



PowerPath for Solaris
Version 4.3

**INSTALLATION AND
ADMINISTRATION GUIDE**

P/N 300-001-681
REV A01

EMC Corporation
Corporate Headquarters:
Hopkinton, MA 01748 -9103
1-508 -435 -1000
www.emc.com

Copyright © 1997-2004 EMC® Corporation. All rights reserved.

Printed July, 2004

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.

Trademark Information

EMC, PowerPath, LEGATO, EMC, EMC Symmetrix, Celerra, CLARION, CLARION, Documentr, VisualSAN, and where Navisphere, PowerPath, ResourcePak, SnapView/IP, SRDF, TimeFinde

Preface	vii
Chapter 1	Installing PowerPath
Before You Install	1-2
Obtain Up-To-Date Information.....	1-3
Choose a Convenient Time	1-3
Locate Your License Key	1-4
Prepare the Host and Storage System	1-4
Prepare for a Clustered Environment.....	1-7
Installing PowerPath 4.3	1-9
Installing PowerPath 4.3 Offline.....	1-14
Offline Installation over NFS	1-14
Upgrading to PowerPath 4.3.....	1-15
About Upgrading	1-15
Upgrade Procedure	1-16
Troubleshooting the Upgrade	1-21
After You Install	1-23
Register PowerPath on the Host.....	1-23
Remove the CD-ROM	1-24
Reboot the Host.....	1-25
Verify the PowerPath Installation	1-26
Verify that PowerPath Commands Are in Your Path	1-27
Verify PowerPath Devices Are Configured on the Host ..	1-28
Reconfigure Applications to Use emcpower Devices	1-29
Disable the Setup Script.....	1-29
Installing the PowerPath Volume Manager VCS Agent	1-30
Error Messages	1-32

Chapter 2	PowerPath in a Cluster Environment	
	PowerPath in a LEGATO 5.1 (or later) AAM Cluster.....	2-2
	Installing PowerPath in a New AAM Cluster.....	2-2
	Integrating/ Upgrading PowerPath into an Existing AAM	
	5.1 (or later) Cluster.....	2-3
	PowerPath in a Sun Cluster 3.x	2-4
	Installing PowerPath in a New Sun Cluster 3.x.....	2-4
	Integrating/ Upgrading PowerPath into an Existing Sun	
	Cluster 3.x.....	2-5
	PowerPath in a VERITAS Cluster Server Cluster	2-6
	Installing PowerPath in a New VCS Cluster.....	2-6
	Integrating/ Upgrading PowerPath into an Existing VCS	
	Cluster	2-8
	Configuring VCS to Recognize PowerPath Volume Manager	
	Resources	2-9
	Setting Major and Minor Numbers.....	2-11
Chapter 3	Configuring a PowerPath Boot Device on Solaris	
	Introduction.....	3-2
	Configuring a PowerPath Native Device as the Boot Device ...	3-3
	Partitioning the Boot Device	3-3
	Creating Filesystems	3-4
	Installing a Bootblk.....	3-5
	Copying Required Files	3-5
	Modifying OpenBoot	3-7
	Setting Up the Boot Alias	3-7
	Moving the Boot Device to an emcpower Device.....	3-8
	Recovery Procedure	3-9
	Removing PowerPath Control over a Boot Device.....	3-11
Chapter 4	Migrating to PowerPath	
	Migrating from HP StorageWorks Secure Path.....	4-2
	Secure Path Operation on Solaris.....	4-2
	Migration Procedure	4-2
	Migrating from IBM Subsystem Device Driver (SDD)	4-4
	SDD Operation on Solaris	4-4
	Migration Procedure	4-4
	Migrating from Hitachi Dynamic Link Manager (HDLM)	4-6
	HDLM Operation on Solaris.....	4-6
	Migration Procedure	4-6

Chapter 5	Removing PowerPath	
	Before You Remove PowerPath.....	5-2
	Removing PowerPath.....	5-3
	After You Remove PowerPath.....	5-4
	Removing PowerPath Files.....	5-4
	Rebooting the Host.....	5-5
	Removing PowerPath 4.3 Offline.....	5-6
	Removing the PowerPath Volume Manager VCS Agent.....	5-7
Chapter 6	PowerPath Administration on Solaris	
	PowerPath and CLARiiON Storage Systems.....	6-2
	Ensuring a Sufficient Stack Size.....	6-2
	Booting a Host with Built-In Fibre Channel Ports.....	6-3
	Rebooting and Custom Settings.....	6-4
	R1/R2 Boot Failover Support.....	6-4
	R1/R2 Supported Configurations.....	6-4
	R1/R2 Boot Procedure.....	6-5
	Device Naming.....	6-5
	Native Devices.....	6-6
	Pseudo Devices.....	6-6
	Selecting a Device Naming Convention.....	6-6
	Reconfiguring PowerPath Devices Online.....	6-10
	Dynamic Reconfiguration.....	6-11
	Adding an HBA to a PowerPath Configuration.....	6-11
	Removing an HBA from a PowerPath Configuration.....	6-12
	Upgrading Solaris.....	6-14
	Upgrading to Solaris 9 with Solaris Live Upgrade.....	6-14
	powercf Configuration Utility.....	6-16
	File Location.....	6-16
	Executing powercf.....	6-16
	emcp.conf File.....	6-16
	Syntax.....	6-17
	Arguments.....	6-17
	Error Messages.....	6-17
Appendix A	Files Changed By PowerPath	
	Solaris Files Modified by PowerPath Installation.....	A-2
	/etc/system.....	A-2
	/etc/profile.....	A-5
	/etc/.login.....	A-5
	/etc/syslog.conf.....	A-5

/etc/driver_classes	A-5
/etc/name_to_major	A-5
Files Created by PowerPath Installation	A-6
/etc	A-6
/usr/lib	A-7
/basedir/EMCpower/bin	A-9
/basedir/EMCpower/driver	A-9
/kernel/drv	A-12
kernel/misc	A-12
/basedir/EMCpower/lib	A-13
/basedir/EMCpower/man/man1	A-15
/basedir/EMCpower/scripts	A-16
/usr/bin	A-16
/usr/man/man1	A-16
Files Created by VCS Agent Installation	A-17
/basedir/VRTSvc/EMC	A-17
/basedir/VRTSvc/bin/EMCvg	A-17
/etc/VRTSvc/conf/config	A-17
/usr/lib	A-17

Appendix B Upgrading from Early Versions of PowerPath

Upgrading from PowerPath 2.1.x	B-2
Upgrading from PowerPath 2.0	B-2
Upgrading from PowerPath 1.5	B-3

Index	i-1
-------------	-----

As part of its effort to continuously improve and enhance the performance and capabilities of the EMC product line, EMC periodically releases new versions of the PowerPath software. Therefore, some functions described in this manual may not be supported by all versions of the PowerPath products or the storage-system hardware they support. For the most up-to-date information on product features, see the PowerPath release notes.

If a feature does not function properly or as described in this manual, please contact the EMC Customer Support Center for assistance. Refer to Where to Get Help on page xi for contact information.

This guide describes the features and functionality of EMC PowerPath for Solaris Version 4.3.

Audience and Prerequisites

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.

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

Organization Here is an overview of the information contained in this guide:

- ◆ Chapter 1, *Installing PowerPath*, describes how to install PowerPath 4.3 on a Solaris host; how to upgrade from PowerPath 4.2.x, 4.1.x, 4.0.x, 3.0.x, and some versions of 2.1.x; and how to install the PowerPath Volume Manager VCS Agent.
- ◆ Chapter 2, *PowerPath in a Cluster Environment*, describes how to install and configure PowerPath in Sun Cluster, LEGATO AAM, and VERITAS Cluster Server environments.
- ◆ Chapter 3, *Configuring a PowerPath Boot Device on Solaris*, describes how to configure a PowerPath device as the boot device.
- ◆ Chapter 4, *Migrating to PowerPath*, describes how to migrate to PowerPath from HP StorageWorks Secure Path, Hitachi Dynamic Link Manager (HDLM), and IBM Subsystem Device Driver (SDD).
- ◆ Chapter 5, *Removing PowerPath*, describes how to remove PowerPath, and the PowerPath Volume Manager VCS agent, from a Solaris host.
- ◆ Chapter 6, *PowerPath Administration on Solaris*, discusses Solaris issues and administrative tasks.
- ◆ Appendix A, *Files Changed By PowerPath*, lists files that are created or modified by PowerPath installation and upgrade.
- ◆ Appendix B, *Upgrading from Early Versions of PowerPath*, provides information on upgrading to PowerPath 4.3 from PowerPath 1.5, 2.0, and some versions of 2.1.x.

Related Documentation

Here is the complete set of PowerPath documentation; all manuals are available from EMC Corporation:

- ◆ *PowerPath Product Guide*
- ◆ *PowerPath Volume Manager User's Guide*
- ◆ *PowerPath Volume Mobility User's Guide*
- ◆ *EMC PowerPath Quick Reference*
- ◆ *PowerPath for AIX Installation and Administration Guide*
- ◆ *PowerPath for HP-UX Installation and Administration Guide*
- ◆ *PowerPath for Linux Installation Guide*
- ◆ *PowerPath for Solaris Installation and Administration Guide*
- ◆ *PowerPath for Windows Installation and Administration Guide*
- ◆ *EMC PowerPath for UNIX Release Notes*
- ◆ *EMC PowerPath for Linux Release Notes*
- ◆ *EMC PowerPath for Windows Release Notes*



Conventions Used in This Guide

EMC uses the following conventions for notes, cautions, warnings, and danger notices.

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



CAUTION

A caution contains information essential to avoid damage to the system or equipment. The caution may apply to hardware or software.

Typographical Conventions

EMC uses the following type style conventions in this guide:

AVANT GARDE	Keystrokes
-------------	------------

Palatino, bold	<ul style="list-style-type: none"> ◆ Dialog box, button, icon, and menu items in text ◆ Selections you can make from the user interface, including buttons, icons, options, and field names
-----------------------	---

<i>Palatino, italic</i>	<ul style="list-style-type: none"> ◆ New terms or unique word usage in text ◆ Command line arguments when used in text ◆ Book titles
-------------------------	---

<i>Courier, italic</i>	Arguments used in examples of command line syntax.
------------------------	--

Courier	System prompts and displays and specific filenames or complete paths. For example:
---------	--

```
working root directory [/user/emc]:
```

```
c:\Program Files\EMC\Symapi\db
```

Courier, bold	<ul style="list-style-type: none"> ◆ User entry. For example: sympoll -p
----------------------	--

- ◆ Options in command line syntax

Where to Get Help

For questions about technical support, call your local sales office or service provider.

If you have a valid EMC service contract, contact EMC Customer Service at:

United States: (800) 782-4362 (SVC-4EMC)

Canada: (800) 543-4782 (543-4SVC)

Worldwide: (508) 497-7901

Follow the voice menu prompts to open a service call and select the applicable product support.

Sales and Customer Service Contacts

For the list of EMC sales locations, please access the EMC home page at:

<http://www.emc.com/contact/>

For additional information on the EMC products and services available to customers and partners, refer to the EMC Powerlink website at:

<http://powerlink.emc.com>

EMC Support Matrix

For information about products and technologies qualified for use with the EMC software described in this manual, go to this EMC Web page:

<http://www.emc.com/horizontal/interoperability>

Choose the link to EMC Interoperability Support Matrices, and then the link to *EMC Support Matrix*.

Your Comments

Your suggestions will help us continue to improve the accuracy, organization, and overall quality of the user publications. Please send a message to techpub_comments@EMC.com with your opinions of this guide.

Installing PowerPath

This chapter describes how to install PowerPath 4.3 on a Solaris host and how to upgrade from an earlier version of PowerPath. The chapter also describes how to install the PowerPath Volume Manager VCS Agent for a VERITAS Cluster Server (VCS) environment.

The PowerPath Volume Manager LEGATO Agent is not part of the PowerPath 4.3 package. See the appropriate LEGATO documentation for information on installing and configuring that agent.

- ◆ Before You Install 1-2
- ◆ Installing PowerPath 4.3 1-9
- ◆ Installing PowerPath 4.3 Offline 1-14
- ◆ Upgrading to PowerPath 4.3 1-15
- ◆ After You Install 1-23
- ◆ Installing the PowerPath Volume Manager VCS Agent..... 1-30
- ◆ Error Messages 1-32

Before You Install

The sequence in which you configure the storage and install PowerPath® depends on the storage system you use:

Symmetrix, Hitachi Lightning, HP StorageWorks xp, and IBM ESS systems

Install PowerPath *after* you set up the storage system and verify that it is working properly.

CLARiiON storage systems

PowerPath installation is an integral part of a CLARiiON® setup and configuration procedure. To properly install PowerPath on a host with a CLARiiON array, refer to the *EMC Installation Roadmap for CX-Series and FC-Series Storage Systems*. There you will find step-by-step instructions for installing PowerPath and other CLARiiON software.

To obtain the latest versions of CLARiiON documentation, log in to the Powerlink™ website (<http://powerlink.emc.com>). Then click **Support > Document Library**.

First-time PowerPath installations —Remove any version of Navisphere® Application Transparent Failover (ATF) installed on the host.

The procedure for migrating from ATF or CDE to PowerPath is not straightforward and could result in data loss if not performed correctly. We strongly recommend that EMC Professional Services perform the migration. If you nevertheless decide to perform the migration yourself, refer to *Removing ATF or CDE Software Before Installing Other Failover Software* for more information. This document is available only on the EMC Powerlink website.

The rest of this section describes what to do before you install PowerPath on the host.

Obtain Up-To-Date Information

Check the Powerlink website (<http://powerlink.emc.com>) for the most current information:

- ❑ Release notes — We update the PowerPath release notes periodically and post them on the Web.
- ❑ EMC Bug Viewer — EMC Bug Viewer is an application that allows you to search for known problems and defects in EMC software. You can use the Bug Viewer to find:
 - Descriptions of PowerPath bugs existing on any PowerPath-supported host platform
 - Workarounds for existing bugs

We update this database regularly between scheduled releases and patch releases.

- ❑ Patches and notices — Review the patch ReadMe files to determine which patches (if any) you want to install after PowerPath, and whether those patches have any added prerequisites that must be met before you install PowerPath.

Choose a Convenient Time

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

Notes:

If you are upgrading from PowerPath 4.0.2 or later, you may not need to reboot after installing PowerPath 4.3. Refer to *Rebooting after the Upgrade* on page 1-15 for more information.

PowerPath 4.3 supports the `pkgadd -R` command, allowing you to install the package offline. Refer to *Installing PowerPath 4.3 Offline* on page 1-14.

Locate Your License Key

The PowerPath license registration key is on the License Key Card delivered with the PowerPath media kit.

If you are upgrading from an earlier version of PowerPath, you need not reregister; PowerPath will use your old key.

Prepare the Host and Storage System

To prepare the host and storage system:

- ❑ Verify that your environment meets the requirements in:
 - ❑ *Environment and System Requirements* section of the *EMC PowerPath Release Notes*. That section describes minimum hardware and software requirements for the host and supported storage systems.
 - ❑ **Chapter 3, *PowerPath Configuration Requirements*, in the *PowerPath Product Guide***. That chapter describes the host-storage system interconnection topologies that PowerPath supports.
- ❑ Set up the Fibre Channel port and LUN addresses. Refer to the EMC host connectivity guides (Symmetrix® and CLARiiON arrays), *EMC Installation Roadmap for CX-Series and FC-Series Storage Systems* (CLARiiON arrays), or the appropriate documentation from your vendor (third-party arrays).

In the `/etc/system` file, make sure the timeout value is set to 60 seconds. This minimizes path failover time without compromising online storage-system microcode or base code upgrades. The entry must be a hexadecimal number:

```
set sd: sd_io_time = 0x3C
```

- ❑ 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.

- ❑ Configure HBA drivers.



CAUTION

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

- ❑ For hosts connected to storage arrays via a Fibre Channel switch (that is, a fabric), configure the HBAs using persistent binding for SCSI target IDs.



CAUTION

Failure to do so could result in loss or corruption of data.

For information about persistent binding, refer to the *Host Connectivity Guide for Sun Solaris*, the HBA driver documentation, and the `.conf` file for your HBA type.

- ❑ Optionally, configure the host to send warnings to the console.

By default, warning messages are sent only to the log file. To see warnings on the console:

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

```
localX.Warn /dev/sysmsg
```

where `X` is a value that is not used by any other `local` setting in `/etc/syslog.conf`.

For example:

```
local0.Warn /dev/sysmsg
```

The separator between `localX.Warn` and `/dev/sysmsg` must be a tab character.

2. Set the `RAD_LOG_FACILITY` environment variable to `LOG_LOCALX`, where `X` is the number you used in the `local` setting of the `/etc/syslog.conf` file.

For example, for `csch`, enter:

```
setenv RAD_LOG_FACILITY LOG_LOCAL0
```

For `sh`, enter:

```
RAD_LOG_FACILITY=LOG_LOCAL0; export RAD_LOG_FACILITY
```

- ❑ If you plan to install both PowerPath and VxVM, install VxVM first.

If VxVM is already installed, install PowerPath. Once you have installed PowerPath, make sure the entries for the `vx*` drivers are above the entries for the PowerPath drivers in the `/etc/system` file.

- ❑ *If you are upgrading from PowerPath 4.0.2 or later:*

To avoid rebooting the host after upgrading from PowerPath 4.0.2 (or later):

- ❑ If any EMC ControlCenter™ agents are running on the host, stop the agents. Follow the instructions in the EMC ControlCenter documentation.
- ❑ If the Navisphere agent is installed on the host, stop the agent. Follow the instructions in the Navisphere documentation.

Note that the PowerPath 4.3 installation stops the Navisphere agent before installing PowerPath 4.3 and attempts to restart the agent after installation of PowerPath 4.3 completes.

However, in some situations, such as SP reboot during upgrade, PowerPath may have problems restarting the agent. Therefore, we recommend that you stop the Navisphere agent before the upgrade and restart the agent after the upgrade. In the alternative, once the upgrade is complete, make sure the agent is running, and restart it if necessary.

- ❑ Stop all applications that use PowerPath; for example, unmount all file systems mounted on PPVM volumes or PowerPath pseudo devices.

Otherwise, uninstalling PowerPath 4.0.2 or later in preparation for the upgrade will abort with a message indicating that PowerPath volumes or PowerPath pseudo devices are still in use. If this happens, you must either stop applications using PowerPath and try `pkgrm` again, or reboot the host and postpone the upgrade.

- ❑ Uninstall any earlier version of PowerPath.

The PowerPath installation program exits if it detects an existing version of PowerPath.

Important: Before you uninstall the earlier version, save your configuration using the `powermt save` command. Refer to the documentation for the earlier version for further uninstallation instructions.

You need not reboot the host after you uninstall the earlier version of PowerPath.

Prepare for a Clustered Environment

If you are installing PowerPath in a clustered environment:

- ❑ Prepare the cluster environment.

Refer to Chapter 2, *PowerPath in a Cluster Environment*, for information on installing PowerPath in a new or existing cluster.

- ❑ Find a major number that can be assigned to the PowerPath driver, `emcp`, on every host in the cluster. This must be a unique number that is unused by any host in the cluster. The installation will prompt you for this number.

Setting the major number is necessary only if you intend to use NFS mounted file systems on PowerPath Volume Manager volumes in your cluster. For more information, refer to *Setting Major and Minor Numbers* on page 2-11.

To find a unique major number:

1. Enter the following command on every host in the cluster:

```
sort -n +1 /etc/name_to_major
```

2. Note the highest major number on every node.

Then, when the installation prompts for a major number, use a number that is higher than all of these.

Installing PowerPath 4.3

This section describes how to install PowerPath 4.3 when no version of PowerPath has ever been installed on the host.

If you are upgrading to PowerPath 4.3 from an earlier version of PowerPath (or if you are reinstalling PowerPath 4.3) refer instead to *Upgrading to PowerPath 4.3* on page 1-15.

Notes: PowerPath 4.3 supports the `pkgadd -R` command. Refer to *Installing PowerPath 4.3 Offline* on page 1-14.

You can install PowerPath in single-user mode.

Mount the CD-ROM

1. Log in as root.
2. Insert the CD-ROM into the CD-ROM drive.

If the CD mounts automatically, continue with *Install the Software*, which follows.

If the CD does not mount automatically, you must mount it manually. Continue with step 3.

3. Mount the CD on your file system. For example, to mount the CD on `/cdrom/cdrom0`, enter:

```
mount -F hsfs -r /dev/dsk/cxytzyz0 /cdrom/cdrom0
```

where *x*, *y*, and *z* are values specific to the host's CD-ROM drive. For example:

```
mount -F hsfs -r /dev/dsk/c0t2d0s0 /cdrom/cdrom0
```

Install the Software

1. If you do not have a graphics terminal, run the `script filename` command to record `pkgadd` output in the specified file. (After `pkgadd` completes, use CTRL-D to stop recording the output.)
2. Change to the `/mount_point/UNIX/SOLARIS` directory. For example, enter:

```
cd /cdrom/cdrom0/UNIX/SOLARIS
```

3. Start the installation program. Enter:

```
/usr/sbin/pkgadd -d .
```

4. You see the following prompt:

The following packages are available:

```
1 EMCpower  EMC PowerPath
              (all) 4.3.0_bxxx
2 EMCvg     EMC Enterprise Cluster PowerPath Volume
              Manager VCS Agent
              (sparc) 4.3.0
```

Select package(s) you wish to process (or 'all' to process all packages). (default: all) [?,??,q]:

Enter **1** and press ENTER.

5. You are prompted for the directory where the PowerPath program files will be installed:

Enter package base directory (default: /opt):

Press ENTER to accept the default base directory (/opt), or type the path to an alternate base directory and press ENTER.

PowerPath installs its files in `/basedir/EMCpower`; the installation process creates the `EMCpower` subdirectory. Make a note of the name and location of the PowerPath base directory for future reference.

6. You are prompted for the directory where the log files will reside:

Enter directory for logging (default: `/basedir/EMCpower/log`):

Press ENTER to accept the default (`/basedir/EMCpower/log`), or type the path to an alternate log files directory and press ENTER. You must specify a local file system.

7. You see the following prompt:

Install for cluster environment [y,n,q,?] (default: n):

- If you are installing PowerPath on a standalone system, press ENTER. Go to step 10.
- If you are installing PowerPath in a clustered environment, enter **y** and press ENTER. Continue with step 8.

8. You are prompted for a major number that can be assigned to the PowerPath device driver (`emcp`) on every host in the cluster:

```
Assign major number for the device driver [?]:
```

Enter a major number and press **ENTER**. You must specify a major number that is not used by any host in the cluster.

The installation will assign this major number to the PowerPath driver. Specify this major number during PowerPath installation on every other host in the cluster.

9. You are prompted to confirm the major number:

```
Major major will be assigned to the device driver [y,n,q,?] (default: y):
```

Press **ENTER** to confirm the major number (or enter **n** to specify a different major number).

If you specify a major number that exceeds, by more than 25, the largest major number currently in use, the installation terminates with an error message. If this happens, refer to *Troubleshooting the Installation* on page 1-13.

10. You see the following prompt:

```
This package contains scripts which will be executed  
with super-user permission during the process of  
installing this package.
```

```
Do you want to continue with the installation of  
<EMCpower> [y,n,?]
```

Enter **y** and press **ENTER**.

11. The screen displays information about the installation, ending with:

```
-----  
* EMCpower installation:  
*  
* Installation is successful.  
*  
* If you have a license key card, register now. Then  
* REBOOT the host to complete the remaining steps of the  
* installation. Refer to the PowerPath for UNIX  
* Installation and Administration Guide for more  
* information.  
-----
```

Installation of <EMCpower> was successful.

The following packages are available:

```
1  EMCpower  EMC PowerPath  
      (all) 4.3.0_bxxx  
2  EMCvg     EMC Enterprise Cluster PowerPath Volume  
      Manager VCS Agent  
      (sparc) 4.3.0
```

Select package(s) you wish to process (or 'all' to process all packages). (default: all) [?,??,q]:

- If you are installing PowerPath for a VCS cluster and you want to install the PowerPath Volume Manager VCS Agent now, enter **2** and press ENTER. Continue with *Install the VCS Agent Software*, step 6, on page 1-31.
- Otherwise, enter **q** and press ENTER.

PowerPath is now installed on the host. You must register PowerPath, reboot the host, and perform other administrative tasks before PowerPath can run on the host. Refer to *After You Install* on page 1-23 for postinstallation information and procedures.

Troubleshooting the Installation

Solaris maintains a pool of extra major number slots in the kernel. A number that exceeds, by more than 25, the largest major number in use could fall outside this reserve pool and cause the system to panic when rebooted after installation.

Thus, if you specify a major number that is greater than the largest major number currently in use plus 25, the installation terminates with the following error message:

```
Error: The assigned major number is out of the allowable range. For this specific host configuration, the assigned major number should not exceed number.
```

If necessary, you can increase the maximum allowable major number:

1. Using a text editor, add a dummy line to `/etc/name_to_major`, specifying a major number 25 greater than the largest major number in use.
2. Reboot the system.
3. Install PowerPath.
4. Delete the dummy line from `/etc/name_to_major`.

Installing PowerPath 4.3 Offline

PowerPath 4.3 supports the `pkgadd -R` command. Use `pkgadd -R` to install PowerPath:

- ◆ On a second root disk on a single host
- ◆ Via NFS
- ◆ When booted from a CD-ROM

Any base directory and/or logging directory path you enter in response to an installation prompt should be specified relative to the client root, not the root of the host.

Offline Installation over NFS

Offline installation over NFS requires that the root file system of the target system be shared in a way that allows root access over NFS to the installing system. You can issue a command such as the following on *target_sys* to share the root file system on *target_sys* so that *installer_sys* has root access:

```
share -F nfs -d 'root on target_sys' -o ro,rw=installer_sys,root=installer_sys /
```

If the base directory of the package (the default is `/opt`) is not part of the root file system, it will also need to be shared with root access.

Upgrading to PowerPath 4.3

This section describes how to upgrade to PowerPath 4.3 from an earlier version of PowerPath.

About Upgrading

You can upgrade directly to PowerPath 4.3 from the following releases:

- ◆ PowerPath 4.2.x
- ◆ PowerPath 4.1.x
- ◆ PowerPath 4.0.x
- ◆ PowerPath 3.0.x
- ◆ PowerPath 2.1.0 and 2.1.2

If another version of PowerPath is installed on the host (for example, PowerPath 1.5, PowerPath 2.0, PowerPath 2.1.1), refer to Appendix B, *Upgrading from Early Versions of PowerPath*.

Uninstalling the Earlier Version

Before upgrading to PowerPath 4.3, you must uninstall the earlier version of PowerPath. Follow the instructions in the documentation for that version. *Note, however, that you need not reboot the host after you uninstall the earlier version.*

Exception: In a cluster, you must follow the cluster upgrade procedure, which may include booting the host out of the cluster before you upgrade.

Upgrading from PowerPath 2.1.0 or 2.1.2

Before you uninstall PowerPath 2.1.0 or 2.1.2, search the file `/etc/profile` for the following string: `end of additions`. If you find that string, replace it with the following string: `done additions`.

Rebooting after the Upgrade

PowerPath 4.0.2 and later. If you are upgrading from PowerPath 4.0.2 or later, you need not reboot the host after you install PowerPath 4.3. Note, however, that you must shut down all applications before you install PowerPath 4.3. Therefore, if the host boots off a PowerPath pseudo device, you must reboot after the upgrade.

PowerPath 4.0.1 and earlier. If you are upgrading from PowerPath 4.0.1 or earlier, you must reboot the host after you install PowerPath 4.3.

Licenses

The upgrade preserves your existing PowerPath license. You need not reenter license information.

Configuration Files

With the release of PowerPath 4.0, PowerPath configuration files changed in both content and number. The upgrade converts your PowerPath 3.0.x, 2.1.0, or 2.1.2 configuration files to the newer formats.

The old configuration files are saved in `/etc/emc/emcpxarchive`.

Upgrading preserves all your customized settings *except* write throttle queue settings. You must reinstate write throttle queue settings after the upgrade.

Upgrade Procedure

Mount the CD-ROM

1. Log in as root.

2. Insert the CD-ROM into the CD-ROM drive.

If the CD mounts automatically, continue with *Install the Software*, which follows.

If the CD does not mount automatically, you must mount it manually. Continue with step 3.

3. Mount the CD on your file system. For example, to mount the CD on `/cdrom/cdrom0`, enter:

```
mount -F hsfs -r /dev/dsk/cxydzs0 /cdrom/cdrom0
```

where *x*, *y*, and *z* are values specific to the host's CD-ROM drive. For example:

```
mount -F hsfs -r /dev/dsk/c0t2d0s0 /cdrom/cdrom0
```

Install the Software

1. If you do not have a graphics terminal, run the `script filename` command to record `pkgadd` output in the specified file. (After `pkgadd` completes, use CTRL-D to stop recording the output.)

2. Change to the `/mount_point/UNIX/SOLARIS` directory. For example, enter:

```
cd /cdrom/cdrom0/UNIX/SOLARIS
```

3. Start the installation program. Enter:

```
/usr/sbin/pkgadd -d .
```

4. You see the following prompt:

The following packages are available:

```
1 EMCpower  EMC PowerPath
              (all) 4.3.0_bxxx
2 EMCvg     EMC Enterprise Cluster PowerPath Volume
              Manager VCS Agent
              (sparc) 4.3.0
```

Select package(s) you wish to process (or 'all' to process all packages). (default: all) [?,??,q]:

Enter **1** and press ENTER.

5. You are prompted for the directory where the PowerPath program files will be installed:

Enter package base directory (default: /opt,?):

Press ENTER to accept the default base directory (/opt), or type the path to an alternate base directory and press ENTER.

PowerPath installs its files in `/basedir/EMCpower`; the installation process creates the `EMCpower` subdirectory. Make a note of the name and location of the PowerPath base directory for future reference.

6. You are prompted for the directory where the log files will reside:

Enter directory for logging (default: `/basedir/EMCpower/log`):

Press ENTER to accept the default (`/basedir/EMCpower/log`), or type the path to an alternate log files directory and press ENTER. You must specify a local file system.

7. You see the following prompt:

```
Install for cluster environment [y,n,q,?] (default: n):
```

- If you are installing PowerPath on a standalone system, press **ENTER**. Go to step 10.
- If you are installing PowerPath in a clustered environment, enter **y** and press **ENTER**. Continue with step 8.

8. You are prompted for a major number that can be assigned to the PowerPath device driver (emcp) on every host in the cluster:

```
Assign major number for the device driver [?]:
```

Enter a major number and press **ENTER**. You must specify a major number that is not used by any host in the cluster.

The installation will assign this major number to the PowerPath driver. Specify this major number during PowerPath installation on every other host in the cluster.

9. You are prompted to confirm the major number:

```
Major major will be assigned to the device driver [y,n,q,?] (default: y):
```

Press **ENTER** to confirm the major number (or enter **n** to specify a different major number).

If you specify a major number that exceeds, by more than 25, the largest major number currently in use, the installation terminates with an error message. If this happens, refer to *If the major number is out range* on page 1-22.

10. You see the following prompt:

```
Configuration for PowerPath version exists - want to upgrade [y,n,q,?](default:y):
```

where *version* is 4.3.x, 4.2.x, 4.1.x, 4.0.x, 3.0.x, or 2.x.

Press **ENTER** to upgrade.

11. You see the following prompt:

This package contains scripts which will be executed with super-user permission during the process of installing this package.

Do you want to continue with the installation of <EMCpower> [y,n,?]

Enter **y** and press ENTER.

12. The screen displays information about the installation.

- In a successful upgrade from PowerPath 4.0.2 or later on a host that does not boot off a PowerPath pseudo device, the display ends with:

```
-----  
* EMCpower installation:  
*  
* Installation is successful.  
*  
-----
```

The following packages are available:

```
1 EMCpower EMC PowerPath  
   (all) 4.3.0_bxxx  
2 EMCvg    EMC Enterprise Cluster PowerPath Volume  
   Manager VCS Agent  
   (sparc) 4.3.0
```

Select package(s) you wish to process (or 'all' to process all packages). (default: all) [?,??,q]:

- In a successful upgrade from PowerPath 4.0.1 or earlier, or if the host boots off a PowerPath pseudo device, the display includes a prompt to reboot the host:

```
-----
* EMCpower installation:
*
* Installation is successful.
*
* Please REBOOT the host to complete the remaining steps of the installation.
* Refer to the PowerPath for UNIX Installation and Administration Guide for more
* information.
-----
```

The following packages are available:

```
1 EMCpower  EMC PowerPath
           (all) 4.3.0_bxxx
2 EMCvgs    EMC Enterprise Cluster PowerPath Volume
           Manager VCS Agent
           (sparc) 4.3.0
```

Select package(s) you wish to process (or 'all' to process all packages). (default: all) [?,??,q]:

Important: If the display indicates that conversion of 3.0.x or 2.x configuration files failed, refer to *Troubleshooting the Upgrade* on page 1-21.

- If you are installing PowerPath for a VCS cluster and you want to install the PowerPath Volume Manager VCS Agent now, enter **2** and press **ENTER**. Continue with *Install the VCS Agent Software*, step 6, on page 1-31.
- Otherwise, enter **q** and press **ENTER**.

PowerPath 4.3 is now installed on the host. You must perform certain administrative tasks before PowerPath can run on the host. Refer to *After You Install* on page 1-23 for postinstallation information and procedures.

Troubleshooting the Upgrade

Upgrading from PowerPath 3.0.x or 2.x could fail for the following reasons:

- ◆ The PowerPath 3.0.x or 2.1 configuration files are corrupted.
- ◆ You inadvertently chose not to upgrade; that is, you answered **n** instead of **y** to the following prompt:

```
Configuration for PowerPath version exists - want to upgrade[y,n,q,?](default: y):
```

- ◆ You entered a major number that is out of the allowable range.

The following sections describe what to do under these circumstances.

If files are corrupted

If the PowerPath 3.0.x or 2.1 configuration files are corrupted, your custom configuration is no longer available after you install PowerPath 4.3. If any of your applications are configured with PowerPath pseudo devices, you must reconfigure those applications. Refer to *Installing and Configuring emcpower Devices With Solaris Applications*, which is available on the Powerlink website (<http://powerlink.emc.com>).

If you chose not to upgrade

If you inadvertently answer no to the upgrade prompt, before you reboot the host:

1. Run the `emcpupgrade check` command to check whether upgrade is possible. Enter:

```
emcpupgrade check -f /etc/emc/emcpxarchive/powermt.custom.saved
```

If the check fails, the upgrade cannot take place. Contact EMC Customer Support.

If the check succeeds, continue with step 2.

2. Run the `emcpupgrade convert` command to convert the PowerPath 3.0.x or 2.1 configuration files to the PowerPath 4.3 format. Enter:

```
emcpupgrade convert -f /etc/emc/emcpxarchive/powermt.custom.saved
```

For information about the `emcpupgrade` utility, refer to Chapter 5 of the *PowerPath Product Guide*.

If the major number is out range

Solaris maintains a pool of extra major number slots in the kernel. A number that exceeds, by more than 25, the largest major number in use could fall outside this reserve pool and cause the system to panic when rebooted after installation.

Thus, if you specify a major number that is greater than the largest major number currently in use plus 25, the installation terminates with the following error message:

```
Error: The assigned major number is out of the allowable range. For this specific host configuration, the assigned major number should not exceed number.
```

If necessary, you can increase the maximum allowable major number:

1. Using a text editor, add a dummy line to `/etc/name_to_major`, specifying a major number 25 greater than the largest major number in use.
2. Reboot the system.
3. Install PowerPath.
4. Delete the dummy line from `/etc/name_to_major`.

After You Install

After installing the PowerPath software:

- Register PowerPath on the host (first-time installation only).
- Remove the CD-ROM.
- Reboot the host, if necessary.

Important. If you are upgrading from PowerPath 4.0.2 or later, you need not reboot the host in most cases. Refer to *Rebooting after the Upgrade* on page 1-15 for more information.

- Verify the PowerPath installation.
- Verify the PowerPath commands are in your path.
- Verify that PowerPath devices are configured on the host.
- Optionally, reconfigure applications to use emcpower devices.
- If necessary, disable the PowerPath setup script in the `.login` or `profile` file.

The following sections describe these procedures.

In addition, if you are upgrading from PowerPath 4.0.2 or later and you stopped EMC ControlCenter or Navisphere agents before the upgrade, restart those agents now.

Register PowerPath on the Host

If you have previously registered an earlier version of PowerPath on the host, PowerPath will use your old key; you need not reregister. (Note, however, that if your previous registration was for evaluation purposes, and the registration has expired, you must reregister before you reboot the host).

To register the PowerPath software:

1. Enter:

```
/etc/emcprep -install
```

You see the following output:

```
===== EMC PowerPath Registration =====  
Do you have a new registration key or keys to enter? [n]
```

2. Type **y** and press ENTER.

You see the following output:

```
Enter the registration key(s) for your product(s),
one per line, pressing Enter after each key.
After typing all keys, press Enter again.
```

```
Key (Enter if done):
```

3. Enter the PowerPath registration key and press ENTER.

If you entered a valid key, you see the following output:

```
1 key(s) successfully added.
Key successfully installed.
```

```
Key (Enter if done):
```

4. Press ENTER. You see the following output:

```
1 key(s) successfully registered.
```

Registering PowerPath after you install the software and before you reboot the host sets the appropriate default load balancing and failover policy. If you reboot the host before you register PowerPath, you must run the `powermt set policy` to reset the load balancing and failover policy. Refer to the *PowerPath Product Guide* or the `powermt(1)` man page for more information on `powermt set policy`.

Remove the CD-ROM

1. If the CD-ROM volume management daemon `vold` is running, unmount and eject the CD-ROM. Enter:

```
eject
```

and remove it from the drive.

2. If `vold` is not running, unmount the CD-ROM. For example, enter:

```
umount /cdrom/cdrom0
```

and, after the CD-ROM unmounts, eject it and remove it from the drive.

Reboot the Host

Registering PowerPath after you install the software and before you reboot the host sets the appropriate default load balancing and failover policy. (Refer to *Register PowerPath on the Host* on page 1-23.) If you reboot the host before you register PowerPath, you must run the `powermt set policy` to reset the load balancing and failover policy. Refer to the *PowerPath Product Guide* or the `powermt(1)` man page for more information on `powermt set policy`.

If you are upgrading from PowerPath 4.0.2 or later, you need not reboot the host in most cases. Refer to *Rebooting after the Upgrade* on page 1-15.

If you are installing PowerPath for the first time, or if you are upgrading from PowerPath 4.0.1 or earlier, you must reboot now.

To reboot the host, enter:

```
reboot -- -r
```

If the `sd` or `ssd` driver does not exist on the host, you see one of the following messages during boot:

```
WARNING: forceload of drv/sd failed
```

or

```
WARNING: forceload of drv/ssd failed
```

You can safely ignore this warning.

Verify the PowerPath Installation

1. Verify that PowerPath is installed properly on the host. Enter:

```
pkginfo -l EMCpower
```

You should see output similar to this:

```
PKGINST:    EMCpower
NAME:       EMC PowerPath
CATEGORY:   system
ARCH:       sparc
VERSION:    4.3.0_bxxx
BASEDIR:    /opt
VENDOR:     EMC
PSTAMP:     cambridge951018123443
INSTDATE:   Mar 14 2003 08:36
STATUS:     completely installed
FILES:      292 installed pathnames
             5 shared pathnames
             17 directories
             139 executables
             137622 blocks used (approx)
```

2. If you installed the PowerPath Volume Manager VCS Agent, verify that the agent is installed properly on the host. Enter:

```
pkginfo -l EMCvlg
```

You should see output similar to this:

```
PKGINST:    EMCvlg
NAME:       EMC Enterprise Cluster PowerPath Volume Manager VCS Agent
CATEGORY:   system
ARCH:       sparc
VERSION:    4.3.0
BASEDIR:    /opt
VENDOR:     EMC Corporation
DESC:       EMC Enterprise Cluster PowerPath Volume Manager VCS Agent
PSTAMP:     410b089_030904
INSTDATE:   Sep 10 2003 10:39
EMAIL:      support@emc.com
STATUS:     completely installed
FILES:      25 installed pathnames
             1 shared pathnames
             12 directories
             12 executables
             393 blocks used (approx)
```

3. Verify that the PowerPath and PowerPath Volume Manager kernel extensions are loaded on the host. Enter:

```
modinfo | grep emcp
```

You should see output similar to this:

```
23 102697a4 38fe 230 1 emcpsf (PP SF 4.3.0)
41 102ec4f5 57758 229 1 emcp (PP Driver 4.3.0)
42 780e0000 1fc68 - 1 emcpmp (PP MP Ext 4.3.0)
43 78102000 4306d - 1 emcpmpc (PP MPC Ext 4.3.0)
44 78146000 22cd2 - 1 emcpmpaa (PP MPAA Ext 4.3.0)
45 7816a000 2f924 - 1 emcpmpap (PP MPAP Ext 4.3.0)
46 7819c000 d1a59 - 1 emcpsapi (PP SAPI Ext 4.3.0)
47 1033f799 143da - 1 emcpcg (PP CG Ext 4.3.0)
48 10352c9b 3568 - 1 emcphr (PP HR Ext 4.3.0)
49 78270000 5ef42 - 1 emcpsm (PV SM Ext 4.3.0)
50 782f2000 8b196 - 1 emcpsc (PV SC Ext 4.3.0)
51 78380000 72421 - 1 emcpevm (PV EVM Ext 4.3.0)
52 783f4000 3aa88 - 1 emcpdpm (PV DPM Ext 4.3.0)
53 78165d5a 269 - 1 emcpioc (PP PIOC 4.3.0)
```

Verify that PowerPath Commands Are in Your Path

Check that the PowerPath commands are in your path. PowerPath installation modifies the `/etc/profile` and `/etc/.login` files to run the PowerPath setup scripts, which modify the `PATH`, `LD_LIBRARY_PATH`, `LD_LIBRARY_PATH_64`, and `MANPATH` variables on the host. These changes take place after you reboot the host; if reboot was not required (refer to *Rebooting after the Upgrade* on page 1-15), you must log out and then log back in to force these changes.

If you find that the PowerPath commands are not in your path, check the `/.profile`, `/.login`, and `/.cshrc` scripts to see if they are overwriting the changes made by the PowerPath setup scripts.

The installation changes your path to point to the appropriate CLIs (32 bit or 64 bit). We recommend that you accept the path the installation adds; changing the path is not supported.

Verify PowerPath Devices Are Configured on the Host

To verify that PowerPath devices are configured on the host:

1. Enter:

```
powermt display dev=all
```

You should see output like the following:

```
Pseudo name=emcpower0a
Symmetrix ID=000000003269;
Logical device ID=048
state=alive; policy=SymmOpt; priority=0; queued-I/Os=0
=====
----- Hosts ----- - Sys - - I/O Paths - --- Stats ----
###   HW Path           I/O Path Interf.  Mode   State Q-I/Os Errors
=====
2300  sbus@f/QLGC,isp@2  c2t1d0   SA 16aB  active alive   0    0
2301  sbus@f/QLGC,isp@2  c1t1d0   SA 16aA  active alive   0    0
2302  sbus@f/QLGC,isp@2  c3t1d0   SA 16bB  active alive   0    0
```

2. If the output of `powermt display dev=all` indicates that some storage system logical devices are not configured as PowerPath devices:

- a. Configure any missing logical devices. Enter:

```
powercf -q
powermt config
```

- b. Rerun `powermt display dev=all` to confirm that:

- These logical devices are configured as emcpower devices.
- The correct failover and load balancing policy is set. (For information about `powermt` and load-balancing policies, see the *PowerPath Product Guide*.)

If you plan to enable R1/R2 boot disk failover, see *R1/R2 Boot Failover Support* on page 6-4.

- c. If the failover and load balancing policy is not set correctly, run the command `powermt set policy`, setting the policy that applies to your storage systems.

Reconfigure Applications to Use emcpower Devices

If you plan to use native names (recommended) with a volume manager, file system application, or database manager, you can ignore this section. PowerPath load balancing and failover functionality are enabled with no further action on your part.

However, if you plan to use emcpower devices with a third-party volume manager, file system application, or database manager, you must reconfigure the application to use emcpower devices. Refer to *Installing and Configuring emcpower Devices With Solaris Applications*, which is available on the Powerlink website (<http://powerlink.emc.com>).

For more information on native and pseudo names, refer to *Device Naming* on page 6-5. If you are unsure whether to use emcpower or native names, contact the EMC Customer Support Center.

Disable the Setup Script

The PowerPath setup script, `emcpv_setup`, may interfere with other applications in your environment (for example, SAP). If this is the case, depending on your shell, disable the script as follows:

For	Delete	From
csh	<code>source /basedir/EMCpower/scripts/emcpv_setup.csh</code>	<code>/etc/.login</code>
ksh	<code>./basedir/EMCpower/scripts/emcpv_setup.sh</code>	<code>/etc/profile</code>

Note that users who run PowerPath administrative commands must then enable the script elsewhere, for example, in their local `.profile` or `.login` files.

Installing the PowerPath Volume Manager VCS Agent

You can install the VCS Agent at the time you install PowerPath, or you can install the agent separately. This section describes how to install the agent separately.

Mount the CD-ROM

1. Log in as root.
2. Insert the CD-ROM into the CD-ROM drive.

If the CD mounts automatically, continue with *Install the Software*, which follows.

If the CD does not mount automatically, you must mount it manually. Continue with step 3.

3. Mount the CD on your file system. For example, to mount the CD on `/cdrom/cdrom0`, enter:

```
mount -F hsfs -r /dev/dsk/cxydzs0 /cdrom/cdrom0
```

where *x*, *y*, and *z* are values specific to the host's CD-ROM drive. For example:

```
mount -F hsfs -r /dev/dsk/c0t2d0s0 /cdrom/cdrom0
```

Install the VCS Agent Software

1. If you do not have a graphics terminal, run the `script filename` command to record `pkgadd` output in the specified file. (After `pkgadd` completes, use CTRL-D to stop recording the output.)
2. Change to the `/mount_point/UNIX/SOLARIS` directory. For example, enter:

```
cd /cdrom/cdrom0/UNIX/SOLARIS
```

3. Start the installation program. Enter:

```
/usr/sbin/pkgadd -d .
```

4. You see the following prompt:

The following packages are available:

```
1 EMCpower  EMC PowerPath
              (all) 4.3.0_bxxx
2 EMCvg     EMC Enterprise Cluster PowerPath Volume
              Manager VCS Agent
              (sparc) 4.3.0
```

Select package(s) you wish to process (or 'all' to process all packages). (default: all) [?,??,q]:

Enter 2 and press ENTER.

5. You see the following prompt:

Enter path to package base directory [?,q]

Enter the path to the directory in which VCS is installed. We recommend /opt as the base directory.

6. The installation program displays the following prompt if it detects files on the host that conflict with files it will install:

Do you want to install these conflicting files [y, n, ?, q]

Enter y and press ENTER.

7. The screen displays information about the installation, ending with the following prompt:

Installation of <EMCvg> was successful.

The following packages are available:

```
1 EMCpower  EMC PowerPath
              (all) 4.3.0_bxxx
2 EMCvg     EMC Enterprise Cluster PowerPath Volume
              Manager VCS Agent
              (sparc) 4.3.0
```

Select package(s) you wish to process (or 'all' to process all packages). (default: all) [?,??,q]:

8. Enter q and press ENTER.

The PowerPath Volume Manager VCS Agent is now installed on the host.

To verify that the agent is installed properly on the host, enter:

```
pkginfo -l EMCvg
```

You should output similar to this:

```
PKGINST: EMCvg
NAME: EMC Enterprise Cluster PowerPath Volume Manager VCS Agent
CATEGORY: system
ARCH: sparc
VERSION: 4.3.0
BASEDIR: /opt
VENDOR: EMC Corporation
DESC: EMC Enterprise Cluster PowerPath Volume Manager VCS Agent
PSTAMP: 410b089_030904
INSTDATE: Sep 10 2003 10:39
EMAIL: support@emc.com
STATUS: completely installed
FILES: 25 installed pathnames
       1 shared pathnames
       12 directories
       12 executables
       393 blocks used (approx)
```

Refer to *Configuring VCS to Recognize PowerPath Volume Manager Resources* on page 2-9, for information on configuring the agent.

Error Messages

Error, warning, and informational messages returned by the PowerPath installation process are described in the *PowerPath Product Guide*, Chapter 6, *PowerPath Messages*.

PowerPath in a Cluster Environment

This chapter describes how to install and configure PowerPath in Solaris cluster environments. For more general information on clustering, refer to the *Symmetrix High Availability Environment Product Guide* (Symmetrix systems) or the *EMC Installation Roadmap for CX-Series and FC-Series Storage Systems* (CLARiiON systems).

PowerPath Volume Manager is currently supported in LEGATO Automated Availability Manager (AAM) and VERITAS Cluster Server (VCS) clusters only.

- ◆ PowerPath in a LEGATO 5.1 (or later) AAM Cluster 2-2
- ◆ PowerPath in a Sun Cluster 3.x..... 2-4
- ◆ PowerPath in a VERITAS Cluster Server Cluster..... 2-6
- ◆ Setting Major and Minor Numbers 2-11

PowerPath in a LEGATO 5.1 (or later) AAM Cluster

This section describes how to:

- ◆ Install PowerPath and LEGATO® Automated Availability Manager (AAM) in a new cluster, that is, where neither the PowerPath nor the AAM software is installed on any host to be included in the cluster.
- ◆ Integrate PowerPath into an existing AAM cluster.

Installing PowerPath in a New AAM Cluster

To install and configure PowerPath and AAM 5.1 (or later) when neither PowerPath nor AAM is installed:

1. Prepare the cluster environment. Refer to the relevant AAM documentation. (In Symmetrix or CLARiiON environments, refer also to the *EMC Host Connectivity Guide for Sun Solaris* or the *EMC Installation Roadmap for CX-Series and FC-Series Storage Systems*.)
2. Use the Solaris `format` utility to verify that all storage system devices are seen by each host.
3. Install PowerPath on all nodes. Refer to Chapter 1, *Installing PowerPath*.
4. Install AAM 5.1 (or later) on all nodes. Refer to the relevant AAM documentation.
5. Add PowerPath Volume Manager resources to AAM resource groups. Refer to the relevant AAM documentation.
6. If the AAM agent is not already started, start the agent on each node in the cluster, using either the `ft_startup` command or the Management Console.
7. Verify that the resource group is up and running, and use either the `ftcli` command `ListResourceGroups` or the Management Console to verify that the resource group can fail over to all nodes in the cluster.
8. Set a common minor number for every PowerPath Volume Manager volume in the cluster. Refer to *Setting Major and Minor Numbers* on page 2-11 for information on setting the minor number.

Integrating/ Upgrading PowerPath into an Existing AAM 5.1 (or later) Cluster

To integrate PowerPath into an existing AAM 5.1 (or later) cluster, follow these steps on each cluster node, one node at a time:

1. Stop cluster services on the node using either the `ft_shutdown` command or the Management Console.
2. Install or upgrade PowerPath on the node. Refer to Chapter 1, *Installing PowerPath*.
3. Start cluster services on the node using either the `ft_startup` command or the Management Console, and wait for the node to be fully reintegrated into the cluster.

PowerPath in a Sun Cluster 3.x

This section describes how to:

- ◆ Install PowerPath and Sun Cluster 3.x in a new cluster, that is, where neither the PowerPath nor the Sun Cluster 3.x software is installed on any host to be included in the cluster.
- ◆ Integrate PowerPath into an existing Sun Cluster 3.x cluster.

PowerPath supports only native devices with Sun Cluster 3.x. Pseudo (emcpower) devices are not supported.

Installing PowerPath in a New Sun Cluster 3.x

To install and configure PowerPath and Sun Cluster 3.x when neither PowerPath nor Sun Cluster 3.x is installed:

1. Prepare the cluster environment. Refer to the relevant Sun Cluster documentation. (In a Symmetrix or CLARiON environment, refer also to the *EMC Host Connectivity Guide for Sun Solaris* or the *EMC Installation Roadmap for CX-Series and FC-Series Storage Systems*.)
2. Install PowerPath on all nodes. Refer to Chapter 1, *Installing PowerPath*.
3. Install Sun Cluster 3.x on all nodes. Refer to the relevant Sun Cluster documentation.
4. Configure the quorum device on all nodes.

You cannot use gatekeepers and VCMDDB devices as quorum devices.

5. Initialize the root disk group on all nodes.
6. Initialize PowerPath devices on all nodes.
7. Start cluster services on the master node.
8. Designate/create shared disk groups on the master node.
9. Create logical volumes from the designated shared disks.
10. Register the disk group.

Integrating/ Upgrading PowerPath into an Existing Sun Cluster 3.x

To integrate PowerPath into an existing Sun Cluster 3.x, follow these steps on each cluster node, one node at a time.

1. Stop cluster services on the node. Enter:

```
boot -x
```

2. Install or upgrade PowerPath on the node. Refer to Chapter 1, *Installing PowerPath*.
3. Start cluster services on the node. Enter:

```
reboot
```

Wait for the node to be fully integrated into the cluster before continuing.

4. For Sun Cluster 3.1, Update 1 or later, if your environment includes Hitachi Lightning, HP xp, 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 and ESS devices:

```
powermt config  
powermt save
```

- c. Run the following command to verify that the pseudo devices have been created:

```
powermt display dev=all
```

PowerPath in a VERITAS Cluster Server 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.
- ◆ Configure VCS to recognize PowerPath Volume Manager resources.

For new installations of VxVM, use native `c#t#a#` devices. Do not initialize emcpower devices for VxVM using `powervxvm`.

If emcpower devices already 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.

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 or CLARiiON environments, refer also to the *EMC Host Connectivity Guide for Sun Solaris* or the *EMC Installation Roadmap for CX-Series and FC-Series Storage Systems*.)
 - b. Use the Solaris `format` utility to verify that all storage system devices are seen by each host.
 - c. Install PowerPath and, if you are using PowerPath Volume Manager in the cluster, the PowerPath Volume Manager VCS Agent. Refer to Chapter 1, *Installing PowerPath*. 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.

If you are using PowerPath Volume Manager in the cluster, proceed with step 2. If you are not using PowerPath Volume Manager, go to step 3.

2. If you are using PowerPath Volume Manager, configure VCS to recognize the PowerPath Volume Manager VCS Agent as a resource type by editing the `/etc/VRTSvcs/conf/config/main.cf` file on one host in the cluster. Refer to *Configuring VCS to Recognize PowerPath Volume Manager Resources* on page 2-9.
3. On each node in the cluster, define the resources (for example, PowerPath Volume Manager volume groups or VxVM volumes) that make up the service group. (You will configure the service group in step 4.)
4. On one host in the cluster (if you are using PowerPath Volume Manager in the cluster, this should be the host where you edited `main.cf` in step 2; otherwise, this can be any host):
 - a. Configure the service group by adding the resources you defined in step 3 to the `/etc/VRTSvcs/conf/config/main.cf` file. The disk or logical device resources should use native `c#t#a#` devices.

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.
 - c. If you are using PowerPath Volume Manager, verify that VCS recognizes the PowerPath Volume Manager VCS Agent as a resource type. Enter:

```
hatype -list | grep EMC
```

`EMCvg` should be listed in the command output.
5. Start cluster services on each remaining node in the cluster. These hosts rebuild their local configuration files from the `main.cf` file you edited in step 4.a.

6. On each node in the cluster:
 - a. Verify that the service group is up and running, and use either the VCS GUI or the `hagrps` command to verify that the service group can successfully fail over to all hosts in the cluster.
 - b. Set a common minor number for every PowerPath Volume Manager volume in the cluster. Refer to *Setting Major and Minor Numbers* on page 2-11 for information on setting the minor number.
 - c. Add other service groups as needed.

Integrating/ Upgrading PowerPath into an Existing VCS Cluster

To integrate 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, and, if you are using PowerPath Volume Manager, install the PowerPath Volume Manager VCS Agent. Refer to Chapter 1, *Installing PowerPath*.
3. If you are using PowerPath Volume Manager, configure VCS to recognize the PowerPath Volume Manager VCS Agent as a resource type by editing the `/etc/VRTSvcs/conf/config/main.cf` file on one node in the cluster. Refer to *Configuring VCS to Recognize PowerPath Volume Manager Resources* on page 2-9.
4. Use 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.
5. Use `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.

Configuring VCS to Recognize PowerPath Volume Manager Resources

After the PowerPath Volume Manager VCS Agent has been installed on a host, you must configure VCS to recognize the agent. To do so, edit the VCS configuration file, `main.cf`, to include the `EMCTypes.cf` file.

The `EMCTypes.cf` file, which is installed on the host during installation of the PowerPath Volume Manager VCS Agent, defines the `EMCvrg` resource type. The `EMCvrg` resource type describes a PowerPath Volume Manager volume group to VCS.

Refer to the *PowerPath for UNIX Volume Manager User's Guide* for a description of the attributes in the `EMCvrg` resource type.

EMC recommends that you leave the `OnlineTimeout` parameter for VERITAS Cluster Server (VCS) at the default value, 300 seconds. If the volume groups configured as `EMCvrg` resources contain more than approximately 500 volumes, set `OnlineTimeout` to 600 seconds.

These values are guidelines only. The threshold for increasing the value of the `OnlineTimeout` parameter is highly configuration dependent, influenced primarily by the structure of the volumes. Experiment with different values to determine the optimal settings for your configuration.

Editing main.cf to Include EMCTypes.cf

Edit the `/etc/VRTSvcs/conf/config/main.cf` file on one node. The changes are propagated to the other nodes in the cluster when you start VCS on those nodes.

You do not need to perform this procedure if GeoSpan 2.0 is installed on the host.

To edit the `main.cf` file:

1. Verify that the PowerPath Volume Manager VCS Agent is installed on the host. Enter:

```
pkginfo -l EMCvg
```

2. Save the current `/etc/VRTSvcs/conf/config/main.cf` file to a backup location. For example, enter:

```
cp /etc/VRTSvcs/conf/config/main.cf /etc/VRTSvcs/conf/config/main.cf.orig
```

3. Add the following line to the beginning of `/etc/VRTSvcs/conf/config/main.cf`:

```
include "EMCTypes.cf"
```

4. Verify that there are no errors in the edited `main.cf` file. Enter:

```
hacf -verify config
```

If the verification is successful, no output is displayed. Otherwise, note the error and correct it in the `main.cf` file.

Setting Major and Minor Numbers

With NFS in a cluster environment:

- ◆ The PowerPath driver, `emcp`, must use the same major number on every node in the cluster.
- ◆ Every PowerPath Volume Manager volume in the cluster must use the same minor number on every node in the cluster.

Otherwise, you may need to restart NFS clients when failover or failback occurs.

Setting a Common Major number

You set a common major number for every host in the cluster during PowerPath installation (refer to Chapter 1, *Installing PowerPath*).

Setting a Common Minor Number

Minor numbers are assigned as follows:

- ◆ When you create a PowerPath Volume Manager volume, a minor number is assigned to that volume on the host where the volume is created.
- ◆ When you then test failover on the nodes in the cluster, a minor number is established for every failover host in the cluster.

After you test failover, and before you start to use the new volume, you must ensure that the volume has the same minor number on every node in the cluster. To do so, use the following procedure:

The commands in this procedure work whether the volume is imported or deported. Note, however, that depots must be done without using the `-removePaths` option.

1. On each node in the cluster, run the `powervol getminor` command to determine the minor number assigned to the volume on that node. Refer to the `powervol(1)` man page for information on `powervol getminor`.
2. If step 1 reveals inconsistent minor numbers, use the `emcpminor` command to find one minor number that can be used for the volume on every host. Refer to the `emcpminor(1)` man page for information on `emcpminor`.

While running `emcpminor`, avoid running any other commands that might claim minor numbers.

3. On every node in the cluster (except any node that is already using the minor number), run the `powervol setminor` command to set the minor number you chose in step 2. Refer to the `powervol(1)` man page for information on `powervol setminor`.

The `powervol setminor` command fails if the volume is open for I/O. It also fails if the minor number is in use, unless you specify the `-autoselect` option.

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 3-2
- ◆ Configuring a PowerPath Native Device as the Boot Device..... 3-3
- ◆ Moving the Boot Device to an emcpower Device 3-8
- ◆ Removing PowerPath Control over a Boot Device 3-11

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 HBA on the emcpower device must support booting in FCODE. Examples are the Emulex LP8000 and the JNI FCE2-1063 adapters.

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.

Configuring a PowerPath Native Device as the Boot Device

Partitioning the Boot Device

Partition the boot device.

1. Use the `format` command to verify that the sizes of the partitions on the storage system device chosen for the Solaris installation are large enough to copy the current OS partitions. Examine the partitions of the host source drive where the current OS resides:
 - a. At the `%>` prompt, enter `format` and press `ENTER`.
 - b. When prompted choose the internal boot device and press `ENTER`.
 - c. At the `format>` prompt, enter `partition` and press `ENTER`.
 - d. At the `partition>` prompt, enter `print` and press `ENTER`.

Sample output:

Part	Tag	Flag	Cylinders	Size
0	root	wm	0-335	24.69 MB
1	swap	wm	336-792	1169.59 MB
2	backup	um	0-5846	2.12GB
3	unassigned	wm	0	0
4	unassigned	wm	0	0
5	unassigned	wm	0	0
6	usr	wm	739-2305	561.40 MB
7	home	wm	2306-5846	1.28GB

2. Select the storage system device on which to install the Solaris operating system as follows:
 - a. At the `partition` prompt, enter `quit` and press `ENTER` to return to the `format` menu.
 - b. At the `format` prompt, enter `disk` and press `ENTER`.

Information similar to the following is displayed:

AVAILABLE SELECTIONS

```
...
1. c0t0d0 <SUN4,2G cyl 3880 alt 2 hd 16 sec 135>
   /pci@1f,4000/scsi@3/sd@0,0
2. c3t0d0 <EMC-SYMMETRIX-5265 cyl 4088 alt 2 hd 15 sec
   64> /pci@1f,4000/QLGC,qla@4/sd@0,0
...
specify disk (enter its number):
```

- c. Enter **2** and press `ENTER`.
3. Repartition the storage system device (if necessary) to model the partition information from the host source drive in order to match the drive configuration:
 - a. At the `format>` prompt, enter **partition** and press `ENTER`.
 - b. At the `partition>` prompt, enter **modify** and press `ENTER`.
 - c. After the partition in the storage system device has been created, enter **modify** and press `ENTER` at the `format` prompt.
 - d. Verify that the partition size allocated is greater than or equal to the size you specified.

Creating Filesystems

Create filesystems on the required partitions in the designated storage system boot device.

The following example shows a subset of the filesystems you would create when copying the OS to a Symmetrix device of Target 0 LUN 0 on controller 3:

```
newfs /dev/rdisk/c3t0d0s0
newfs /dev/rdisk/c3t0d0s6
newfs /dev/rdisk/c3t0d0s7
```

To use a different storage system device, change `t0d0` to `tStorage_Device_target_numberdLUN_number`. To use a different controller, change `c3` to `ccontroller_number`. You must create both the root partition (slice 0) and the `/usr` partition (slice 6). You should also create other partitions (`/export/home`, `/var`, `/opt`) to mirror the current layout. Use the `newfs` command to create the filesystems.

Installing a Bootblk

Install a *bootblk* (the information required by the host to boot the operating system) on the root partition of the storage system boot device.

1. Use the `uname -a` command to determine the architecture of the workstation. This determines the directory in `/usr/platform` where the *bootblk* will be found. The following example shows `sun4u` as the architecture, and subdirectory for `/usr/platform`.

Example: At the `%>` prompt, enter `uname -a` and press `ENTER` to display a line similar to the following:

```
SunOS patriot 5.8 Generic_108528-15 sun4u Sparc SUNW,Ultra-60
```

2. At the `%>` prompt, enter a command like the following to install a *bootblk* onto the storage system boot device:

```
/usr/sbin/installboot /usr/platform/sun4u/lib/fs/ufs/bootblk\  
/dev/rdisk/ccontroller_numberstorage_system_device_target_numberd0s0
```

For example, enter the following command to install a *bootblk* onto a boot device of Target 0 LUN 0 on controller 3:

```
/usr/sbin/installboot /usr/platform/sun4u/lib/fs/ufs/bootblk /dev/rdisk/c3t0d0s0
```

Copying Required Files

Use `ufsdump` and `ufsrestore` to copy the required files and directories from the host source drive to the storage system boot device.

1. At the `%>` prompt, enter `mount /dev/dsk/c3t0d0s0 /mnt` and press `ENTER` to mount the root directory of the storage system device to `/mnt`.
2. Use the `ufsdump` command to create the required directory structure on the new storage system device and copy the files.

In the following example, the current operating system is located on `/dev/dsk/c0t0d0`.

Example: At the `%>` prompt, enter the following command and press `ENTER`:

```
ufsdump 0f - /dev/dsk/c0t0d0s0 | ( cd /mnt; ufsrestore rf - )
```

When the above command completes, the storage system device (`c3t0d0s0`) will have the complete image of the root partition.

3. In the `/mnt/etc` directory update the `vfstab` file to indicate the storage system boot device address to be mounted during boot. Modify all partitions that will be located on the storage system boot device.
 - a. At the `%>` prompt, enter `cd /mnt/etc` and press `ENTER`.
 - b. At the `%>` prompt, enter `vi vfstab` and press `ENTER`.

Example — Before modification:

```
#device          device          mount FS    fsck  mount  mount
#to mount to    fsck            point type  pass  at boot options

/dev/dsk/c0t0d0s1 -                -      swap  -      no     -
/dev/dsk/c0t0d0s0 /dev/rdisk/c0t0d0s0 /      ufs   1      no     -
/dev/dsk/c0t0d0s6 /dev/rdisk/c0t0d0s6 /usr   ufs   1      no     -
swap - /tmp tmpfs - yes -
```

Example — After modification:

```
#device          device          mount FS    fsck  mount  mount
#to mount to    fsck            point type  pass  at boot options
#
/dev/dsk/c3t0d0s1 -                -      swap  -      no     -
/dev/dsk/c3t0d0s0 /dev/rdisk/c3t0d0s0 /      ufs   1      no     -
/dev/dsk/c3t0d0s6 /dev/rdisk/c3t0d0s6 /usr   ufs   1      no     -
swap - /tmp tmpfs - yes -
```

To increase system performance, you can leave the swap partition on the internal boot drive by leaving the fourth line (`/dev/dsk/c0t0d0s1 - - swap- no -`) unchanged.

4. At the `%>` prompt, enter `umount /mnt` and press `ENTER` to unmount the root partition.
5. Repeat steps 1, 2, and 4 for the `/usr` partition. In the current root directory:
 - a. At the `%>` prompt, enter `mount /dev/dsk/c3t0d0s6 /mnt` and press `ENTER`.
 - b. At the `%>` prompt, enter the following command and press `ENTER`:


```
ufsdump 0f - /dev/dsk/c0t0d0s6 | ( cd /mnt; ufsrestore rf -)
```
 - c. At the `%>` prompt, enter `umount /mnt` and press `ENTER`.
6. Repeat steps 1, 2, and 4 for any other partitions to be mounted from the storage system boot device.

Modifying OpenBoot

You need modify OpenBoot only if you use Fibre Channel HBAs.

The procedure for modifying OpenBoot depends on your HBAs. For CLARiiON and Symmetrix environments, refer to the following documentation:

If you use	Refer to
Emulex HBAs	<i>EMC Fibre Channel with Emulex Host Bus Adapters in the Solaris Host Environment</i> http://www.emulex.com/ts/docoem/framemc.htm
QLogic HBAs	<i>EMC Fibre Channel with QLogic Host Bus Adapters in the Solaris Environment</i> http://download.qlogic.com/drivers/5419/QLogic_Solaris.pdf
JNI HBAs	<i>EMC Symmetrix Fibre Channel with JNI Emerald Host Bus Adapters in the Solaris Environment</i> http://www.jni.com/Drivers/Files/Solaris/FCE2-6412/emc/JNI_EM3_SOLARIS.PDF
	<i>EMC Symmetrix Fibre Channel with JNI Tachyon Host Bus Adapters in the Solaris Environment</i> http://www.jni.com/Drivers/Files/Solaris/FC64-1063/emc/JNI_Tach_Solaris.pdf
	<i>EMC Symmetrix Fibre Channel with JNI Emerald 4 Host Bus Adapters in the Solaris Environment</i> http://www.jni.com/Drivers/Files/Solaris/FCX2-6562/emc/JNI_Em4_Solaris.pdf

For third-party array environments, refer to the appropriate documentation from your vendor.

Setting Up the Boot Alias

Set up a new alias for your new boot disk. At the `OK>` prompt:

1. Enter `nvalias your_alias your_device` and press `ENTER`.
Example:
`nvalias storagedisk /pci@1f,4000/QLGC,qla@4/sd@0,0`
2. Enter `nvstore` and press `ENTER` to store your new boot name.
3. Enter `setenv boot-device storagedisk` and press `ENTER` to change the boot device to the new alias.
4. Enter `setenv auto-boot? false` and press `ENTER`.
5. Enter `reset -all`.
6. Enter `boot`.

Moving the Boot Device to an emcpower Device

Important: Do not boot from an emcpower device on an unsupported HBA. For information on HBAs that are supported with Symmetrix and CLARiiON storage systems, refer to the *EMC Support Matrix* (<http://powerlink.emc.com>). For information on HBAs supported with other storage systems, refer to the appropriate documentation from your vendor.

To move a boot device from a native device to an emcpower device:

1. Boot from the native device.
2. Install PowerPath. Refer to Chapter 1, *Installing PowerPath*.
3. Locate the native device from which you are booting, and correlate this device to an emcpower device. Enter:

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

The output is similar to the following:

```
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. Enter:

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

The output is similar to the following:

```
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. Enter:

```
cp /etc/system /etc/system.no_EMCPower
cp /etc/vfstab /etc/vfstab.no_EMCPower
```

6. Using a text editor such as `vi`, add the following line below the `forceload: drv/emcp` statement:

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

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

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



CAUTION

4. Enter:

```
TERM=sun-cmd  
export TERM
```
5. 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.
6. Shut down the host. Enter:

```
shutdown -y -g5 -i0
```
7. At the `ok` prompt, enter:

```
eject
```
8. Remove the Solaris Operating System CD-ROM from the host's CD-ROM drive.
9. Reboot the host. Enter:

```
boot
```

Removing PowerPath Control over a Boot Device

If your PowerPath installation uses a storage system device as the boot device, use the following procedure to remove PowerPath 4.3 control over the boot device:

1. Determine whether the boot path from the host to the storage system has changed since the storage system boot device was configured.

- a. Examine the `/etc/vfstab.no_EMCPower` file to identify the native name for the boot device; for example, `c3t0d0`. Enter:

```
cat /etc/vfstab.no_EMCPower
```

- b. Determine the boot path associated with this native device. For example, enter:

```
ls -al /dev/dsk/c3t0d0s0
```

An example of a boot path is

```
pci@1f,4000/QLGC,qla@4/sd@0,0.
```

- c. Compare the boot path from step 1b. with the original boot path. Enter:

```
cat /etc/nvramrc.orig
```

2. If the boot path listed in step 1b. differs from that in the `nvramrc.orig` file, update the boot path to reflect the value listed in step 1b. For example, enter:

```
eeeprom nvramrc="devalias storagedisk /pci@1f,4000/QLGC,qla@4/sd@0,0"
```

3. Verify that the boot path was changed to the new value. Enter:

```
eeeprom
```

4. Restore the versions of `/etc/system` and `/etc/vfstab` that do not contain references to PowerPath. Enter:

```
cp /etc/system.no_EMCPower /etc/system
```

```
cp /etc/vfstab.no_EMCPower /etc/vfstab
```

5. If you did not reboot the host after removing PowerPath, do so now. Enter:

```
reboot -- -r
```

If the host fails to boot, refer to *Recovery Procedure* on page 3-9 for suggested actions.

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 4-2
- ◆ Migrating from IBM Subsystem Device Driver (SDD) 4-4
- ◆ Migrating from Hitachi Dynamic Link Manager (HDLM) 4-6

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 `cXtYdZ` devices belonging to a particular EVA or HSG80 LUN (there is one such native device for each path) with one unique `cPt0dZ` 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.

Migration Procedure

To migrate 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 and come up unmanaged on the HP HSx devices.

3. Uninstall Secure Path.

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; run the following command:

```
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`). Refer to the *EMC PowerPath for UNIX Release Notes* for details.
 - Make sure the `sd.conf` file has LUN number entries for EVA or HSG80 targets.
6. Reboot the host with device reconfiguration. Enter:

```
reboot -- -r
```

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 the command `powermt display dev=all class=hphsx` 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`.

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 `cXtYdZ` 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 `cXtYdZ` 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 `cXtYdZ` 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 `cXtYdZ` path. Since VERITAS/DMP does support ESS arrays, DMP can provide multipathing for ESS LUNs.

Migration Procedure

To migrate 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 and come up unmanaged on the ESS devices.
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 the command `powermt manage class=ess`.
 - c. Run the command `powermt display dev=all class=ess` to verify that PowerPath can now 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 cPtQdZ 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 ChPtQdZ 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.

Migration Procedure

To migrate 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 and come up unmanaged on the Hitachi Lightning devices.
3. Uninstall HDLM, which requires a reboot with device reconfiguration (`reboot -- -r`).
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 the command `powermt manage class=hitachi`.

- c. Run the command `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`.

Removing PowerPath

This chapter describes how to remove PowerPath and the PowerPath Volume Manager VCS Agent from a Solaris host.

- ◆ Before You Remove PowerPath 5-2
- ◆ Removing PowerPath..... 5-3
- ◆ After You Remove PowerPath 5-4
- ◆ Removing PowerPath 4.3 Offline 5-6
- ◆ Removing the PowerPath Volume Manager VCS Agent 5-7

Before You Remove PowerPath

Before you remove PowerPath from the host:

- ❑ Check the Powerlink website (<http://powerlink.emc.com>) for the most current information. We update the PowerPath release notes periodically and post them on the Powerlink website.
- ❑ Stop any application and shut down any database that is using emcpower devices or a PowerPath Volume Manager volume. Unmount any file system mounted on a PowerPath Volume Manager volume.
- ❑ Unmount any HighRoad filesystem.

If you fail to unmount HighRoad filesystem before you uninstall PowerPath 4.3, you may see the following error message:

```
can't unload the module: Unknown error
Unload of modules failed.
```

- ❑ *If you are removing PowerPath from the host entirely (that is, you are not planning to reinstall PowerPath), remove all PowerPath Volume Manager volumes and volume groups allocated on the host:*

Run the `powervadm list` command to determine whether any volumes and volume groups are allocated on the host. If volumes and volume groups are allocated, destroy them. Refer to *Destroying a Volume* and *Destroying a Volume Group* in the *PowerPath for UNIX Volume Manager User's Guide*.

- ❑ If you have a database partition, discontinue use of 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. Refer to *Removing PowerPath Control over a Boot Device* on page 3-11.

Removing PowerPath

To remove PowerPath from a Solaris host:

1. Log in as root.
2. Start the uninstall program. Enter:

```
/usr/sbin/pkgrm EMCpower
```

The screen displays information like this:

```
The following package is currently installed:  
EMCpower EMC PowerPath  
      (all) 4.3.0_bxxx
```

```
Do you want to remove this package? [y,n,?,q]
```

3. Enter **y** and press ENTER to remove the package.

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 `.430.saved`:

- ◆ `/kernel/drv/emcp.conf`
- ◆ `/kernel/drv/emcpsf.conf`
- ◆ `/etc/powermt.custom`
- ◆ `/etc/emcp_registration`
- ◆ `/etc/emcp_devicesDB.dat`
- ◆ `/etc/emcp_devicesDB.idx`
- ◆ `/etc/PPVM_config`
- ◆ `/etc/PPVM_config_bak`

If the removal program detects existing files with the `.4x.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), remove these files from the host:

1. Enter the following command:

```
/etc/emcpv_cleanup
```

This command removes the PowerPath license and all supporting files.

`emcpv_cleanup` issues the following warning:

```
Warning: Attempt to remove saved configuration files
for PowerPath versions. Removing these files will
result in loss of saved configuration and upgrade
features. Proceed with caution!!
```

```
Do you want to continue [y,n,?] (default: n):
```

2. Enter **y** and press ENTER.

The uninstall program asks for confirmation:

```
Are you sure [y,n,?] (default: n):
```

If Saved Files Remain on the Host

3. Enter **y** and press ENTER.

`emcpv_cleanup` takes an optional argument: a revision number, without periods.

Normally, `emcpv_cleanup` without arguments removes all saved files. However, if after running the command without arguments, you discover saved files remaining on the host, you can run `emcpv_cleanup` again, specifying a revision number.

For example, if after removing PowerPath 4.3.0 and running `emcpv_cleanup` without arguments, you see saved files from PowerPath 4.1.0, run the following command:

```
emcpv_cleanup 410
```

Rebooting the Host

You need *not* reboot the host after you uninstall PowerPath if:

- ◆ The message at the end of the removal process indicated that reboot was not necessary.
- ◆ You plan to reinstall PowerPath 4.3
- ◆ You plan to upgrade to a later version of PowerPath

Otherwise, reboot the host by entering the following command:

```
reboot -- -r
```

If the `sd` or `ssd` driver does not exist on the host, you see one of the following messages during boot:

```
WARNING: forceload of drv/sd failed
```

or

```
WARNING: forceload of drv/ssd failed
```

You can safely ignore this warning.

Removing PowerPath 4.3 Offline

PowerPath 4.3 supports the `pkgrm -R` command. Use `pkgrm -R` to remove PowerPath:

- ◆ From a second root disk on a single host
- ◆ Via NFS
- ◆ When booted from a CD-ROM

Removing the PowerPath Volume Manager VCS Agent



CAUTION

Do not try to remove the EMCvg package while VCS with PPVM is up and running.

Complete the following steps to remove the PowerPath Volume Manager VCS Agent:

1. Log in as root.
2. Close the VCS configuration file. Enter:

```
haconf -dump -makero
```

3. Save the current `/etc/VRTSvcs/conf/config/main.cf` file to a backup location. For example, enter:

```
cp /etc/VRTSvcs/conf/config/main.cf /etc/VRTSvcs/conf/config/main.cf.orig
```

4. Identify all EMCvg resources. Enter:

```
hares -display -type EMCvg
```

5. Delete each EMCvg resource that you identified. Enter:

```
hares -delete resource_name
```

6. Delete the EMCvg resource type. Enter:

```
hatype -delete EMCvg
```

7. Verify the configuration changes. Enter:

```
hacf -verify config
```

If the verification is successful, no output is displayed. Otherwise, note the error and correct it in the `main.cf` file.

8. Remove the PowerPath Volume Manager VCS agent package. Enter:

```
/usr/sbin/pkgrm EMCvg
```

The screen displays information like this:

```
The following package is currently installed:
1 EMCvg      EMC Enterprise Cluster PowerPath
              Volume Manager VCS Agent
              (sparc) 4.3.0
```

```
Do you want to remove this package? [y,n,?,q]
```

9. Enter **y** and press ENTER to remove the package.

The screen displays the following information:

```
## Removing installed package instance <EMCvg>
```

```
Do you want to continue with the removal of this package
[y,n,?,q]
```

10. Enter **y** and press ENTER to continue with deinstallation.

The screen displays information about the removal process, ending with:

```
Removal of <EMCvg> was successful.
```

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

Throughout this chapter, many procedural steps use `powermt` commands. For detailed descriptions of these commands, refer to the *PowerPath Product Guide*.

- ◆ PowerPath and CLARiON Storage Systems..... 6-2
- ◆ Ensuring a Sufficient Stack Size 6-2
- ◆ Booting a Host with Built-In Fibre Channel Ports 6-3
- ◆ Rebooting and Custom Settings..... 6-4
- ◆ R1/R2 Boot Failover Support..... 6-4
- ◆ Device Naming 6-5
- ◆ Reconfiguring PowerPath Devices Online..... 6-10
- ◆ Dynamic Reconfiguration..... 6-11
- ◆ Upgrading Solaris 6-14
- ◆ `powercf` Configuration Utility..... 6-16
- ◆ Error Messages 6-17

PowerPath 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 CLARiiON configuration requirements, refer to the *EMC Installation Roadmap for CX-Series and FC-Series Storage Systems* and the *EMC Host Connectivity Guide for Sun Solaris*, available on the Powerlink website.

The *Host Connectivity Guide* 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 larger than the default provided by Solaris. The stack size is controllable through the `lwp_default_stksize` and `rpcmod:svc_run_stksize` variables in `/etc/system`.

Other products also require a change to the stack size. To ensure that PowerPath does not impact other products, the PowerPath installation checks `/etc/system` for a `lwp_default_stksize` value. If a value is present and is greater than the PowerPath requirement, the PowerPath installation does not modify the value. If a value is not present, or is present and is less than the PowerPath requirement, the installation changes the value to the minimum required by PowerPath:

- ◆ 0x6000 on a Solaris 7, 8, or 9 system
- ◆ 0x4000 on a Solaris 2.6 system

If you subsequently install another application that resets the stack size (for example VxVM or VxFS), make sure that `lwp_default_stksize` and `rpcmod:svc_run_stksize` are set to at least 0x6000 on a Solaris 7, 8, or 9 system, and at least 0x4000 on a Solaris 2.6 system. Other applications may reset the stack size to a value that is insufficient for PowerPath.

Moreover, if the `lwp_default_stksize` variable set in `/etc/system` contains an illegal value—one that is larger than the system-specific maximum or that is not aligned to page boundary—Solaris reverts to its default value, which, again, is insufficient for PowerPath.

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 `lwp_default_stksize` in `/etc/system` both are legal 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 social host adapters to storage system volumes disconnected or dysfunctional, PowerPath will not configure any social host adapter paths. After physically restoring the social connections, run the following commands to restore the paths in PowerPath:

On hosts running this OS	Run these commands
Solaris 7, 8, and 9	<pre>devfsadm powercf -q (optional) powermt config</pre>
Solaris 2.6	<pre>drvconfig; disks; devlinks powercf -q (optional) powermt config</pre>

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 (`/etc/powermt.custom.[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, enter the following commands:

```
/etc/powermt load file=/etc/powermt.custom.<desired version>
/etc/powermt save
```

Although you can restore an earlier `powermt.custom` file, it is not currently possible to restore an earlier `emcp_devicesDB` file.

R1/R2 Boot Failover Support

If a storage system device corresponding to a bootable emcpower device is mirrored via SRDF[®], it is possible in the event of a server failure at the local storage system 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

EMC supports the following specific R1/R2 configurations:

- ◆ Each boot host is connected to only one Symmetrix.
- ◆ The two hosts must have identical hardware.
- ◆ All R1 devices reside on one Symmetrix, *Symmetrix A*, and are visible only to a single host, *Host A*.
- ◆ All R2 devices reside on a separate Symmetrix, *Symmetrix B*, and are visible only to the identical host in reserve, *Host B*.
- ◆ Each R1 device has only one mirror. (Concurrent SRDF is not supported.)

- ◆ Neither host has non-mirrored devices, BCVs, or gatekeepers.
- ◆ SRDF is managed from either of the following two facilities:
 - EMC ControlCenter Management Server
 - Symmetrix Service Processor

R1/R2 Boot Procedure

R1/R2 boot support assumes that the systems are configured to boot from an emcpower device. If you plan to enable R1/R2 boot disk failover, after you install PowerPath, run the `powercfg -Z` command while booted on the R1 copy of the boot disk. This will update the `emcp.conf` file so that each entry contains both an R1 and an R2 Symmetrix volume ID for the pseudo (emcpower) device.

When the host on the R2 side boots, it is connected to a different Symmetrix system and set of volume IDs. Therefore, the `emcp.conf` and `powermt.custom` files (which are identical to the R1 files since the boot disk is identical) are modified to create a valid mapping between the emcpower device and native path device for both R1 and R2 locations. Having both the R1 and R2 Symmetrix volume IDs in the `emcp.conf` file ensures a valid mapping between the pseudo devices and the underlying native path device. PowerPath will determine which Symmetrix volume IDs are valid (that is, the visible ones) and will act accordingly when either the R1 or the R2 host is booted.

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.

Native Devices

A *native device* describes a device special file of one of the following forms:

- ◆ Block device — `/dev/dsk/c#t#d#s#`
- ◆ Raw device — `/dev/rdisk/c#t#d#s#`

where:

- ◆ The `c #` is the instance number for the interface card.
- ◆ The `t #` is the target address of the storage system logical device on the bus.
- ◆ The `d #` is the storage system logical device at the target.
- ◆ The `s #` is the slice, ranging from 0 to 7.

Pseudo Devices

A *pseudo device* describes a device special file of one of the following forms:

- ◆ Block device — `/dev/dsk/emcpower#[a-h]`
- ◆ Raw device — `/dev/rdisk/emcpower#[a-h]`

where:

- ◆ `#` is the disk number.
- ◆ `[a-h]` is the slice.

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`.

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.

Native devices are preferable for most installations. Native devices offer the following advantages:

- ◆ If PowerPath is installed, VxVM automatically scans for and recognizes native devices when it (VxVM) is installed. (Pseudo devices must be referenced manually when initializing disks for use with VxVM.)

- ◆ If both PowerPath and VxVM are installed, VxVM automatically scans for and recognizes native devices when volumes are imported. (With pseudo devices, extra manual steps are required to set up disk groups that can be imported.)
- ◆ VxVM provides unqualified support for native names. (VxVM supports pseudo names with the following qualification: pseudo names cause VxVM 3.2 to generate warnings in some circumstances. This is a recognized VERITAS bug, 85455, and EMC has a documented workaround. For details, see <http://seer.support.veritas.com/docs/242612.htm>.)
- ◆ 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.)

Pseudo (emcpower) devices offer the following advantages:

- ◆ There is only one pseudo device name for each multipathed logical device. (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#ã#s#` paths are removed, which can be disruptive to applications using those paths.) Refer to *Dynamic Reconfiguration* on page 6-11 for information on using DR to add and remove HBAs in a PowerPath environment.

Table 6-1 summarizes the functional differences between native devices and emcpower devices in the Solaris environment.

Table 6-1 Native Devices versus emcpower Devices

Function	Native Device	Pseudo Device
I/O failover	✓	✓
I/O load balancing	✓	✓
Bootting: boot-path failover	No	✓
Reboot (reconfiguration)	✓ (Partial support) If a path is missing, PowerPath does not create a "replacement" <code>c#t#d#</code> device.	✓ (Full support)
Support for VxVM sliced disks	✓	No
Support for VxVM simple disks	✓	✓
Support for Solaris disk partitions (slices)	✓	✓
Support for interaction with VxVM DMP (Dynamic MultiPathing) (When using PowerPath 4.3.0, you can manage a CLARiiON® storage system with either PowerPath or DMP, but not both.)	✓	✓
DR transparency	Limitations — <code>c#t#d#s#</code> 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.)
PowerPath's No Redirect load-balancing and failover policy (transparent mode)	Native devices deliver I/O to the path where it would go if PowerPath were not installed. If that path fails, I/Os fail.	Pseudo devices select a configured path for all subsequent I/O. If that path fails, I/O to the pseudo device fails.

Table 6-2 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 6-2 Support for Native and Pseudo Devices

Software Features	Are Native Devices Supported?	Are Pseudo Devices Supported?
Boot Requirements		
Boot-time failover—Symmetrix and CLARiiON	No	✓ <hr/> Failover is not supported during the initial phase of boot; after boot, failover is supported.
Volume Managers		
VxVM 3.1 and 3.2 <ul style="list-style-type: none"> • New installations of PowerPath (DMP always is fully configured and enabled) • Existing installations of PowerPath 	✓ (Preferred) ✓	✓ ✓
VxVM 3.1 and earlier <ul style="list-style-type: none"> • New installations of PowerPath (DMP may be enabled or disabled) • Existing installations of PowerPath 	✓ (Preferred) ✓	✓ ✓
Other volume managers (including Sun's DiskSuite, raw devices, filesystems on raw devices, and raw table spaces) <ul style="list-style-type: none"> • If DR is used • If DR is not used 	✓ ✓	✓ (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:

- ◆ 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

To reconfigure PowerPath devices:

1. Update the `/kernel/drv/sd.conf` file to include target/logical device entries for all multipath storage system logical devices.
2. Create the device nodes. Run the appropriate command:

On hosts running	Run
Solaris 7, 8, and 9	<code>devfsadm -C</code>
Solaris 2.6	<code>drvconfig;disks;devlinks</code>

3. Optionally, run the `powercf -q` command.

PowerPath displays the following message as it creates new devices:

```
Creating new device nodes
```

This step is no longer necessary and is included for backward compatibility only.

4. Run the `powermt config` command.

It is important to run `powermt config` before sending any I/O to the new devices.

5. Run the `format` command and look for `emcpower` devices.
6. Display the new device. Enter:

```
powermt display dev=all
```

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 add and remove HBAs in a PowerPath environment.

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

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.

Adding an HBA to a PowerPath Configuration

To use DR to add an HBA to a Solaris system in a PowerPath configuration, follow these steps:

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

```
powermt config
```

Removing an HBA from a PowerPath Configuration

To use DR to remove an HBA from a Solaris host in a PowerPath configuration, follow these steps:

1. Correlate the `c#t#d#s#` device special files of the HBA being removed with the PowerPath adapter number for that HBA. The PowerPath adapter number is used in the `powermt remove adapter` command later in this procedure.

On 10000 class systems:

- Start the `dr` shell. Enter:

```
dr
```

The prompt changes to `dr>`.

- From within the `dr` shell, list the devices and corresponding `c#t#d#s#` device special files on the I/O board being removed. For example, for I/O board 1, enter:

```
drshow 1
```

The command returns output like the following:

```
I/O Bus Controllers and Devices for Board 1
----- I/O Bus 1 : Slot 0 : esp0 -----

device  opens  name                usage
-----  -
sd838   4        /dev/dsk/c2t0d1s2  /test
...
```

`drshow` displays all device special files that point to HBAs on the I/O board. In the example above, board 1 Slot 0 has a single device (HBA) attached named `c2t0d1s2`.

- Press `CTRL-D` to exit the `dr` shell.
- Associate the device special file (identified above with `drshow`) with a PowerPath adapter number. Enter:

```
powermt display dev=all
```

The command returns output like the following:

```
pseudo name=emcpower0a
Symmetrix ID=000000003269;
logical device ID=048
state=alive; policy=SymmOpt; priority=0; queued-I/Os=0
=====
----- Hosts ----- - Sys - - I/O Paths - --- Stats --
### HW Path          I/O Path Interf.  Mode   State Q-I/Os Errors
=====
0  psbus@b/QLGC,isp@0 c2t0d1  SA 16aB  active alive   0     0
```

In the output, locate device `c2t0d1`. Notice that the adapter number for `c2t0d1` is `0`; therefore, `0` is the adapter number you use as an argument to `powermt remove hba` in step 2.

2. Use `powermt remove hba` to remove the HBA from the PowerPath configuration. Enter:

```
powermt remove hba=#
```

where `#` corresponds to the PowerPath adapter number identified in step 1.

3. Disconnect the HBA, following the instructions in the Sun Dynamic Reconfiguration documentation.

Upgrading Solaris

PowerPath 4.3 supports Solaris Live Upgrade, which lets you upgrade the operating system without uninstalling PowerPath.

Solaris Live Upgrade is supported for upgrades to Solaris 9 only. If you are upgrading to an earlier version of Solaris, you must uninstall PowerPath before you upgrade the OS and then reinstall PowerPath after the upgrade.

Upgrading to Solaris 9 with Solaris Live Upgrade

The PowerPath 4.3 package includes a script, `sol9_liveup.sh`, located in the `scripts` subdirectory of the PowerPath 4.3 installation directory (for example, `/opt/EMCpower/scripts`). You must run this script before you activate Solaris 9.

Solaris Live Upgrade does not work when booting off emcpower devices.

To upgrade to Solaris 9:

1. Ensure that the HBA driver is supported with Solaris 9. Otherwise, Solaris Live Upgrade will not work.
2. Install PowerPath 4.3 on the host.
3. Optionally, save the `sd.conf` file.

Solaris Live Upgrade may replace `sd.conf` with a default `sd.conf` file, causing devices to become inaccessible. If those devices are system critical devices, the new boot image will not work.

Solaris Live Upgrade saves a copy of the `sd.conf` file in `/kernel/drv/sd.conf.~version` (where *version* is the Solaris version from which you are upgrading; for example, if you are upgrading from Solaris 8, the file name is `sd.conf.~8`.) Nevertheless, you might want to save `sd.conf` before you start the upgrade.

4. Install Solaris Live Upgrade on the host.
5. Create an inactive boot environment. Follow the instructions in the *Solaris 9 Installation Guide*.

6. Upgrade the inactive boot environment. Follow the instructions in the *Solaris 9 Installation Guide*.
7. Run the `sol9_liveup.sh` script and then activate the inactive boot environment:
 - a. Mount the root of the inactive boot environment to a directory such as `/mnt`.
 - b. Run the script `sol9_liveup.sh`. Enter:

```
cd /install_dir/EMCpower/scripts
./sol9_liveup.sh /mnt
```

where `install_dir` is the base install location of the EMCpower package. For example, if the base install directory is `opt`, enter:

```
cd /opt/EMCpower/scripts
./sol9_liveup.sh /mnt
```
 - c. Continue with the procedure documented in the *Solaris 9 Installation Guide*.

Troubleshooting

If Solaris Live Upgrade replaces the `sd.conf` file with a default `sd.conf`, the new boot image may not work. In this case, replace the default `sd.conf` with either the file Solaris Live Upgrade automatically saved in `/kernel/drv/sd.conf.~version` or the file you saved before starting the upgrade procedure. (See step 3, above.)

If you fail to run `sol9_liveup.sh` before activating Solaris 9, the new boot image will not work. In this case:

1. Boot from the older Solaris boot area.
2. Mount the Solaris 9 boot area.
3. Run `sol9_liveup.sh` in the Solaris 9 boot area.
4. Boot Solaris 9.

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

Refer to *Reconfiguring PowerPath Devices Online* on page 6-10 for instructions on reconfiguring PowerPath devices on Solaris.

File Location

The `powercf` utility resides in the `/etc` directory.

Executing `powercf`

You must have superuser privileges to use `powercf`.

To run `powercf` on a Solaris host, type the command, plus any options, at the shell prompt.

`emcp.conf` File

The `/kernel/drv/emcp.conf` file lists the primary and alternate path to each storage system logical device and the storage system device serial number for that logical device. The `powercf -q` command updates the existing `emcp.conf` file or creates a new one if it does not already exist.

Syntax

```
powercf -q|-Z
```

Arguments

`powercf` scans HBAs for single-ported and multiported storage system logical devices and compares those logical devices with PowerPath device entries in `emcp.conf`.

`-q`

Runs `powercf` in quiet mode.

`powercf -q` updates the `emcp.conf` file by removing PowerPath devices not found in the HBA scan and adding new PowerPath devices that were found. It saves a primary and an alternate path to each PowerPath device.

`powercf -q` runs automatically during system boot.

`-Z`

Configures an SRDF-enabled server to be bootable from an R2 mirror of a Symmetrix-based `emcpower` boot disk by a remote host.

`powercf -Z` should be run manually whenever such a server's Symmetrix volume configuration changes due to the addition or deletion of volumes.

Error Messages

PowerPath reports any errors, diagnostic messages, and failover recovery messages to the system console and to the file `/var/adm/messages`. Refer to the *PowerPath Product Guide* for a complete list of PowerPath error messages.

**Files Changed By
PowerPath**

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

- ◆ Solaris Files Modified by PowerPath Installation A-2
- ◆ Files Created by PowerPath Installation..... A-6
- ◆ Files Created by VCS Agent Installation A-17

Solaris Files Modified by PowerPath Installation

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

/etc/system

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

System V Semaphore Values

The PowerPath installation sets the semaphore values shown here. You can safely change them, as long as you do not set them to a value lower than the value required by PowerPath. See the individual semaphore descriptions for details.



CAUTION

If after PowerPath is installed, you (or another application) reset a value lower than the value required by PowerPath, PowerPath may not work correctly.

semsys:seminfo_semmns

PowerPath alone needs a value of 257. Therefore, the installation adds at least 257 to the current value in `/etc/system`:

- ◆ If the current value in `/etc/system` is greater than 60, the PowerPath installation adds 257.
- ◆ If the current value in `/etc/system` is less than or equal to 60, the installation adds 317 to allow room for expansion.
- ◆ If the value of `semsys:seminfo_semmns` is not set in `/etc/system`, the PowerPath installation sets the value at 317.

semsys:seminfo_semmni

PowerPath alone needs a value of 2. Therefore, the installation adds at least 2 to the current value in `/etc/system`:

- ◆ If the current value in `/etc/system` is greater than 10, the PowerPath installation adds 2.
- ◆ If the current value in `/etc/system` and is less than or equal to 10, the installation adds 12.
- ◆ If the value of `semsys:seminfo_semmni` is not set in `/etc/system`, the PowerPath installation sets the value at 12.

semsys:seminfo_semmsl

The maximum value allowed for this semaphore is 129. The PowerPath installation changes this value as follows:

- ◆ If the current value in `/etc/system` is 129, it is not reset.
- ◆ If the current value in `/etc/system` is less than 129, it is set to 129.
- ◆ If the value of `semsys:seminfo_semmsl` is not set in `/etc/system`, it is set to 129.

semsys:seminfo_semmap

`semsys:seminfo_semmap` is set to the largest value of `semsys:seminfo_semmns` found in `/etc/system`, or to 317 if `semsys:seminfo_semmns` is not in `/etc/system`.

semsys:seminfo_semmnu

`semsys:seminfo_semmnu` is set to the largest value of `semsys:seminfo_semmns` found in `/etc/system`, or to 317 if `semsys:seminfo_semmns` is not in `/etc/system`.

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/emcp
forceload: misc/emcpmp
forceload: misc/emcpmpc
forceload: misc/emcpmpaa
forceload: misc/emcpsapi
forceload: misc/emcpcg
forceload: misc/emcphr
forceload: misc/emcpsm
forceload: misc/emcpsc
forceload: misc/emcpevm
forceload: misc/emcpdpm
forceload: misc/emcpioc
```

PowerPath also adds these forceload statements if they are not already present:

```
forceload: drv/sd
forceload: drv/ssd
```

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 emcp:bPxEnableInit=1
set lwp_default_stksize=0x6000
set rpcmod:svc_run_stksize=0x6000
```

Other (Solaris 5.6 Only)

On hosts running Solaris 5.6, PowerPath sets the following parameters in `/etc/system`:

- ◆ `maxusers` to `0x800`
- ◆ `kobj_map_space_len` to `0x400000`

/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/emcpv_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 /opt/EMCpower/scripts/emcpv_setup.csh
```

/etc/syslog.conf

PowerPath adds to `/etc/syslog.conf` lines of the following form:

```
localn.info /var/adm/messages
```

where n is a digit 0-7.

The number of such lines added by PowerPath depends upon whether lines of this form already exist in the file. If this line exists:

```
localm.info /var/adm/messages
```

PowerPath adds entries for all digits less than m .

/etc/driver_classes

/etc/name_to_major

Files Created by PowerPath Installation

The following files are created when PowerPath is installed on a Solaris host. Symbolic links are indicated by ->.

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

/etc

This directory contains PowerPath CLI commands and utilities:

- ◆ S87powervxvm
- ◆ cgmt
- ◆ emc/bin/emcp_discover
- ◆ emc/bin/emcp_ini
- ◆ emc/bin/emcp_purge
- ◆ emc/bin/emcpdiscover
- ◆ emc/bin/emcpmgr
- ◆ emc/bin/emcppurge
- ◆ emc/bin/inquiry.pp
- ◆ emc/mpaa.excluded
- ◆ emc/mpaa.lams
- ◆ emcpmgr
- ◆ emcpower_mode_dir
- ◆ emcpreg
- ◆ emcpupgrade
- ◆ emcpv_cleanup
- ◆ init.d/emcpvStartup
- ◆ powercf
- ◆ powermt
- ◆ powerprotect
- ◆ powervshm
- ◆ powervsvs
- ◆ powervxvm
- ◆ rc2.d/S02configcgs
- ◆ rc3.d/S98emcpv -> ../init.d/emcpvStartup
- ◆ rcS.d/S24powerstartup
- ◆ rcS.d/S63powershift
- ◆ K1lemcpv -> ../init.d/emcpvStartup
- ◆ emcpcvt -> /etc/emcpupgrade
- ◆ powermt.custom
- ◆ emcp_registration

- ◆ emcp_devicesDB.dat
- ◆ emcp_devicesDB.idx
- ◆ PPVM_config
- ◆ PPVM_config_bak

/usr/lib

This directory contains symbolic links to installed PowerPath libraries:

- ◆ libcg.so -> libemcpcg.so
- ◆ libemcppn.so -> libpn.so
- ◆ libdm.so -> ../../opt/EMCpower/lib/libdm_32.so
- ◆ libdm_32.so -> ../../opt/EMCpower/lib/libdm_32.so
- ◆ libemcp_32.so -> libemcp.so
- ◆ libemcp_lic_rtl_32.so -> libemcp_lic_rtl.so
- ◆ libemcp_mp_rtl_32.so -> libemcp_mp_rtl.so
- ◆ libemcpmt_32.so -> libemcpmt.so
- ◆ libemcpvapi.so ->
 - ../../opt/EMCpower/lib/libemcpvapi_32.so
- ◆ libemcpvapi_32.so ->
 - ../../opt/EMCpower/lib/libemcpvapi_32.so
- ◆ libemcpvapimt.so ->
 - ../../opt/EMCpower/lib/libemcpvapimt_32.so
- ◆ libemcpvapimt_32.so ->
 - ../../opt/EMCpower/lib/libemcpvapimt_32.so
- ◆ libmp_32.so -> libemcnpmp.so
- ◆ libom.so -> ../../opt/EMCpower/lib/libom_32.so
- ◆ libom_32.so -> ../../opt/EMCpower/lib/libom_32.so
- ◆ libommt.so -> ../../opt/EMCpower/lib/libommt_32.so
- ◆ libommt_32.so -> ../../opt/EMCpower/lib/libommt_32.so
- ◆ libpn_32.so -> libpn.so
- ◆ libradcommon.so ->
 - ../../opt/EMCpower/lib/libradcommon_32.so
- ◆ libradcommon_32.so ->
 - ../../opt/EMCpower/lib/libradcommon_32.so
- ◆ libradcommonmt.so ->
 - ../../opt/EMCpower/lib/libradcommonmt_32.so
- ◆ libradcommonmt_32.so ->
 - ../../opt/EMCpower/lib/libradcommonmt_32.so
- ◆ libraddebug.so ->
 - ../../opt/EMCpower/lib/libraddebug_32.so
- ◆ libraddebug_32.so ->
 - ../../opt/EMCpower/lib/libraddebug_32.so
- ◆ libraddebugmt.so ->
 - ../../opt/EMCpower/lib/libraddebugmt_32.so

- ◆ libraddebugmt_32.so ->
../opt/EMCpower/lib/libraddebugmt_32.so
- ◆ libsc.so -> ../opt/EMCpower/lib/libsc_32.so
- ◆ libsc_32.so -> ../opt/EMCpower/lib/libsc_32.so
- ◆ libscmt.so -> ../opt/EMCpower/lib/libscmt_32.so
- ◆ libscmt_32.so -> ../opt/EMCpower/lib/libscmt_32.so
- ◆ libsm.so -> ../opt/EMCpower/lib/libsm_32.so
- ◆ libsm_32.so -> ../opt/EMCpower/lib/libsm_32.so
- ◆ libsmmt.so -> ../opt/EMCpower/lib/libsmmt_32.so
- ◆ libsmmt_32.so -> ../opt/EMCpower/lib/libsmmt_32.so
- ◆ libdpmh.so -> ../opt/EMCpower/lib/libdpmh_32.so
- ◆ libdpmh_32.so -> ../opt/EMCpower/lib/libdpmh_32.so
- ◆ libdpmhmt.so ->
../opt/EMCpower/lib/libdpmhmt_32.so
- ◆ libdpmhmt_32.so ->
../opt/EMCpower/lib/libdpmhmt_32.so
- ◆ libdpmu.so -> ../opt/EMCpower/lib/libdpmu_32.so
- ◆ libdpmu_32.so -> ../opt/EMCpower/lib/libdpmu_32.so
- ◆ libdpmumt.so ->
../opt/EMCpower/lib/libdpmumt_32.so
- ◆ libdpmumt_32.so ->
../opt/EMCpower/lib/libdpmumt_32.so
- ◆ libevm.so -> ../opt/EMCpower/lib/libevm_32.so
- ◆ libevm_32.so -> ../opt/EMCpower/lib/libevm_32.so
- ◆ libevmmt.so -> ../opt/EMCpower/lib/libevmmt_32.so
- ◆ libevmmt_32.so ->
../opt/EMCpower/lib/libevmmt_32.so
- ◆ libgm.so -> ../opt/EMCpower/lib/libgm_32.so
- ◆ libgm_32.so -> ../opt/EMCpower/lib/libgm_32.so
- ◆ libgmmt.so -> ../opt/EMCpower/lib/libgmmt_32.so
- ◆ libgmmt_32.so -> ../opt/EMCpower/lib/libgmmt_32.so
- ◆ sparcv9/libcpg.so -> libemcpg.so
- ◆ sparcv9/libemcppn.so -> libpn.so
- ◆ sparcv9/libemcpmt.so
- ◆ sparcv9/libemcp_lic_rtl.so
- ◆ sparcv9/libemcp.so
- ◆ sparcv9/libemcp_mp_rtl.so
- ◆ sparcv9/libemcpg.so
- ◆ sparcv9/libemcpmp.so
- ◆ sparcv9/libpn.so

/basedir/EMCpower/bin

This directory contains 32- and 64-bit versions of the PowerPath Volume Manager CLI commands:

- ◆ 64 -> ./sparcv9
- ◆ emcpminor -> ./emcpminor_32
- ◆ emcpminor_32
- ◆ powerdm -> ./powerdm_32
- ◆ powerdm_32
- ◆ powervadm -> ./powervadm_32
- ◆ powervadm_32
- ◆ powervg -> ./powervg_32
- ◆ powervg_32
- ◆ powervmeta -> ./powervmeta_32
- ◆ powervmeta_32
- ◆ powervol -> ./powervol_32
- ◆ powervol_32
- ◆ powervperf -> ./powervperf_32
- ◆ powervperf_32
- ◆
- ◆ sparcv9/emcpminor -> ./emcpminor_64
- ◆ sparcv9/emcpminor_64
- ◆ sparcv9/powerdm -> ./powerdm_64
- ◆ sparcv9/powerdm_64
- ◆ sparcv9/powervadm -> ./powervadm_64
- ◆ sparcv9/powervadm_64
- ◆ sparcv9/powervg -> ./powervg_64
- ◆ sparcv9/powervg_64
- ◆ sparcv9/powervmeta -> ./powervmeta_64
- ◆ sparcv9/powervmeta_64
- ◆ sparcv9/powervol -> ./powervol_64
- ◆ sparcv9/powervol_64
- ◆ sparcv9/powervperf -> ./powervperf_64
- ◆ sparcv9/powervperf_64

/basedir/EMCpower/driver

This directory contains 32- and 64-bit versions of the PowerPath kernel files:

- ◆ 64 -> ./sparcv9
- ◆ emcp -> ./emcp_32
- ◆ emcp_32

- ◆ emcp_32_5_9
- ◆ emcp_5_9 -> ./emcp_32_5_9
- ◆ emcpcg -> ./emcpcg_32
- ◆ emcpcg_32
- ◆ emcpcg_32_5_9
- ◆ emcpcg_5_9 -> ./emcpcg_32_5_9
- ◆ emcpdpm -> ./emcpdpm_32
- ◆ emcpdpm_32
- ◆ emcpdpm_32_5_9
- ◆ emcpdpm_5_9 -> ./emcpdpm_32_5_9
- ◆ emcpevm -> ./emcpevm_32
- ◆ emcpevm_32
- ◆ emcpevm_32_5_9
- ◆ emcpevm_5_9 -> ./emcpevm_32_5_9
- ◆ emcphr -> ./emcphr_32
- ◆ emcphr_32
- ◆ emcphr_32_5_9
- ◆ emcphr_5_9 -> ./emcphr_32_5_9
- ◆ emcpioc -> ./emcpioc_32
- ◆ emcpioc_32
- ◆ emcpioc_32_5_9
- ◆ emcpioc_5_9 -> ./emcpioc_32_5_9
- ◆ emcpmp -> ./emcpmp_32
- ◆ emcpmp_32
- ◆ emcpmp_32_5_9
- ◆ emcpmp_5_9 -> ./emcpmp_32_5_9
- ◆ emcpmpaa -> ./emcpmpaa_32
- ◆ emcpmpaa_32
- ◆ emcpmpaa_32_5_9
- ◆ emcpmpaa_5_9 -> ./emcpmpaa_32_5_9
- ◆ emcpmpc -> ./emcpmpc_32
- ◆ emcpmpc_32
- ◆ emcpmpc_32_5_9
- ◆ emcpmpc_5_9 -> ./emcpmpc_32_5_9
- ◆ emcpsapi -> ./emcpsapi_32
- ◆ emcpsapi_32
- ◆ emcpsapi_32_5_9
- ◆ emcpsapi_5_9 -> ./emcpsapi_32_5_9
- ◆ emcpssc -> ./emcpssc_32
- ◆ emcpssc_32
- ◆ emcpssc_32_5_9
- ◆ emcpssc_5_9 -> ./emcpssc_32_5_9
- ◆ emcpsm -> ./emcpsm_32
- ◆ emcpsm_32
- ◆ emcpsm_32_5_9
- ◆ emcpsm_5_9 -> ./emcpsm_32_5_9
- ◆ emcpte -> ./emcpte_32

- ◆ emcpte_32
- ◆ emcpte_32_5_9
- ◆ emcpte_5_9 -> ./emcpte_32_5_9
- ◆ sparcv9
- ◆ sparcv9/emcp -> ./emcp_64
- ◆ sparcv9/emcp_5_9 -> ./emcp_64_5_9
- ◆ sparcv9/emcp_64
- ◆ sparcv9/emcp_64_5_9
- ◆ sparcv9/emcpcg -> ./emcpcg_64
- ◆ sparcv9/emcpcg_5_9 -> ./emcpcg_64_5_9
- ◆ sparcv9/emcpcg_64
- ◆ sparcv9/emcpcg_64_5_9
- ◆ sparcv9/emcpdpm -> ./emcpdpm_64
- ◆ sparcv9/emcpdpm_5_9 -> ./emcpdpm_64_5_9
- ◆ sparcv9/emcpdpm_64
- ◆ sparcv9/emcpdpm_64_5_9
- ◆ sparcv9/emcpevm -> ./emcpevm_64
- ◆ sparcv9/emcpevm_5_9 -> ./emcpevm_64_5_9
- ◆ sparcv9/emcpevm_64
- ◆ sparcv9/emcpevm_64_5_9
- ◆ sparcv9/emcphr -> ./emcphr_64
- ◆ sparcv9/emcphr_5_9 -> ./emcphr_64_5_9
- ◆ sparcv9/emcphr_64
- ◆ sparcv9/emcphr_64_5_9
- ◆ sparcv9/emcpioc -> ./emcpioc_64
- ◆ sparcv9/emcpioc_5_9 -> ./emcpioc_64_5_9
- ◆ sparcv9/emcpioc_64
- ◆ sparcv9/emcpioc_64_5_9
- ◆ sparcv9/emcpmp -> ./emcpmp_64
- ◆ sparcv9/emcpmp_5_9 -> ./emcpmp_64_5_9
- ◆ sparcv9/emcpmp_64
- ◆ sparcv9/emcpmp_64_5_9
- ◆ sparcv9/emcpmpaa -> ./emcpmpaa_64
- ◆ sparcv9/emcpmpaa_5_9 -> ./emcpmpaa_64_5_9
- ◆ sparcv9/emcpmpaa_64
- ◆ sparcv9/emcpmpaa_64_5_9
- ◆ sparcv9/emcpmpc -> ./emcpmpc_64
- ◆ sparcv9/emcpmpc_5_9 -> ./emcpmpc_64_5_9
- ◆ sparcv9/emcpmpc_64
- ◆ sparcv9/emcpmpc_64_5_9
- ◆ sparcv9/emcpsapi -> ./emcpsapi_64
- ◆ sparcv9/emcpsapi_5_9 -> ./emcpsapi_64_5_9
- ◆ sparcv9/emcpsapi_64
- ◆ sparcv9/emcpsapi_64_5_9
- ◆ sparcv9/emcpsc -> ./emcpsc_64
- ◆ sparcv9/emcpsc_5_9 -> ./emcpsc_64_5_9
- ◆ sparcv9/emcpsc_64

- ◆ sparcv9/emcpsc_64_5_9
- ◆ sparcv9/emcpsm -> ./emcpsm_64
- ◆ sparcv9/emcpsm_5_9 -> ./emcpsm_64_5_9
- ◆ sparcv9/emcpsm_64
- ◆ sparcv9/emcpsm_64_5_9
- ◆ sparcv9/emcpte -> ./emcpte_64
- ◆ sparcv9/emcpte_5_9 -> ./emcpte_64_5_9
- ◆ sparcv9/emcpte_64
- ◆ sparcv9/emcpte_64_5_9

/kernel/drv

This directory contains the PowerPath driver:

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

kernel/misc

This directory contains the PowerPath driver extensions:

- ◆ emcpdpm
- ◆ emcpsm
- ◆ emcpsc
- ◆ emcpevm
- ◆ emcploc
- ◆ emcpsapi
- ◆ emcphr
- ◆ emcpmpc
- ◆ emcpcg
- ◆ emcmpaa
- ◆ emcpmp
- ◆ sparcv9/emcpdpm
- ◆ sparcv9/emcpsm
- ◆ sparcv9/emcpsc
- ◆ sparcv9/emcpevm
- ◆ sparcv9/emcploc
- ◆ sparcv9/emcpsapi
- ◆ sparcv9/emcphr
- ◆ sparcv9/emcpmpc
- ◆ sparcv9/emcpcg
- ◆ sparcv9/emcmpaa
- ◆ sparcv9/emcpmp

/basedir/EMCpower/lib

This directory contains 32- and 64-bit versions of the PowerPath libraries:

- ◆ 64 -> ./sparcv9
- ◆ libdm.so -> ./libdm_32.so
- ◆ libdm_32.so
- ◆ libdpmh.so -> ./libdpmh_32.so
- ◆ libdpmh_32.so
- ◆ libdpmhmt.so -> ./libdpmhmt_32.so
- ◆ libdpmhmt_32.so
- ◆ libdpmu.so -> ./libdpmu_32.so
- ◆ libdpmu_32.so
- ◆ libdpmumt.so -> ./libdpmumt_32.so
- ◆ libdpmumt_32.so
- ◆ libemcp.so -> ./libemcp_32.so
- ◆ libemcp_32.so
- ◆ libemcp_lam.so -> ./libemcp_lam_32.so
- ◆ libemcp_lam_32.so
- ◆ libemcp_lic_rtl.so -> ./libemcp_lic_rtl_32.so
- ◆ libemcp_lic_rtl_32.so
- ◆ libemcp_mp_rtl.so -> ./libemcp_mp_rtl_32.so
- ◆ libemcp_mp_rtl_32.so
- ◆ libemcpcg.so -> ./libemcpcg_32.so
- ◆ libemcpcg_32.so
- ◆ libemcpmp.so -> ./libemcpmp_32.so
- ◆ libemcpmp_32.so
- ◆ libemcpmt.so -> ./libemcpmt_32.so
- ◆ libemcpmt_32.so
- ◆ libemcpvapi.so -> ./libemcpvapi_32.so
- ◆ libemcpvapi_32.so
- ◆ libemcpvapimt.so -> ./libemcpvapimt_32.so
- ◆ libemcpvapimt_32.so
- ◆ libevm.so -> ./libevm_32.so
- ◆ libevm_32.so
- ◆ libevmmt.so -> ./libevmmt_32.so
- ◆ libevmmt_32.so
- ◆ libgm.so -> ./libgm_32.so
- ◆ libgm_32.so
- ◆ libgmmt.so -> ./libgmmt_32.so
- ◆ libgmmt_32.so
- ◆ libmp_32.so -> ./libemcpmp_32.so
- ◆ libom.so -> ./libom_32.so
- ◆ libom_32.so
- ◆ libommt.so -> ./libommt_32.so
- ◆ libommt_32.so

- ◆ libpn.so -> ./libpn_32.so
- ◆ libpn_32.so
- ◆ libradcommon.so -> ./libradcommon_32.so
- ◆ libradcommon_32.so
- ◆ libradcommonmt.so -> ./libradcommonmt_32.so
- ◆ libradcommonmt_32.so
- ◆ libraddebug.so -> ./libraddebug_32.so
- ◆ libraddebug_32.so
- ◆ libraddebugmt.so -> ./libraddebugmt_32.so
- ◆ libraddebugmt_32.so
- ◆ libsc.so -> ./libsc_32.so
- ◆ libsc_32.so
- ◆ libscmt.so -> ./libscmt_32.so
- ◆ libscmt_32.so
- ◆ libsm.so -> ./libsm_32.so
- ◆ libsm_32.so
- ◆ libsmmt.so -> ./libsmmt_32.so
- ◆ libsmmt_32.so
- ◆ sparcv9/libdm.so -> ./libdm_64.so
- ◆ sparcv9/libdm_64.so
- ◆ sparcv9/libdpmh.so -> ./libdpmh_64.so
- ◆ sparcv9/libdpmh_64.so
- ◆ sparcv9/libdpmhmt.so -> ./libdpmhmt_64.so
- ◆ sparcv9/libdpmhmt_64.so
- ◆ sparcv9/libdpmu.so -> ./libdpmu_64.so
- ◆ sparcv9/libdpmu_64.so
- ◆ sparcv9/libdpmumt.so -> ./libdpmumt_64.so
- ◆ sparcv9/libdpmumt_64.so
- ◆ sparcv9/libemcp.so -> ./libemcp_64.so
- ◆ sparcv9/libemcp_64.so
- ◆ sparcv9/libemcp_lam.so -> ./libemcp_lam_64.so
- ◆ sparcv9/libemcp_lam_64.so
- ◆ sparcv9/libemcp_lic_rtl.so -> ./libemcp_lic_rtl_64.so
- ◆ sparcv9/libemcp_lic_rtl_64.so
- ◆ sparcv9/libemcp_mp_rtl.so -> ./libemcp_mp_rtl_64.so
- ◆ sparcv9/libemcp_mp_rtl_64.so
- ◆ sparcv9/libemcpcg.so -> ./libemcpcg_64.so
- ◆ sparcv9/libemcpcg_64.so
- ◆ sparcv9/libemcpmp.so -> ./libemcpmp_64.so
- ◆ sparcv9/libemcpmp_64.so
- ◆ sparcv9/libemcpmt.so -> ./libemcpmt_64.so
- ◆ sparcv9/libemcpmt_64.so
- ◆ sparcv9/libemcpvapi.so -> ./libemcpvapi_64.so
- ◆ sparcv9/libemcpvapi_64.so
- ◆ sparcv9/libemcpvapimt.so -> ./libemcpvapimt_64.so
- ◆ sparcv9/libemcpvapimt_64.so
- ◆ sparcv9/libevm.so -> ./libevm_64.so

- ◆ sparcv9/libevm_64.so
- ◆ sparcv9/libevmmt.so -> ./libevmmt_64.so
- ◆ sparcv9/libevmmt_64.so
- ◆ sparcv9/libgm.so -> ./libgm_64.so
- ◆ sparcv9/libgm_64.so
- ◆ sparcv9/libgmmt.so -> ./libgmmt_64.so
- ◆ sparcv9/libgmmt_64.so
- ◆ sparcv9/libmp_64.so -> ./libemcpmp_64.so
- ◆ sparcv9/libom.so -> ./libom_64.so
- ◆ sparcv9/libom_64.so
- ◆ sparcv9/libommt.so -> ./libommt_64.so
- ◆ sparcv9/libommt_64.so
- ◆ sparcv9/libpn.so -> ./libpn_64.so
- ◆ sparcv9/libpn_64.so
- ◆ sparcv9/libradcommon.so -> ./libradcommon_64.so
- ◆ sparcv9/libradcommon_64.so
- ◆ sparcv9/libradcommonmt.so -> ./libradcommonmt_64.so
- ◆ sparcv9/libradcommonmt_64.so
- ◆ sparcv9/libraddebug.so -> ./libraddebug_64.so
- ◆ sparcv9/libraddebug_64.so
- ◆ sparcv9/libraddebugmt.so -> ./libraddebugmt_64.so
- ◆ sparcv9/libraddebugmt_64.so
- ◆ sparcv9/libsc.so -> ./libsc_64.so
- ◆ sparcv9/libsc_64.so
- ◆ sparcv9/libscmt.so -> ./libscmt_64.so
- ◆ sparcv9/libscmt_64.so
- ◆ sparcv9/libsm.so -> ./libsm_64.so
- ◆ sparcv9/libsm_64.so
- ◆ sparcv9/libsmmt.so -> ./libsmmt_64.so
- ◆ sparcv9/libsmmt_64.so

/basedir/EMCpower/man/man1

This directory contains the PowerPath Volume Manager and Volume Mobility man pages:

- ◆ powerdm.1
- ◆ powervadm.1
- ◆ powervg.1
- ◆ powervintro.1
- ◆ powervmeta.1
- ◆ powervol.1
- ◆ powervperf.1

/basedir/EMCpower/scripts

This directory contains PowerPath scripts:

- ◆ emcpv_cleanup
- ◆ emcpv_cron_remove.sh
- ◆ emcpv_cron_setup.sh
- ◆ emcpv_logchecker
- ◆ emcpv_setup.csh
- ◆ emcpv_setup.sh
- ◆ emcpv_slv2devinfo
- ◆ sol9_liveup.sh

/usr/bin

This directory contains `powervini`, the PowerPath Volume Manager initialization file.

/usr/man/man1

This directory contains the PowerPath man pages:

- ◆ emcpminor.1
- ◆ emcpreg.1
- ◆ emcpupgrade.1
- ◆ powercf.1
- ◆ powermt.1

Files Created by VCS Agent Installation

The following files are created when the PowerPath Volume Manager VCS Agent is installed on a Solaris host.

/basedir/VRTSvcs/EMC

- ◆ bin/EMCvgAgent
- ◆ lib/libEMCvg.so
- ◆ messages.EMC
- ◆ scripts/_abortOnline
- ◆ scripts/_getResAttr
- ◆ scripts/_getResAttrLocal
- ◆ scripts/_getTypeAttr
- ◆ scripts/_isResExist
- ◆ scripts/_setChldResAttr
- ◆ scripts/_setResAttr
- ◆ scripts/_setTypeAttr
- ◆ types/EMCVgTypes.cf

/basedir/VRTSvcs/bin/EMCvg

- ◆ EMCvgAgent (symbolic link to
/base_directory/VRTSvcs/EMC/bin/EMCvgAgent)

/etc/VRTSvcs/conf/config

- ◆ EMCTypes.cf

/usr/lib

- ◆ libEMCvg.so (symbolic link to
/base_directory/VRTSvcs/EMC/lib/libEMCvg.so)

Upgrading from Early Versions of PowerPath

This appendix describes what to do if you are upgrading to PowerPath 4.3 from PowerPath 2.1.x, 2.0, or 1.5.

- ◆ Upgrading from PowerPath 2.1.x..... B-2
- ◆ Upgrading from PowerPath 2.0..... B-2
- ◆ Upgrading from PowerPath 1.5..... B-3

Upgrading from PowerPath 2.1.x

You can upgrade directly to PowerPath 4.3 from PowerPath 2.1.0 and 2.1.2. Refer to *Upgrading to PowerPath 4.3* on page 1-15.

To upgrade from any of the following releases, however, you must first upgrade to PowerPath 3.0.x and then upgrade to PowerPath 4.3:

- ◆ PowerPath 2.1.1
- ◆ PowerPath 2.1.3
- ◆ PowerPath 2.1.4
- ◆ PowerPath 2.1.5
- ◆ PowerPath 2.1.6

The general upgrade procedure for these releases is as follows:

1. Search the file `/etc/profile` for the following string:

```
end of additions
```

If you find that string, replace it with the following string:

```
done additions
```

2. Uninstall PowerPath 2.1.x, following the instructions in the PowerPath 2.1 documentation *with this exception: Do not reboot the host after you uninstall PowerPath 2.1.x.*
3. Install PowerPath 3.0.x, following the instructions in the PowerPath 3.0.x documenting.
4. Uninstall PowerPath 3.0.x, following the instructions in the PowerPath 3.0 documentation *with this exception: Do not reboot the host after you uninstall PowerPath 3.0.x.*
5. Install PowerPath 4.3, following the instructions in this manual.

Upgrading from PowerPath 2.0

To upgrade from PowerPath 2.0:

1. Uninstall PowerPath 2.0, following the instructions in the PowerPath 2.0 documentation *with this exception: Do not reboot the host after you uninstall PowerPath 2.0.*
2. Install PowerPath 2.1.0 or PowerPath 3.0.x, following the instructions in the PowerPath 2.1 or 3.0 documentation.

3. Restore your configuration. Enter the following command:

```
powermt load file=/etc/powermt.custom.pre-pp2.1.0
```
4. Uninstall PowerPath 2.1.0 or PowerPath 3.0.x, following the instructions in the PowerPath 2.1 or 3.0 documentation with this exception: *Do not reboot the host after you uninstall PowerPath 2.1.0 or 3.0.x.*
5. Install PowerPath 4.3, following the instructions in this manual.

Upgrading from PowerPath 1.5

Upgrading from PowerPath 1.5 to PowerPath 4.3 requires multiple intermediate steps. Contact your Customer Support representative for information, citing EMC Knowledge Base Article emc75763.

Symbols

/etc/.login A-5
/etc/.login file 1-27
/etc/profile A-5
/etc/profile file 1-27
/etc/syslog.conf A-5
/etc/system A-2

A

AAM. *See* LEGATO AAM
Adding devices online 6-10

B

Boot device, removing PowerPath control 3-11

C

CLARiiON storage systems
 setting up 1-2
Clusters. *See* LEGATO AAM, Sun Cluster 2.2, Sun
 Cluster 3.0, VCS
Configuration, restoring 6-4
Configuring PowerPath devices using powercf
 6-16

D

Device naming conventions 6-5
Devices, adding online 6-10
DR 6-11
Dynamic Reconfiguration 6-11

E

emcp.conf file 6-16
emcpower devices 6-6
 and VERITAS Volume Manager 6-7
 configuring applications for 1-29
Error messages 1-32, 6-17
Error messages, ODS device overlap 6-2

F

Forceload statements set by PowerPath A-4
format utility, Solaris 1-4

H

Host bus adapter
 adding 6-11
 removing 6-12

I

Installing PowerPath
 files created or modified A-6
 installation procedure 1-9
 postinstallation checklist 1-23
 preinstallation checklist 1-2
 upgrading 1-15
Installing PowerPath Volume Manager VCS
 Agent 1-30, A-17

K

Kernel parameters set by PowerPath A-2
Kernel stack size settings A-4

L

LEGATO AAM
 installing PowerPath into 2-2
 integrating PowerPath into 2-3
 PowerPath Volume Manager LEGATO
 Agent 1-1

M

Major number 1-7, 1-11, 1-18, 2-11
 Minor number 2-11

N

Naming conventions 6-6
 Native devices 6-6

P

powercf configuration utility 6-16
 PowerPath boot device
 configuring 3-2
 troubleshooting 3-9
 PowerPath devices
 emcpower 6-6
 native 6-6
 reconfiguring online 6-10
 PowerPath error messages 6-17
 PowerPath Volume Manager LEGATO Agent 1-1
 PowerPath Volume Manager VCS Agent
 configuring VCS for 2-9
 files A-17
 installing 1-30
 uninstalling 5-7
 Pseudo devices 6-6

R

R1/R2 boot failover support 6-4
 Reconfiguring PowerPath devices online 6-10
 Removing PowerPath Volume Manager VCS
 Agent 5-7

S

semsys
 seminfo_semmap A-3
 seminfo_semmni A-3
 seminfo_semmns A-2

seminfo_semmnu A-3
 seminfo_semmsl A-3
 Solaris format utility 1-4
 Solaris Live Upgrade 6-14
 Solaris OS, upgrading 6-14
 Stack size 6-2
 Sun Cluster 3.0
 installing with PowerPath 2-4
 integrating PowerPath into 2-5
 System V semaphore values A-2

T

Troubleshooting
 PowerPath upgrade 1-21
 Solaris Live Upgrade 6-15

U

Uninstalling PowerPath Volume Manager VCS
 Agent 5-7
 Upgrading PowerPath 1-15
 Upgrading Solaris 6-14

V

VCS
 agent, PowerPath Volume Manager. *See*
 PowerPath Volume Manager VCS
 Agent
 installing with PowerPath 2-6
 integrating PowerPath into 2-8
 VERITAS Cluster Server. *See* VCS
 VERITAS Volume Manager and emcpower
 devices 6-7

W

Warning messages, displaying 1-5