IT.UTF-8/LC\_MESSAGES/EMCpower.mo
- ◆ catalog/ja\_JP.UTF-8/LC\_MESSAGES/EMCpower.mo
- ◆ catalog/ko\_KR.UTF-8/LC\_MESSAGES/EMCpower.mo
- ◆ catalog/pt\_BR.UTF-8/LC\_MESSAGES/EMCpower.mo
- ◆ catalog/zh\_CN.UTF-8/LC\_MESSAGES/EMCpower.mo
- ◆ emcpxarchive/
- ◆ mpaa.excluded
- ◆ mpaa.lams
- ◆ ppme/
- ◆ server.pem

**/etc**

- ◆ rc2.d/S99powermig
- ◆ emcp\_devicesDB.dat
- ◆ emcp\_devicesDB.idx
- ◆ emcpcvt
- ◆ emcpmgr
- ◆ emcpreg
- ◆ emcpsf\_unload

- ◆ emcpupgrade
- ◆ powercf
- ◆ powermt
- ◆ powermt\_custom.xml
- ◆ system.pre-EMCpower

**/lib/svc/method**

- ◆ emcp\_mond
- ◆ powerconfig
- ◆ powershift
- ◆ powershutdown
- ◆ powerstartup

**/kernel/drv**

- ◆ sparcv9/emcp
- ◆ sparcv9/emcpsf
- ◆ emcp
- ◆ emcp.conf
- ◆ emcpsf
- ◆ emcpsf.conf

**/kernel/misc**

- ◆ sparcv9/emcpdm
- ◆ sparcv9/emcpgpx
- ◆ sparcv9/emcpioc
- ◆ sparcv9/emcpx
- ◆ emcpdm
- ◆ emcpgpx
- ◆ emcpioc
- ◆ emcpx

**/opt/EMCpower/bin**

- ◆ sparcv7/emcpminor
- ◆ sparcv9/emcpminor
- ◆ emcpadm
- ◆ emcpminor
- ◆ rkmc\_cache\_conv

**/opt/EMCpower/kernel/drv**

- ◆ sparcv9/emcp
- ◆ sparcv9/emcpsf
- ◆ emcp
- ◆ emcpsf

**/opt/EMCpower/kernel/misc**

- ◆ sparcv9/emcpdm
- ◆ sparcv9/emcpgpx
- ◆ sparcv9/emcpioc
- ◆ sparcv9/emcpx
- ◆ emcpdm
- ◆ emcpgpx
- ◆ emcpioc
- ◆ emcpx

**/opt/EMCpower/lib**

- ◆ sparcv9/libemcp.so
- ◆ sparcv9/libemcp\_64.so
- ◆ sparcv9/libemcp\_core.so
- ◆ sparcv9/libemcp\_core\_64.so
- ◆ sparcv9/libemcp\_lam.so
- ◆ sparcv9/libemcp\_lam\_64.so
- ◆ sparcv9/libemcp\_lic\_rtl.so
- ◆ sparcv9/libemcp\_lic\_rtl\_64.so
- ◆ sparcv9/libemcp\_mp\_rtl.so
- ◆ sparcv9/libemcp\_mp\_rtl\_64.so
- ◆ sparcv9/libemcpcg\_64.so
- ◆ sparcv9/libemcpmp.so
- ◆ sparcv9/libemcpmp\_64.so
- ◆ sparcv9/libmp\_64.so
- ◆ sparcv9/libpn.so
- ◆ sparcv9/libpn\_64.so
- ◆ libemcp.so
- ◆ libemcp\_32.so
- ◆ libemcp\_core.so
- ◆ libemcp\_core\_32.so
- ◆ libemcp\_lam.so
- ◆ libemcp\_lam\_32.so
- ◆ libemcp\_lic\_rtl.so
- ◆ libemcp\_lic\_rtl\_32.so
- ◆ libemcp\_mp\_rtl.so
- ◆ libemcp\_mp\_rtl\_32.so
- ◆ libemcpcg\_32.so
- ◆ libemcpmp.so
- ◆ libemcpmp\_32.so
- ◆ libmp\_32.so
- ◆ libpn.so
- ◆ libpn\_32.so

**/opt/EMCpower/tools**

- ◆ sparcv9/emcpserialfix
- ◆ sparcv9/emcpserialfix\_64
- ◆ emcpserialfix
- ◆ emcpserialfix\_32

**/opt/EMCpower/scripts**

- ◆ emcp\_cleanup
- ◆ emcp\_mond.conf
- ◆ emcp\_mond\_configure
- ◆ emcp\_setup.csh
- ◆ emcp\_setup.sh
- ◆ snmpd.local.conf
- ◆ sol\_liveup.sh

**/opt/EMCpower/install**

- ◆ .PPconfig
- ◆ checkinstall



- ◆ installUtils.sh
- ◆ pkgParms
- ◆ pkginfo
- ◆ pkgremove
- ◆ postinstall
- ◆ postremove
- ◆ preremove

#### **/opt/EMCpower**

- ◆ smf/powerpath\_services.xml
- ◆ Postinstall.log

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

- ◆ sparcv9/libcg.so
- ◆ sparcv9/libemcp.so
- ◆ sparcv9/libemcp\_lic\_rtl.so
- ◆ sparcv9/libemcp\_mp\_rtl.so
- ◆ sparcv9/libemcp\_mpapi\_rtl.so
- ◆ sparcv9/libemcpcg.so
- ◆ sparcv9/libemcpmp.so
- ◆ sparcv9/libemcppn.so
- ◆ sparcv9/libpn.so
- ◆ libcg.so
- ◆ libemcp.so
- ◆ libemcp\_32.so
- ◆ libemcp\_lam.so
- ◆ libemcp\_lam\_32.so
- ◆ libemcp\_lic\_rtl.so
- ◆ libemcp\_lic\_rtl\_32.so
- ◆ libemcp\_mp\_rtl.so
- ◆ libemcp\_mp\_rtl\_32.so
- ◆ libemcp\_mpapi\_rtl.so
- ◆ libemcpcg.so
- ◆ libemcpmp.so
- ◆ libemcppn.so
- ◆ libmp\_32.so
- ◆ libpn.so
- ◆ libpn\_32.so

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

- ◆ man1/emcpadm.1
- ◆ man1/emcpminor.1
- ◆ man1/emcpreg.1
- ◆ man1/emcupgrade.1
- ◆ man1/powercf.1
- ◆ man1/powerformat.1
- ◆ man1/powermig.1
- ◆ man1/powermt.1



# INDEX

## Symbols

/etc/.login 87  
/etc/profile 87  
/etc/system 86

## A

adding devices online 70

## B

boot device  
removing PowerPath control 37

## C

cat etc/name\_to\_major | grep vxio 27  
configuration  
R1/R2 65  
restoring custom 64

## D

device labeling 12  
devices  
adding devices online 70  
naming conventions 66  
documentation, related 7  
DR (Dynamic Reconfiguration) 71

## E

EFI labeling 12  
EMC Online Support website 7  
emcp\_mond.conf 60  
error messages  
logging 80  
ODS device overlap 63  
event monitoring daemon  
configuring 61  
definition 60

## F

forceload statements set by PowerPath 86  
format utility, Solaris 12

## H

HBA (Host Bus Adapter)  
adding 72  
removing 71  
Hitachi Dynamic Link Manager (HDLM), migrating from 50  
HP StorageWorks Secure Path, migrating from 48 to 49

## I

IBM Subsystem Device Driver (SDD), migrating from 49

installing PowerPath  
files created or modified 87  
installation procedure 13  
postinstallation checklist 44  
pre-installation instructions 10

## K

kernel stack size settings 86

## L

labeling, devices 12

## M

migrating from  
Hitachi Dynamic Link Manager (HDLM) 50  
HP StorageWorks Secure Path 48 to 49  
IBM Subsystem Device Driver (SDD) 49  
monitored events (SNMP) 60

## N

name\_to\_major file 27  
Navisphere agent 13

## P

powercf configuration utility 76  
PowerPath boot device  
configuring 32  
troubleshooting 37  
PowerPath Configuration Checker (PPCC) 11  
PowerPath devices  
reconfiguring online 70  
pre-installation instructions 10

## R

R1/R2 supported configurations 65  
reconfiguring PowerPath devices online 70  
removal  
post-removal 56 to 57

## S

SMI labeling 12  
SNMP event monitoring daemon 60  
SNMP trap packet format 62  
Solaris format utility 12  
Solaris Live Upgrade  
troubleshooting 45  
SRDF (Symmetrix Remote Data Facility) 65  
stack size 63  
Sun Cluster 3.0  
installing PowerPath with 26  
integrating PowerPath into 26

## Index

syslog.conf file 80

## T

trap packet format 62

troubleshooting

host hang after installing VxVM 81

Navisphere agent does not restart 13

pkgm fails 80

powermt display shows HBA as UNKNOWN 81

Solaris disk driver logging warning messages 63

Solaris Live Upgrade 45

vxdisk list displays multiple entries for one device 82

## U

unique host ID 27

## V

VCS (Veritas Cluster Server)

installing PowerPath with 27

integrating PowerPath into 28

vxio number 27