
EMC Symmetrix, CLARiiON, and IBM AIX Fibre Boot

Feature/Functionality, Requirements, Limitations

Abstract

This white paper contains technical information and instruction on configuring and troubleshooting an AIX host when booting from the Symmetrix storage array over a fibre connection.

Published 4/29/2005

Copyright © 2003 - 2005 EMC Corporation. All rights reserved.

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.

EMC², EMC, EMC ControlCenter, ApplicationXtender, Celerra, CentraStar, CLARalert, CLARiiON, Connectrix, Dantz, Direct Matrix Architecture, DiskXtender, Documentum, EmailXtender, EmailXtract, HighRoad, Legato, Navisphere, OpenScale, PowerPath, RepliStor, ResourcePak, Retrospect, Smarts, SnapView/IP, SRDF, Symmetrix, TimeFinder, VisualSAN, and where information lives are registered trademarks and EMC Developers Program, EMC OnCourse, EMC Proven, EMC Snap, EMC Storage Administrator, Access Logix, ArchiveXtender, Authentic Problems, Automated Resource Manager, AutoSwap, AVALONidm, C-Clip, Celerra Replicator, Centera, CLARevent, Codebook Correlation Technology, Common Information Model, CopyCross, CopyPoint, DatabaseXtender, Direct Matrix, DiskXtender 2000, EDM, E-Lab, EmailXaminer, Enginuity, eRoom, FarPoint, FLARE, FullTime, Graphic Visualization, InfoMover, Invista, MirrorView, NetWin, NetWorker, OnAlert, Powerlink, RepliCare, SafeLine, SAN Advisor, SAN Copy, SAN Manager, SDMS, SnapSure, SnapView, StorageScope, SupportMate, SymmAPI, SymmEnabler, Symmetrix DMX, Viewlets, and VisualSRM are trademarks of EMC Corporation. All other trademarks used herein are the property of their respective owners.

Part Number H572.2

Table of Contents

| | |
|---|-----------|
| Introduction | 5 |
| Operating System Requirements | 5 |
| Symmetrix Models Supported | 5 |
| CLARiiON Models Supported | 5 |
| Symmetrix Director Bit Settings | 5 |
| Minimum Enginuity Storage Operating Environment Version for Each Family | 5 |
| Minimum FLARE Base Operating Environment Versions | 6 |
| Supported HBAs | 6 |
| Supported System Firmware | 7 |
| Supported EMC Switch Models | 7 |
| Installation | 7 |
| New Installation Overview | 7 |
| Cloned Installation Overview | 8 |
| Preparing Your System for Booting from the Symmetrix or CLARiiON Storage Array | 8 |
| Prerequisites | 8 |
| Recommendations | 9 |
| Adding the Adapters to Your Host | 9 |
| Host with Existing Adapters | 10 |
| Updating the System/Service Processor Combined Microcode | 12 |
| Prerequisites | 12 |
| Connecting to the CLARiiON Array | 12 |
| Prerequisites | 12 |
| Adding Physical Connections | 13 |
| Completing Your Configuration | 13 |
| Installing to the Symmetrix or CLARiiON Disk | 14 |
| Prerequisites | 14 |
| Recommendations | 15 |
| New Installation | 15 |
| Migration Installation | 16 |
| Cloned Installation Using alt_disk_install | 17 |
| Installing the EMC ODM Support Package | 18 |
| Prerequisites | 18 |
| ODM Download and Installation for Symmetrix and CLARiiON Release 13 or Higher | 18 |
| Pre-Release 13: Installing CLArrayS3 from the EMC Utilities Kit | 19 |
| Prerequisites | 19 |
| CLArrayS3 Installation | 19 |
| Special Considerations | 20 |
| Volume Logix | 20 |
| Prerequisites for Volume Logix | 20 |

| | |
|---|----|
| Recommendations for Volume Logix | 20 |
| PowerPath..... | 22 |
| Considerations for CLARiiON Storage | 23 |
| The pprootdev Tool | 23 |
| Trespasses | 23 |
| Configuring a New PowerPath Installation | 24 |
| Configuring an Existing PowerPath Installation | 24 |
| The lsvg Command | 26 |
| Actions That Use bosboot..... | 26 |
| Special Recovery Procedures: Path Failure during Bootup | 26 |
| rootvg Resides in the Symmetrix Storage Array But Is Mirrored across PowerPath Adapters . | 27 |
| Recommendations | 27 |
| Recommendations | 28 |
| Special Recovery Procedures: Path Failure | 30 |
| Using MKSYSB for Backup and Restore | 30 |
| Prerequisites..... | 30 |
| SRDF | 31 |
| SRDF Installation Overview | 32 |
| Adding PowerPath | 43 |
| Potential Problems | 45 |
| BOOT Device LEDs | 45 |
| 552/554 LED..... | 49 |
| Failure Recovery Scenarios..... | 55 |
| Scheduled Shutdown of Production Local Server..... | 55 |
| Power Loss or Hang of Production Local Server | 56 |
| Sysdump of Production Local Server | 56 |
| Extended Outages | 57 |
| Return the Servers to Production State | 57 |

Introduction

The following release notes provide general guidelines, requirements, and instructions for booting AIX from Symmetrix[®] and CLARiiON[®] systems in an FC-AL or FC-SW environment.

Operating System Requirements

- AIX 4.3.3: Release Name 4330-09, VRMF 4.3.3.75, APAR IY22024.
- AIX 5.1: Release Name 5100-01, VRMF 5.1.0.10, APAR IY21957.
- AIX 5.2: Release Name 5200-01, VRMF 5.2.0.10, APAR IY44479.
- One of the above versions of AIX must be installed on an internal drive to facilitate microcode updates if needed.

Symmetrix Models Supported

- Symmetrix Models 3330/5330, 3430/5430, 3700/5700
- Symmetrix Models 3630/5630, 3830/5830, 3930/5930
- Symmetrix 8000 models
- Symmetrix DMX[™] models

CLARiiON Models Supported

- CLARiiON CX series
- CLARiiON FC4700 series

Symmetrix Director Bit Settings

- For FC-AL (RPQ required), the following fiber director bits should be enabled: T, SC3, H, and UWN.
- For FC-SW, the following fiber director bits should be enabled: T, SC3, PP, and UWN.
- For using the Volume Logix[™] technology that exists in ESN Manager, in addition to bits set for topology, the volume configuration management (VCM) bit must be enabled for the appropriate Fibre Channel Director/Interface Adapter (FA), and the VCM flag must be enabled on the device that is to be the storage device access control database device.

Minimum Enginuity Storage Operating Environment Version for Each Family

- V5266.43 or higher
- V5566.45 or higher
- V5267.35 or higher
- V5567.42 or higher
- V5268.11 or higher
- V5568.36 or higher
- V5669.36 or higher
- V5670.73 or higher

Minimum FLARE Base Operating Environment Versions

- CX400 Series with Array Software Base – 02.06.400.5.004
- CX500 Series with Array Software Base – 02.06.500.5.003
- CX600 Series with Array Software Base – 02.06.600.5.003
- CX700 Series with Array Software Base – 02.06.700.5.003
- FC4700 Series with Array Software Base – 8.51.52

Supported HBAs

- FC 6227 with adapter firmware revision 3.22a1; this comes with boot BIOS 1.31a4. The FC 6227 is a 1 Gigabit FC PCI adapter Type (4-S) with an SC connector. Supported with AIX 4.3.3 and later.
- FC 6228 with adapter firmware revision 3.82a1; this comes with boot BIOS 1.31a4. The FC 6228 is a 2 Gigabit FC PCI adapter Type (4-W) with an LC connector. Supported with AIX 4.3.3 and later.
- FC 6239 with adapter firmware revision 1.00X5; this comes with boot BIOS. The FC 6239 is a 2 Gigabit FC PCI-X adapter Type (5704) with an LC connector. Supported with AIX 5.1 ML04 or AIX 5.2 ML01 and later.
- Either adapter shipping in new systems after November 30, 2001, will have these levels of firmware or higher.

Supported System Firmware

You must obtain the most recent firmware level for the intended pSeries or RS/6000 system from IBM. The website where the microcode updates can be found is:

<http://techsupport.services.ibm.com/server/mdownload/download.html>

Supported EMC Switch Models

Refer to the *EMC Support Matrix* for supported switch configurations.

Installation

This portion of the document describes three methods of installation. The first installation method assumes that you have an existing AIX installation and will be migrating that installation to the Symmetrix or CLARiiON boot device. The second installation method assumes that this will be a brand new installation of AIX directly to the external boot device. In the final method the internal installation will be cloned to the external device using IBM's `alt_disk_install` package leaving two installations to alternate between. The alternate disk installation is the preferred method since it leaves a backup installation of AIX on the host.

Prior to configuring a system for booting from the Symmetrix system, microcode updates are required in two areas. The first required update is for the adapter microcode. This is the code that will be placed on the adapter that is to become your bootable adapter. There are different microcode levels based on the adapter that you will use for booting. The adapters currently supported for fibre boot are adapter feature codes 6227, 6228, and 6239. For the FC 6227, the minimum required microcode level is 3.22a1; for the FC 6228, the minimum required microcode level is 3.82a1; and for the 6239, the minimum required microcode level is 1.00X5.

The second microcode update that is required is the system/service processor combined microcode. This is based on the model number of the host that will be used for fibre boot (for example, 7017-S85). All firmware updates can be obtained from the following website:

<http://techsupport.services.ibm.com/server/mdownload/download.html>

This link also provides complete instructions on updating the firmware and verifying that the process was successful. Many pSeries hosts now come with supported revisions of firmware already installed on both the hosts and the adapters. If your hosts already meet minimum requirements, any steps for installing host or adapter firmware can be safely skipped.

New Installation Overview

1. Identify the host to be used for fibre boot.
2. Identify the type of host adapter to be used for fibre boot and install it.
3. Update the adapter microcode as described in *Host with Existing Adapters* on page 10.
4. Update the system/service processor microcode as described in *Updating the System/Service Processor Combined Microcode* on page 12.
5. Perform the installation to the external boot device as described in *New Installation* on page 15.
6. Install the EMC Object Data Manager Support Package as described in *EMC ODM Support Package Installation* on page 18.

Migration Installation Overview

1. Identify the host to be used for fibre boot.
2. Identify the type of host adapter to be used for fibre boot and install it.
3. Update the adapter microcode as described in *Host with Existing Adapters* on page 10.
4. Update the system/service processor microcode as described in *Updating the System/Service Processor Combined Microcode* on page 12.
5. Perform a migration install to the AIX installation currently installed to your internal disk.
6. Use migration procedures to migrate from your internal hdisk device to the external Symmetrix device as described in *Migration Installation* on page 16.
7. Install the EMC ODM Support Package as described in *EMC ODM Support Package Installation* on page 18.

Cloned Installation Overview

1. Identify the host to be used for fibre boot.
2. Identify the type of host adapter to be used for fibre boot and install it.
3. Update the adapter microcode as described in *Host with Existing Adapters* on page 10.
4. Update the system/service processor microcode as described in *Updating the System/Service Processor Combined Microcode* on page 12.
5. Perform the clone from the current boot volume to the external disk as described in *Cloned Installation Using alt_disk_install* on page 17.
6. Install the EMC ODM Support Package as described in *EMC ODM Support Package Installation* on page 18.

Preparing Your System for Booting from the Symmetrix or CLARiiON Storage Array

Prerequisites

- A supported pSeries or RS/6000 system with AIX installed to the internal disk drive.
- A minimum of one host adapter, FC 6227, FC 6228, or FC 6239, to be used as the boot adapter.
- An Internet connection or method for getting the microcode update files onto the host.
- For AIX 4.3.3, you need the AIX 4.3.3 installation CD (which contains the FC 6227 drivers) or APAR IY20367 (which contains the FC 6227 and FC 6228 drivers).
- For AIX 5.1, you need the AIX 5.1 installation CD (which contains drivers for both the FC 6227 and FC 6228).

- For AIX 5.2, you need the AIX 5.2 installation CD (which contains drivers for the 6227, 6228, and 6239) or APAR IY45849 (which adds support for FC 6239 adapters).

Note: To get APAR IY20367 or APAR IY45849, go to IBM's fix central and use the search feature:

<http://www-1.ibm.com/servers/eserver/support/pseries/index.html>

Recommendations

- If the host does not have the adapter(s) installed, install only the adapter to be used for booting at this time.
- If the host already has multiple adapters installed, to avoid confusion, you may want to remove all fibre adapters except the one intended to be used as the boot adapter.

Adding the Adapters to Your Host

1. Identify the feature code of the adapter you will use for fibre boot: 6227, 6228, or 6239.
2. Insert the adapter into the host. If the system is not hot-pluggable, you will have to power down the host to insert the adapter. Once the adapter is installed, you must install the drivers for it. If the adapter has already been installed in the system and you know the drivers are loaded, you may skip to the section *Host with Existing Adapters*, and proceed directly to page 10, step 2.

If you have an FC 6228 adapter, go to step 4. For the FC 6227 adapter, install the drivers as follows:

- a. Insert the AIX 4.3.3 installation CD (or point to the location of the APAR IY20367 files) and run **smitty install**.
 - b. Select **Install and Update Software**.
 - c. Select **Install and Update from LATEST Available Software**.
 - d. Select your installation media.
 - e. The fileset you need to install is devices.pci.df1000f7. You can search for this fileset using the menu. Once you find it, mark it and complete the install.
 - f. For complete functionality, a reboot is recommended.
3. For the FC 6228 adapter, install the drives as follows:
 - a. Point to the location of the APAR IY20367 files and run smitty install.
 - b. Select Install and Update Software.
 - c. Select Install and Update from LATEST Available Software.
 - d. Select your installation media.
 - e. The fileset you need to install is devices.pci.df1000f9. You can search for this fileset using the menu. Once you find it, mark it and complete the install. If there is a failure while looking for prerequisite software, mark the prerequisite software as well and retry the install.

- f. For complete functionality, a reboot is recommended.

Note: The instructions in steps 3 and 4 are for AIX 4.3.3. If you have AIX 5.1, the drivers for the FC 6227 and FC 6228 are included on the installation CD. In this scenario, use the 5.1 installation CD instead of the media listed in the steps, and then follow the steps for the appropriate adapter. This statement also applies to the AIX 5.2 installation CD and the 6239 adapter.

4. After a successful reboot, the operating system detects the fibre adapter(s). Use the `lsdev` command as follows to display the fibre adapters. Note the name of the adapter on your system.

```
# lsdev -Cc adapter | grep fcs
fcs0    Available 10-68    FC Adapter
fcs1    Available 10-78    FC Adapter
```

5. Proceed to the next section, *Host with Existing Adapters*, and begin with step 2.

Host with Existing Adapters

1. Enter the command `lsdev -Cc adapter | grep fcs`. If the command returns no results, refer to *Adding the Adapters to Your Host*. In our example, we will be using two installed adapters. FCS0 will be used as the boot adapter. Note the name of the adapters.

```
# lsdev -Cc adapter | grep fcs
fcs0    Available 10-68    FC Adapter
fcs1    Available 10-78    FC Adapter
```

Note: If there are devices allocated to the adapter with volume groups varied on, they must be varied off before continuing.

2. Go to:

```
http://techsupport.services.ibm.com/server/mdownload/download.html
```

3. Click the **Adapter Microcode** link, and locate the feature code for the adapter you will be using. You can first click on the description file, which contains instructions for *Determining Adapter Microcode Levels and Updating the Adapter's Microcode*. You should either print the file or make sure you have access to it throughout the update process.
4. After you have familiarized yourself with the update procedure, download the adapter microcode.
5. Use the description file containing the instructions to perform the update.
6. After the update is complete, look at the configuration for your adapter and verify that the levels are as follows:

- For the FC 6227, field Z9 should be as follows:

```
# lscfg -vl fcs4
DEVICE                LOCATION              DESCRIPTION
fcs4                  80-60                FC Adapter

Part Number.....09P1162
EC Level.....D
Serial Number.....KT10504591
Manufacturer.....0010
FRU Number.....09P1173
```

```

Network Address.....10000000C926710B
ROS Level and ID.....02903291
Device Specific.(Z0).....4002206D
Device Specific.(Z1).....10020193
Device Specific.(Z2).....3001506D
Device Specific.(Z3).....02000909
Device Specific.(Z4).....FF101450
Device Specific.(Z5).....02903291
Device Specific.(Z6).....06113291
Device Specific.(Z7).....07113291
Device Specific.(Z8).....20000000C926710B
Device Specific.(Z9).....SS3.22A1
Device Specific.(ZA).....S1F3.22A1
Device Specific.(ZB).....S2F3.22A1
Device Specific.(YL).....U0.2-P1-I10/Q1

```

- For the FC 6228, field Z9 should be as follows:

```

# lscfg -vl fcs0
DEVICE          LOCATION          DESCRIPTION
fcs0            10-58             FC Adapter

```

```

Network Address.....10000000C928C6EF
ROS Level and ID.....02C03891
Device Specific.(Z0).....2002606D
Device Specific.(Z1).....00000000
Device Specific.(Z2).....00000000
Device Specific.(Z3).....02000909
Device Specific.(Z4).....FF401050
Device Specific.(Z5).....02C03891
Device Specific.(Z6).....06433891
Device Specific.(Z7).....07433891
Device Specific.(Z8).....20000000C928C6EF
Device Specific.(Z9).....CS3.82A1
Device Specific.(ZA).....C1D3.82A1
Device Specific.(ZB).....C2D3.82A1
Device Specific.(YL).....U0.1-P1-I5/Q1

```

- For the FC 6239, field Z9 should be as follows:

```

# lscfg -vl fcs0
fcs0            U0.1-P1-I5/Q1    FC Adapter

Part Number.....00P4295
EC Level.....A
Serial Number.....1E338093A7
Manufacturer.....001E
Feature Code/Marketing ID...5704
FRU Number.....00P4297
Device Specific.(ZM).....3
Network Address.....10000000C9388375
ROS Level and ID.....02E01035

```

```

Device Specific.(Z0) .....2003806D
Device Specific.(Z1) .....00000000
Device Specific.(Z2) .....00000000
Device Specific.(Z3) .....03000909
Device Specific.(Z4) .....FF601032
Device Specific.(Z5) .....02E01035
Device Specific.(Z6) .....06631035
Device Specific.(Z7) .....07631035
Device Specific.(Z8) .....20000000C9388375
Device Specific.(Z9) .....HS1.00X5
Device Specific.(ZA) .....H1D1.00X5
Device Specific.(ZB) .....H2D1.00X5
Device Specific.(YL) .....U0.1-P1-I5/Q1

```

Updating the System/Service Processor Combined Microcode

Prerequisites

- A supported pSeries or RS/6000 system with AIX installed to the internal disk drive.
- Host model number (for example, 7017-S85).
- An Internet connection or method for getting the microcode update files onto the host.

1. Go to:

<http://www.rs6000.ibm.com/support/micro/download.html>

2. Under system/service processor-combined microcode, locate the model number of the host you will be using. You can first click on the description file, which contains instructions for downloading microcode to the system/service processor. You should either print the file or make sure you have access to it throughout the update process.
3. After you have familiarized yourself with the update procedure, download the system/service processor combined microcode for your host.
4. Use the description file containing the instructions to perform the update.
5. A reboot is required for this microcode update to take full effect.

Connecting to the CLARiiON Array

Prerequisites

- A supported pSeries or RS/6000 system with AIX installed to the internal disk drive with SAN boot approved host microcode installed and active. See the *EMC Support Matrix* for a list of supported hosts.
- An FC 6227 adapter with 322.a1 microcode level, FC 6228 with 3.82a1 microcode level, or FC 6239 with 1.00X5 microcode level
- Appropriate operating system drivers for your FC 6227, FC 6228, or FC 6239 adapters installed

Adding Physical Connections

1. Start by adding all desired physical connections between the host and the array.
2. For switched fabric, zones must be created to allow the host's HBA(s) to log into the array's SP ports. If a host is to have the visibility into multiple SP ports, make sure that the zones include all of the expected members.
3. Using whatever management utility appropriate for your topology and equipment, verify that all host and array components have logged in and all zones have been created allowing the host's visibility into the array. To verify this visibility:
 - a. Log in to Navisphere[®] Manager and right-click the **Array** component under the **Storage** tab.
 - b. Select **Connectivity Status**. If this is the first time connecting the host to the array, you should see the WWN for each HBA showing up as Fibre Logged In -> Yes and Registered -> No. (If this is a host that had existing connections, this does not apply.)
4. If you see all the hosts HBA(s) logged in, your physical connections are complete. If all of the adapters have not logged in, you must figure out where the breakdown is and address it before continuing.

Completing Your Configuration

(If the host agent is loaded, some of this process may take place automatically. If the host agent is not loaded, you must register your adapters manually.)

1. If your adapters have not been registered before, register them as follows:
 - a. Return to the **Connectivity Status** screen.
 - b. Right-click the **Array** component under the **Storage** tab, and then select **Connectivity Status**.
2. Select each adapter and, in turn, click the **Register** button at the bottom of the window.
3. In the following window, select **AIX hosts Initiator Type=CLARiiON Open, Failover Mode=1, ArrayCommPath=disabled**. These settings are the appropriate settings for PowerPath[®] configurations. Enter any other information such as hostname and IP address.

Note: For a more complete explanation of these settings and how to set them, refer to *Setting Storage System Properties* in the *EMC Storage-System Host Utilities for AIX Administrator's Guide*, P/N 069001137, REV A02.

4. Run configuration manager against each adapter.

```
# cfigmgr -vl fcs0
Time: 1 LEDS: 0x539
Number of running methods: 0
-----
attempting to configure device 'hdisk1'
Time: 1 LEDS: 0x626
invoking /usr/lib/methods/cfgscsidisk -l hdisk1
Number of running methods: 1
```

5. For each initiator you should be able to use the **lsdev** command as follows. Verify your connections to the CLARiiON SP.

```
# lsdev -Cc array
sp0 Available CLARiiON Storage Processor
sp1 Available CLARiiON Storage Processor
```

6. CLARiiON arraycomppath behavior results in devices appearing visible to a host when no LUNs are bound. There is one device visible per initiator (path). Each device has an ID of drivetype unknown. For AIX, as mentioned above, the arraycomppath setting should be disabled. In the event that you do see LUNZ devices, you must **rmdev** these in order to properly utilize your storage groups once they are configured.

```
# lsdev -Cc disk
hdisk1 Available 30-68-00-10,0 16 Bit SCSI Disk Drive
hdisk0 Available 30-68-00-9,0 16 Bit LVD SCSI Disk Drive
hdisk2 Available 50-58-01 FC CLArray LUNZ Disk
hdisk3 Available 80-60-01 FC CLArray LUNZ Disk
```

7. You must now go through the standard CLARiiON practice:
- Create your RAID group(s).
 - Create LUNs within your RAID group(s).
 - Create storage group(s).
 - Assign your LUNs into your storage group(s).
 - Add your host to your storage group(s) for access to your LUNs.

For detailed instructions on each of these, refer to the *EMC ControlCenter Navisphere Manager Administrator's Guide*, P/N 069001125.

8. After you complete step 16 a through e, be sure to use the **rmdev** command to remove any LUNZ devices prior to configuring you real LUNs on the host.
9. When all of these steps are complete, you can create your AIX hdisk devices from the LUNs that are presented to the host by the CLARiiON. These steps are outlined in *Installing the Disk*.

Installing to the Symmetrix or CLARiiON Disk

Even in an environment where you plan on doing a new installation to the external boot device, you may have devices capable of being used as a boot device already allocated to the host. For this reason, and since SANs by nature allow access to a large number of devices, identifying the hdisk to install to can be difficult. Carefully review the following recommendations so you can easily discover the `lun_id` to hdisk correlation.

Prerequisites

- A supported pSeries or RS/6000 system with AIX installed to the internal disk drive.
- An FC 6227 adapter with 322.a1 microcode level, FC 6228 with 3.82a1 microcode level, or FC 6239 with 1.00X5 microcode level or higher.
- Physical connections in place to the array or switch, depending on which topology you will be using.

Recommendations

- Zone the switch or disk array such that the machine being installed can only discover the disk(s) to be installed to. After the installation has completed, you can re-open the zoning so the machine can discover all necessary devices.
- Assign PVID (physical volume identifiers) to all disks from an already-installed AIX system that can access the disks. To do this, use the command `chdev -l hdiskX -a pv=yes`, where *X* is the appropriate number. Create a table mapping PVIDs to physical disks. The PVIDs will be visible from the install menus by selecting the alternate attributes option when selecting the disk to install to.
- Use ODM commands to locate the hdisk to install to. From the main install menu, select **Start Maintenance Mode for System Recovery, Access Advanced Maintenance Functions**, and enter the **Limited Function Maintenance Shell**. At the prompt, you can use the command `odmget -q"attribute=lun_id AND value=0xNN..N" CuAt`, where *0xNN..N* is the lun_id you are looking for. This command will print out the ODM stanzas for the hdisks that have this lun_id. Enter **exit** to return to the installation menus.
- As of this writing, AIX 5.1 does not support multiple paths to the boot device without third-party software. Attach the external disk and the adapter into the SAN Switch. Verify that the adapters have logged in to the switch and the World Wide Names match the new adapters. These may be labeled on the back of the physical adapter, or can be found with the `lscfg -pv | pg` command and search for the `fcs#` stanza of the adapter you are using. The Network Address field in the `lscfg` output is equal to the World Wide Name. Once you have attached the disks, run `cfgmgr` on the AIX host to add the newly discovered disks into the servers ODM. It is a good idea to only attach one physical connection at this time, even if you are later going to use a multipathing solution (e.g., PowerPath) due to problems that could occur during installation.

New Installation

1. If a boot device has already been assigned and discovered by the host via the internal installation of AIX, skip to step 3.
2. After the boot device has been assigned to the host, you must discover the boot device. To do so, use the command `cfgmgr -vl fcsX` where *X* is the number of the adapter that will be used for fibre boot. As the following output shows, unless you already have the EMC ODM Support Package installed, the new external devices will appear as Other FC SCSI Disk Drive. Only the relevant portions of the output are shown.

```
# lsdev -Cc disk
hdisk0 Available 30-68-00-10,0 16 Bit LVD SCSI Disk Drive

# cfgmgr -vl fcs0
Time: 1 LEDS: 0x539
Number of running methods: 0
-----
attempting to configure device 'hdisk1'
Time: 1 LEDS: 0x626
invoking /usr/lib/methods/cfgscsidisk -l hdisk1
Number of running methods: 1

# lsdev -Cc disk
hdisk0 Available 30-68-00-10,0 16 Bit LVD SCSI Disk Drive
hdisk1 Available 10-58-01      Other FC SCSI Disk Drive
```

- The easiest way to track your external boot device is to assign a PVID to the intended installation disk. This is done using the command `chdev -l hdiskX -a pv=yes`.

```
# lspv
hdisk0          000c53cdd955af6e    rootvg
hdisk1          none                None

# chdev -l hdisk1 -a pv=yes
hdisk1 changed

# lspv
hdisk0          000c53cdd955af6e    rootvg
hdisk1          000c53cde980ba9d    None
```

- Write down the PVID; you will then use this to identify your installation disk.
- Place the AIX 4.3.3, 5.1, or 5.2 installation media into the host and shut down. After the shutdown, you can remove the internal disk if you want.
- Boot from the installation media. If the installation media was not a part of your bootlist, you can select it by entering the SMS menus and selecting your installation device as a boot device. Typically, the CD-ROM is a part of your bootlist.
- Select the display to be used as a terminal.
- Select the language to be used.
- Select option 2, **Change/Show Installation Settings and Install**.
- Select option 1, **System Settings**, and ensure that **New and Complete Overwrite** is selected.
- If you have more than one disk allocated to the adapter, to verify your installation disk you can select option 77, **Display More Disk Information**, to see the PVID of the hdisks. Depending on the method you used to identify the installation disk, as you continue selecting 77, additional information such as WWPN, SCSI ID, and LUN ID will appear. Also, on the initial disk selection screen, verify that **Yes** appears under the field marked bootable for the hdisk you will install to. Once you have identified the correct hdisk device(s), select it and then continue. After the installation is complete, the host will automatically reboot from your external boot device.

Note: In a Volume Logix environment, you may see the Volume Logix Database Device. It is typically 7 MB in size. Even though this device is read-only, be sure that you do not accidentally select it as your installation device.

Migration Installation

- Unlike the new installation procedure, because you will be migrating from an internal hdisk there are several ways to update your existing version of AIX. Using your preferred method, update or migrate your current version of AIX to either AIX VRMF 4.3.3.75, VRMF 5.1.0.10, or VRMF 5.2.0.10. The update must be made prior to attempting to boot the operating system from the external device.
- After the migration or update is complete and the system rebooted, begin the migration to external disk by adding the external boot device to the root volume group. For the remaining commands, `hdisk0` is the internal boot device we are migrating from and `hdisk1` is the external boot device we will migrate to:

```
# extendvg -f rootvg hdisk1
```

3. To determine if the boot image is on the disk you want to migrate, run the `lslv -l bootlv` command:


```
# lslv -l hd5
hd5:N/A
PV                COPIES          IN BAND          DISTRIBUTION
hdisk0            001:000:000    100%            001:000:000:000:000
```
4. Migrate the boot logical volume to the external boot device:


```
# migratepv -l hd5 hdisk0 hdisk1
```
5. After you have moved the boot image, clear the boot record on the original disk to prevent possible problems should the host accidentally attempt to boot from that disk:


```
# mkboot -cd /dev/hdisk0
```
6. Update the boot image on your external boot device:


```
# bosboot -ad /dev/hdisk1
```
7. At this point, you can migrate the rest of the contents of the source physical disk to the external boot device:


```
# migratepv hdisk0 hdisk1
```
8. Use the `reducevg` command to remove the original disk device from the root volume group:


```
# reducevg -d rootvg hdisk0
```
9. Use the `bootlist` command to change the bootlist to reflect the new boot device:


```
# bootlist -m normal hdisk1
```
10. Reboot the host:


```
# shutdown -Fr now
```

Cloned Installation Using alt_disk_install

1. If you have not already installed the `alt_disk_install` LPP, begin by selecting that fileset and installing it. It can be located on the AIX installation CD.


```
# installp -a -d/dev/cd0 bos.alt_disk_install
```
2. After the fileset has been successfully installed, `rootvg` can now be cloned to the external disk as opposed to moved, leaving the internal original copy installed to the internal disk. You should notice that this process automatically updates the bootlist for you to the device that you cloned to.


```
# alt_disk_install -C hdisk1
Calling mkszfile to create new /image.data file.
Checking disk sizes.
Creating cloned rootvg volume group and associated logical volumes.
Creating logical volume alt_hd5
Creating logical volume alt_hd8
Creating logical volume alt_hd6
Creating logical volume alt_hd4
Creating logical volume alt_hd1
Creating logical volume alt_hd3
Creating logical volume alt_hd2
Creating logical volume alt_hd9var
Creating /alt_inst/ file system.
Creating /alt_inst/home file system.
Creating /alt_inst/tmp file system.
```

```

Creating /alt_inst/usr file system.
Creating /alt_inst/var file system.
Generating a list of files
for backup and restore into the alternate file system...
Backing-up the rootvg files and restoring them to the
alternate file system...
Modifying ODM on cloned disk.
Building boot image on cloned disk.
forced unmount of /alt_inst/var
forced unmount of /alt_inst/usr
forced unmount of /alt_inst/tmp
forced unmount of /alt_inst/home
forced unmount of /alt_inst
forced unmount of /alt_inst
Changing logical volume names in volume group descriptor area.
Fixing LV control blocks...
Fixing file system superblocks...
Bootlist is set to the boot disk: hdisk1

```

- From this point forward, whichever hdisk you actively boot from will contain `rootvg`. Whenever you boot from your original installation, the alternate hdisk will contain a volume group called `altinst_rootvg`.

```

# lspv
hdisk0          0002583f112f6d71    rootvg
hdisk1          0002583f69d5822b    altinst_rootvg

```

- However, after booting from the cloned installation, you will notice that the original installation will be called `old_rootvg`. Reboot the host now to use the external boot device. Notice we have now booted from `hdisk1`. Updating the bootlist to use whichever device you want to boot from does switching between the two installations.

```

# shutdown -Fr now
# lspv
hdisk0          0002583f112f6d71    old_rootvg
hdisk1          0002583f69d5822b    rootvg

```

Installing the EMC ODM Support Package

Prerequisites

- A supported pSeries or RS/6000 system with AIX installed to the internal or external boot device.
- An Internet connection or method for getting the ODM installation files onto the host.

ODM Download and Installation for Symmetrix and CLARiiON Release 13 or Higher

- From the FTP server, `ftp.EMC.com`, in `/pub/elab/aix/ODM_DEFINITIONS`, download the most recent ODM fileset. Currently for AIX 4.3.3, the fileset is `EMC.AIX.4.3.3.4.tar.Z`, for AIX 5.1, it is `EMC.AIX.5.1.0.1.tar.Z`, and for AIX 5.2.0.0, it is `EMC.AIX.5.2.0.1.tar.Z`. If a more recent fileset is available, use it as a substitute.
- In the `/tmp` directory, uncompress the fileset. Substitute the appropriate filename revision for the command to work properly:

```
uncompress EMC.AIX.X.X.X.tar.Z
```

3. After you uncompress the files, you will be left with a file with the `.tar` extension. Untar the resulting file. Substitute the appropriate filename revision for the command to work properly:

```
tar -xvf EMC.AIX.X.X.X.tar
```
4. If the SMIT menu interface is preferred, invoke `smit installp` from the `/tmp` directory.
Select **Install and Update from LATEST Available Software**.
5. Use the List function to select `/tmp` as the installation directory.
6. Use the List function to select EMC Symmetrix AIX Support Software, EMC CLARiiON Fibre Channel Support Software, or both depending on your configuration.
7. Press **Enter** *after* making all desired changes. You can use the default options.
8. Be sure to scroll all the way down to the bottom of the window to see the Installation Summary, and verify that the `SUCCESS` message is displayed.
9. Reboot the host for all changes to take effect.

Pre-Release 13: Installing CLArrayS3 from the EMC Utilities Kit

Prerequisites

- A supported pSeries or RS/6000 system with AIX installed to the internal or external boot device
- *EMC Storage System Host Utilities for AIX* CD-ROM

CLArrayS3 Installation

1. If the Host Utilities for AIX CD-ROM is not in the server's CD-ROM drive, insert it.
2. If **SMIT** is not running, start it by entering:

```
smit
```

SMIT starts and displays the **System Management** menu.

3. From the **SMIT** menu, follow the menu path:

Software Installation and Maintenance

Install and Update Software

Install and Update from LATEST Available Software

4. Select the CD-ROM drive as follows:
 - a. At the line, **INPUT device/directory for software**, select the option to **List**.
 - b. Select `/dev/cd0 (CD-ROM Drive)` from the list that appears.The **SMIT** window to Install and Update opens automatically.

5. In the **SMIT** window, select the software to install as follows:
 - a. At the **SOFTWARE** to install line, select the option to List. The Multi-select List window appears.
 - b. Select the **CLArrayS3** software.
 - c. Press Enter to continue the install once all selections are made.
6. Be sure to scroll all the way down to the bottom of the window to see the Installation Summary, and verify that the **SUCCESS** message is displayed.
7. Reboot the host for all changes to take effect.

Special Considerations

Volume Logix

Because Volume Logix uses the Volume Logix Database Device to allow or deny access to disk devices, there are a few things that you should be aware of during the installation. To avoid attempting an installation to the Volume Logix Database Device, it is highly recommended that you assign a PVID to your installation disk prior to starting the install. The steps are outlined in *Installing the Disk* on page 10.

Prerequisites for Volume Logix

- An FC 6227, FC 6228, or FC 6239 logged in to the switch port or fibre port of the Symmetrix storage array.
- An ESN Manager/Volume Logix administrative control workstation.
- AIX installation to internal disk drive *or* a host with AIX installed that the boot adapter can be temporarily installed in to.

Recommendations for Volume Logix

On the host with the bootable adapter installed, you should install the Volume Logix utilities and run the `vcmfind` command.

1. On the host with the host bus adapter, look at the configuration of your intended boot adapter using the `lscfg` command. The Network Address field contains the WWN of the adapter. Write the WWN down.

```
# lscfg -vl fcs0
DEVICE                LOCATION                DESCRIPTION
fcs0                   10-58                   FC Adapter

Network Address.....10000000C928C6EF
ROS Level and ID.....02C03891
Device Specific.(Z0).....2002606D
Device Specific.(Z1).....00000000
Device Specific.(Z2).....00000000
Device Specific.(Z3).....02000909
Device Specific.(Z4).....FF401050
Device Specific.(Z5).....02C03891
Device Specific.(Z6).....06433891
Device Specific.(Z7).....07433891
Device Specific.(Z8).....20000000C928C6EF
```

```
Device Specific.(Z9).....CS3.82A1
Device Specific.(ZA).....C1D3.82A1
Device Specific.(ZB).....C2D3.82A1
Device Specific.(YL).....U0.1-P1-I5/Q1
```

The `vcmfnd` command serves three purposes. It will determine the WWN of each host HBA (for hosts that do not have `fpath` installed), assign a recognizable alias known as an ASCII WWN (AWWN) to each WWN (if one is not already present), and enable the GUI to draw and display the ESN topology. See the included screenshot (Figure 1) for an example of this. Using the *ESN Manager/Volume Logix Administrator's Guide*, you will simplify configuration by running `vcmfnd` on the host.

2. At any time after the adapter has logged in to the switch or Symmetrix fibre adapter, you can refresh the display in the ESN Manager/Volume Logix interface. Refreshing the display, among other things, will search the fabric for new WWN(s). After the display has been drawn, you should be able to locate the WWN that you wrote down in step 1 or the ASCII WWN that has been assigned by the `vcmfnd` utility.
3. Using your preferred method as outlined in the *ESN Manager/Volume Logix Administrator's Guide*, allocate a device to the *named* adapter or WWN of the adapter that you will use for booting. Figure 1 shows some of the terminology used.
4. Follow the instructions in *Installing to the Symmetrix or CLARiiON Disk* on page 14. Using the steps outlined in *Installing the Disk* will discover the newly allocated disk down the path it has been assigned. Assign the PVID as well and write it down so that you can select the correct installation disk during the install.

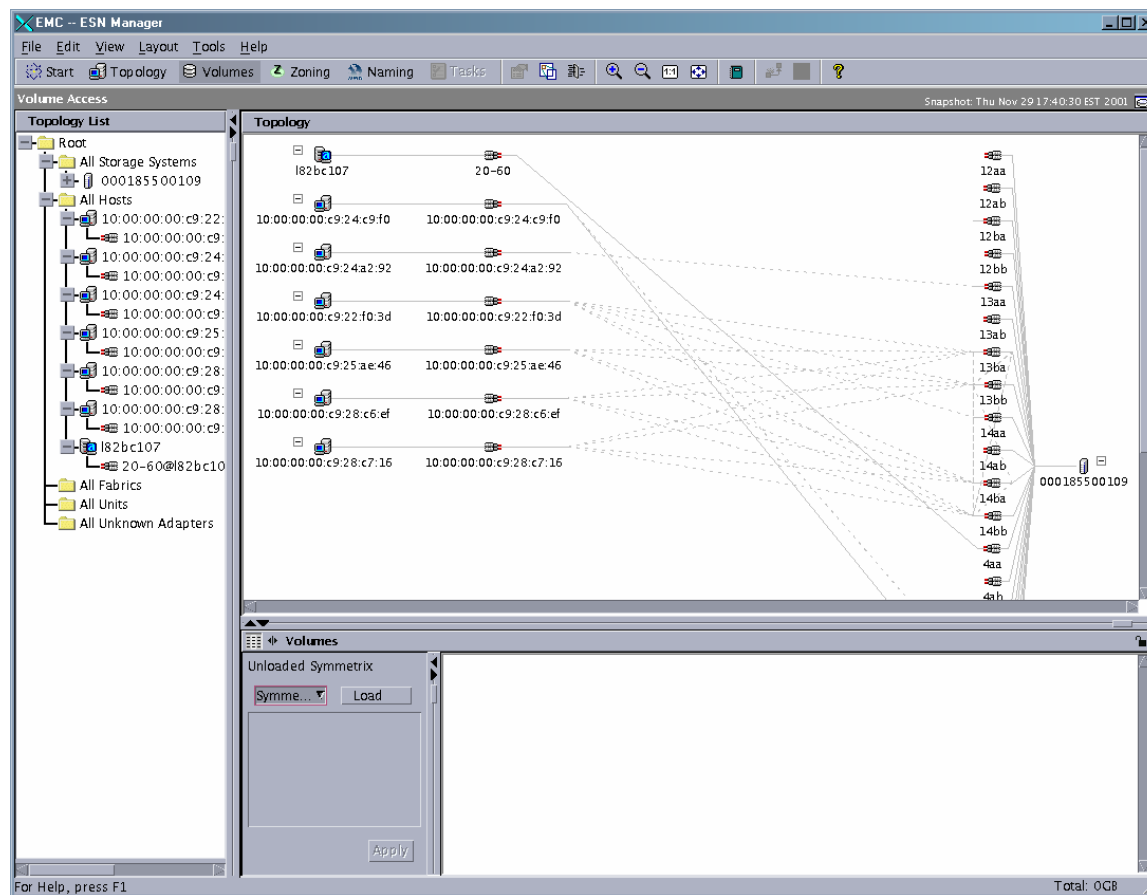


Figure 1. This Shows the ESN Topology Including Host Base WWNs

Note: In Figure 1, notice that some hosts and HBAs are named (have a logical hostname and adapter location code) while others are not named (are depicted only by WWN). If you install and run the Volume Logix Utilities on the host, then your host and HBA will appear as a named host and HBA, making management easier. This is not a requirement, as you can manage device allocation as well by using the WWN to locate the adapter to which you want to allocate devices.

- Proceed with the installation as described in *Installing to the Symmetrix or CLARiiON Disk* on page 14. Take care to follow the precautions outlined to avoid attempting the installation to the Volume Logix Database Device.

PowerPath

On some storage systems, you can use a PowerPath hdiskpower device as a boot device—the device that contains the startup image. (Consult the *EMC Support Matrix* to find out if your storage system supports PowerPath boot devices.) Using a PowerPath hdiskpower device as a boot device provides load balancing and path failover for the boot device.

On AIX, booting from a PowerPath device is supported in SCSI and Fibre Channel environments that include specific versions of EMC Enginuity™ software. Refer to the *EMC Support Matrix* for details. Contact your EMC Customer Support Representative for information about installing Enginuity software.

When you set up a Symmetrix or CLARiiON PowerPath boot device, consider the following:

- All path devices that make up the `hdiskpower` device must be valid AIX boot devices.
- The boot device should not be visible to any other host attached to the same storage system. If using a storage system device as a boot device in an HACMP environment (with or without PowerPath), other hosts should not be able to address the boot device.
- The host's boot list must contain all hdisks that compose the `hdiskpower` device being used as the boot device. Otherwise, the host may fail to boot if one or more paths is disabled while the machine tries to boot.
- At startup, the system searches for an AIX boot image in the boot list, which is a list of hdisks stored in the hardware's NVRAM. If the system fails to boot, you can change the boot list. Use one of the following methods:
 - Boot the system from an installation device (CDROM or tape) into Maintenance Mode. Select the option to access the root volume group, and then run the AIX `bootlist` command from the shell.
 - Enter the **System Management Services** menu when the system starts, and use the **Multiboot** menu options to change the boot list. This method is faster, but it is more difficult to determine the correspondence between devices listed in the menu and the storage-system device you want to add to or remove from the boot list.

Considerations for CLARiiON Storage

PowerPath can be used to enable multipathing and failover to an external boot device on a CLARiiON array, but such a configuration has some functional limitations and extra configuration steps. The primary limitation to CLARiiON boot is that the AIX host must have access to the active SP for the boot LUN in order to be able to boot without user intervention. If all paths to the boot LUN's default SP are failed or the LUN has trespassed to the alternate SP, the host will not boot until access to the LUN through the default SP is restored. Using a switched environment where all HBAs have access to the default SP of the boot LUN can reduce the effects of this restriction.

In a CLARiiON environment, the bootlist for the AIX host should contain only hdisks that correspond to active paths of the default SP. If PowerPath is not configured, these devices will show a PVID in the output of the `lspv` command. Passive hdisks will show a PVID of `None`. If PowerPath is configured, use the `powermt display dev=n` command to examine the boot device and determine which hdisks are active to the default SP.

Once the AIX host is up and running, PowerPath will enable it to survive path failures and trespasses of the boot device.

The pprootdev Tool

The `pprootdev` tool changes AIX configuration rules and updates the boot image so the AIX Logical Volume Manager uses `hdiskpower` devices to vary on the `rootvg` the next time the system boots. The `pprootdev` tool cannot change the state of `rootvg` on a running system. It does, however, modify ODM data that other tools use to determine what devices `rootvg` is using. For this reason, some commands report information that may appear to be incorrect if they are run after `pprootdev` is run and before a system reboot.

Trespasses

When a trespass occurs in a CLARiiON environment, a passive interface becomes the active interface. In this situation, `bosboot` will fail unless you transfer the `rootvg` PVID to the newly active interface. To do so, run the command `emcpassive2active`. Once you have run the command, `bosboot` will succeed.

Run `emcpassive2active` whenever a trespass occurs.

Configuring a New PowerPath Installation

If the system contains sufficient internal storage, install and configure the operating system on the internal device(s). Use the procedure described in *Configuring an Existing PowerPath Installation* to clone the operating system image on the storage system. If there is insufficient internal storage, use the following procedure to install AIX directly onto a storage-system device and use PowerPath to manage multiple paths to the root volume group:

1. Start with a single connection to the storage subsystem. If you are using a switch, only one logical path should be configured.
2. Install AIX on a storage system device that is accessed via fibre adapter.
3. Install the current storage-system drivers.
4. Reboot the host.
5. Use `rmdev -d` to delete any hdisks in the `Defined` state.
6. Install PowerPath. Refer to Chapter 2 of the *PowerPath for UNIX Installation and Administration Guide*.
7. Connect remaining physical connections between the host and the storage system. If you are using a switch, update the zone definitions to the new configuration.
8. Configure an hdisk for each path. In Chapter 2 of the *Powerpath for UNIX Installation and Administration Guide*, refer to the section *Before You Install on an AIX Host*.
9. Run the `powermt config` command.
10. Use the `pprootdev` command to set up multipathing to the root device. Enter: `pprootdev on`
11. Use the `bootlist` command to add all active `rootvg` storage devices to the boot list.
12. Reboot the host.

Configuring an Existing PowerPath Installation

This section describes the process for converting a system that has AIX installed on an internal disk to boot from a logical device on a storage system. The steps in this process are:

- Transferring a complete copy of the operating system from an internal disk to a logical device on the storage system.
- Configuring PowerPath so the root volume group takes advantage of multipathing and failover capabilities.

EMC recommends that you use this procedure. In the event of a problem, you can revert operations to the host's internal disks. First, ensure that the AIX disk includes:

- A copy of the AIX `alt_disk_install` LPP. (The LPP is on the AIX installation CD.)
- The `rte` and `boot_images` filesets.

Then follow these steps:

1. Verify that all device connections to the storage system are established.
2. Verify that all hdisks are configured properly. (Refer to *Before You Install on an AIX Host* in the *PowerPath for UNIX Installation and Administration Guide*.)

3. Locate drives with adequate space. Enter:

```
bootinfo -s hdiskx
```

For this example, assume hdisks 132-134 are adequate with 8 GB total space.

4. Use the `powermt display` command to determine which `hdiskpower` device contains `hdisk132` (the first `hdisk` identified in step 2) as well as all the path `hdisks` for that `hdiskpower`. Enter:

```
powermt display dev=hdisk132
```

The output looks similar to this:

```
Pseudo name=hdiskpower38
Symmetrix ID=000100006216
Logical device ID=006C
state=alive; policy=SymmOpt; priority=0; queued-IOs=0
=====
----- Host ----- - Stor - -- I/O Path - -- Stats ---
### HW Path I/O Paths Interf. Mode State Q-IOs Errors
=====
0 fscsi0 hdisk132 FA 14bA active alive 0 0
1 fscsi1 hdisk223 FA 14bB active alive 0 0
1 fscsi1 hdisk314 FA 14bA active alive 0 0
0 fscsi0 hdisk41 FA 14bB active alive 0 0
```

You see that `hdiskpower38` contains `hdisk132` and that the path `hdisks` for `hdiskpower38` are `hdisk132`, `hdisk223`, `hdisk314`, and `hdisk41`.

5. Run the `powermt config` command.
6. Run the following command to ensure there are PVIDs on all devices:

```
lsdev -Ct power -c disk -F name | xargs -n1 chdev -apv=yes -l
```

7. Use the `rmdev` command to remove all PowerPath devices. PowerPath should remain installed, but all PowerPath devices must be deleted.

```
lsdev -Ct power -c disk -F name | xargs -n1 rmdev -l
```

8. Remove the `powerpath0` device.

```
rmdev -l powerpath0
```

9. Run the `lsdev -Ct power` command. The command output should not list any devices.

10. Determine which `hdisk(s)` on the storage system will receive a copy of the operating system. Run this command to create the copy on the storage system `hdisk(s)`:

```
alt_install_disk -C hdisk_list
```

11. Reboot the system.

12. Specify that all path `hdisks` identified in step 3 are included in the bootlist. Enter:

```
bootlist -m normal hdisk132 hdisk223 hdisk314 hdisk41
```

If you are booting from a CLARiiON storage system, include only hdisks that correspond to active paths of the default SP. Refer to *Considerations for CLARiiON Storage* on page 22 of this document for more information.

13. Run the `pprootdev on` command.

14. Reboot the system.

The `lsvg` Command

The AIX `lsvg` command, when used with the `-p` flag, displays devices in use by the specified volume group. This command, however, is not designed to operate with PowerPath or with storage-system logical devices that are addressable as different hdisk devices. In general, the output of `lsvg -p vgroupname` shows correct information, but several administrative tasks change the ODM and could cause `lsvg` to show misleading information. These tasks include:

- Use of the `pprootdev` tool. This tool changes the ODM and is intended to be used when you expect to reboot the system soon after using `pprootdev`. The `lsvg` command shows misleading device information when run after `pprootdev`. This is not an indication that something is wrong. A reboot corrects the `lsvg` output but it is not required.
- Use of `cfgmgr` to create new hdisk devices after PowerPath is already configured. Always run `powermt config` after adding new devices to include them in PowerPath's configuration.

Actions That Use `bosboot`

After a system boots from a PowerPath device, the `bosboot` command cannot function correctly. This happens because of the state of the configuration after booting from a PowerPath device and the fact that `bosboot` expects the boot device to be an hdisk, not an hdiskpower device. Several system administrative tasks require the system boot image to be rebuilt, and these will fail if `bosboot` cannot run successfully. These tasks include applying certain software patches and using the `mksysb` utility to create a system backup. To address this limitation, the `pprootdev` tool provides a `fix` option that corrects the configuration to allow `bosboot` to work. Run `pprootdev fix` before undertaking any administrative task that runs `bosboot`. This corrects the configuration for `bosboot` but does not change the PowerPath boot switch; the next system boot still uses PowerPath. You need to run `pprootdev fix` only once after a system boots using PowerPath; after that, `bosboot` should function correctly until the system boots again.

Special Recovery Procedures: Path Failure during Bootup

For most situations involving error recovery, refer to the PowerPath manual to seek resolution. The only situation that can occur when booting over fibre (that may not be covered in the PowerPath manual) involves booting the system with a failed path. If the system boots up with a failed path, all devices that were available down the failed path will show up as Defined. In this situation, once the problem with the path has been corrected, the following procedure applies to both CLARiiON and Symmetrix storage arrays:

1. To determine which adapters have failed, use the commands in the following example. From this output we know that the adapter at 20-68 had the problem.

```
(1) (AIX4.3) > lsdev -Cc disk | grep Defined
hdisk4          Defined    20-68-01      EMC Symmetrix FCP Raid1
hdisk6          Defined    20-68-01      EMC Symmetrix FCP Raid1
hdisk7          Defined    20-68-01      EMC Symmetrix FCP Raid1
hdisk5          Defined    20-68-01      EMC Symmetrix FCP Raid1
hdisk8          Defined    20-68-01      EMC Symmetrix FCP Raid1
hdisk9          Defined    20-68-01      EMC Symmetrix FCP Raid1
hdisk10         Defined    20-68-01      EMC Symmetrix FCP Raid1
hdisk11         Defined    20-68-01      EMC Symmetrix FCP Raid1
```

```

hdisk12      Defined    20-68-01      EMC Symmetrix FCP Raid1
hdisk13      Defined    20-68-01      EMC Symmetrix FCP Raid1
hdisk14      Defined    20-68-01      EMC Symmetrix FCP Raid1
hdisk15      Defined    20-68-01      EMC Symmetrix FCP Raid1
hdisk16      Defined    20-68-01      EMC Symmetrix FCP Raid1

```

```
(0) (AIX4.3)> lsdev -Cc adapter | grep fcs
```

```

fcs0      Available 20-68      FC Adapter
fcs1      Available 40-58      FC Adapter

```

2. Once the problem with the failed path has been corrected, place the devices that are in a Defined state into an Available state using the `cfgmgr` command with the `-vl` parameters. Do not use the `emc_cfgmgr` script. As an alternative, the `mkdev -l` command can be used as well.

```
(0) (AIX4.3)> cfgmgr -vl fcs0
```

3. Verify that those devices that were in a Defined state are now in the Available state using the `lsdev` command.

```
(0) (AIX4.3)1> lsdev -Cc disk | grep 20-68
```

```

hdisk4      Available 20-68-01      EMC Symmetrix FCP RAID1
hdisk6      Available 20-68-01      EMC Symmetrix FCP RAID1
hdisk7      Available 20-68-01      EMC Symmetrix FCP RAID1
hdisk5      Available 20-68-01      EMC Symmetrix FCP RAID1
hdisk8      Available 20-68-01      EMC Symmetrix FCP RAID1
hdisk9      Available 20-68-01      EMC Symmetrix FCP RAID1
hdisk10     Available 20-68-01      EMC Symmetrix FCP RAID1
hdisk11     Available 20-68-01      EMC Symmetrix FCP RAID1
hdisk12     Available 20-68-01      EMC Symmetrix FCP RAID1
hdisk13     Available 20-68-01      EMC Symmetrix FCP RAID1
hdisk14     Available 20-68-01      EMC Symmetrix FCP RAID1
hdisk15     Available 20-68-01      EMC Symmetrix FCP RAID1
hdisk16     Available 20-68-01      EMC Symmetrix FCP RAID1

```

4. If you use the `lspv` command at this time, you may notice that the PVIDs for the now-available hdisks may reside on the underlying hdisk devices instead of the `hdiskpower` devices. Run the `powermt config` command to place all of the PVIDs back on to the `hdiskpower` devices.

```
(0) (AIX4.3)> powermt config
```

5. Run `powermt restore` to bring the devices back online if necessary and available for PowerPath for load balanced I/O.

```
(0) (AIX4.3)> powermt restore
```

rootvg Resides in the Symmetrix Storage Array But Is Mirrored across PowerPath Adapters

In some situations, it may be desired to not place `rootvg` under PowerPath control, even though you are booting from the storage array. To help protect your environment from path failure in this scenario, `rootvg` can be mirrored across two different disks across your PowerPath-controlled HBA(s).

Recommendations

- AIX operating system installed and booting properly from the Symmetrix storage array (one cable connected)
- An approved PowerPath configuration
- Most recent EMC ODM definitions for the appropriate OS level of AIX
- PowerPath Version 3.0 or higher with a valid license

As with any procedure involving `rootvg`, an `mksysb` should be performed to back up any critical data. To make sure that you understand the configuration used, included is a brief explanation of the complete configuration used for this section of the documentation. This configuration has a total of eight hdisks. Four belong to the adapter at location code 10-58-01, and four belong to the adapter at 40-60-01. If you look at the output from the second command (`inq`), which is short for inquiry, you will see that devices 01A and 019 appear only once, whereas devices 01B, 01C, and 01D appear twice. You can see that the set of devices that appear twice will be under the control of PowerPath, where the set that appears once will not be. The first set of devices will contain `rootvg` and instead of being controlled by PowerPath, they will be protected by LVM mirroring.

```
(0) (AIX5.1)> lsdev -Cc disk
```

```
hdisk1 Available 10-58-01 EMC Symmetrix FCP Raid1
hdisk0 Available 40-60-01 EMC Symmetrix FCP Raid1
hdisk2 Available 10-58-01 EMC Symmetrix FCP Raid1
hdisk3 Available 10-58-01 EMC Symmetrix FCP Raid1
hdisk4 Available 10-58-01 EMC Symmetrix FCP Raid1
hdisk5 Available 40-60-01 EMC Symmetrix FCP Raid1
hdisk6 Available 40-60-01 EMC Symmetrix FCP Raid1
hdisk7 Available 40-60-01 EMC Symmetrix FCP Raid1
```

```
(127) (AIX5.1)>inq
```

```
Inquiry utility, Version V7.2-154 (Rev 16.0) (SIL Version V4.2-154)
Copyright (C) by EMC Corporation, all rights reserved.
For help type inq -h.
```

```
.....
```

```
-----
DEVICE          :VEND      :PROD          :REV   :SER NUM   :CAP (kb)
-----
/dev/rhdisk0    :EMC        :SYMMETRIX     :5568  :0901A140  :2208960
/dev/rhdisk1    :EMC        :SYMMETRIX     :5568  :09019130  :2208960
/dev/rhdisk2    :EMC        :SYMMETRIX     :5568  :0901B130  :2208960
/dev/rhdisk3    :EMC        :SYMMETRIX     :5568  :0901C130  :2208960
/dev/rhdisk4    :EMC        :SYMMETRIX     :5568  :0901D130  :2208960
/dev/rhdisk5    :EMC        :SYMMETRIX     :5568  :0901B140  :2208960
/dev/rhdisk6    :EMC        :SYMMETRIX     :5568  :0901C140  :2208960
/dev/rhdisk7    :EMC        :SYMMETRIX     :5568  :0901D140  :2208960
```

Recommendations

1. Use the commands shown in the following example to verify your configuration by making sure that you have two unique disks allocated to a minimum of two unique Fibre Channel paths. It is crucial that these devices are unique as they will not be under the control of the PowerPath product, and they must be unique to allow for LVM mirroring.

```
(0) (AIX5.1)>lsdev -Cc disk
```

```
hdisk1 Available 10-58-01 EMC Symmetrix FCP Raid1
hdisk0 Available 40-60-01 EMC Symmetrix FCP Raid1
```

```
(0) (AIX5.1)>inq
```

```
Inquiry utility, Version V7.2-154 (Rev 16.0) (SIL Version V4.2-154)
Copyright (C) by EMC Corporation, all rights reserved.
For help type inq -h.
```

```
..
```

```

-----
DEVICE          :VEND      :PROD          :REV   :SER NUM   :CAP (kb)
-----
/dev/rhdisk0    :EMC        :SYMMETRIX    :5568  :0901A140  :2208960
/dev/rhdisk1    :EMC        :SYMMETRIX    :5568  :09019130  :2208960
-----

```

Notice that there are currently two disks available, hdisk0 and hdisk1, each available down two different physical paths. The output from the `inq` command shows that in the field labeled SERIAL NUM there are two unique devices, one device 01A and the other device 019. Notice as well that the capacities are the same.

2. Since the installation started out by having a properly installed fibre boot system, we now have to extend `rootvg` to include the hdisk to which we will mirror. In this configuration, the installation was performed to hdisk1 so `rootvg` will be extended onto hdisk0.

```

(0) (AIX5.1)>extendvg rootvg hdisk0
0516-1254 extendvg: Changing the PVID in the ODM.
0516-014 installpv: The physical volume appears to belong to another
                    volume group.
000c53cd00004c00
0516-631 extendvg: Warning, all data belonging to physical
                    volume hdisk0 will be destroyed.
extendvg: Do you wish to continue? y(es) n(o)? yes

```

3. Tell the LVM to create a mapped mirror onto the new destination hdisk.

```

(0) (AIX5.1)>mirrorvg -m rootvg hdisk0
0516-1124 mirrorvg: Quorum requirement turned off, reboot system for this
                    to take effect for rootvg.
0516-1126 mirrorvg: rootvg successfully mirrored, user should perform
                    bosboot of system to initialize boot records. Then, user must modify
                    bootlist to include: hdisk0 hdisk1.

```

4. Use the `bosboot` command to reinitialize the boot sectors of the hdisks.

```

(0) (AIX5.1)>bosboot -ad /dev/ipldevice

bosboot: Boot image is 14300 512 byte blocks.

```

5. Set and then verify the bootlist order using the `bootlist` command. Be sure to include any additional boot devices that are specific to your environment.

```

(0) (AIX5.1)>bootlist -m normal hdisk0 hdisk1

(0) (AIX5.1)>bootlist -m normal -o
hdisk0
hdisk1

```

6. Verify the mirroring. Because there are two hdisks in the mirror, there will be twice the number of physical partitions in use.

```

(0) (AIX5.1)>lsvg -l rootvg
rootvg:
  LV NAME      TYPE      LPs   PPs   PVs   LV STATE      MOUNT POINT
  hd5          boot      2     4     2     closed/syncd  N/A
  hd6          paging   16    32    2     open/syncd    N/A
  hd8          jfslog   1     2     2     open/syncd    N/A
  hd4          jfs      2     4     2     open/syncd    /
  hd2          jfs     116   232   2     open/syncd    /usr
  hd9var       jfs      3     6     2     open/syncd    /var

```

| | | | | | | |
|-----------|---------|-----|-----|---|------------|-------|
| hd3 | jfs | 6 | 12 | 2 | open/syncd | /tmp |
| hd1 | jfs | 1 | 2 | 2 | open/syncd | /home |
| hd10opt | jfs | 5 | 10 | 2 | open/syncd | /opt |
| lg_dump1v | sysdump | 256 | 256 | 1 | open/syncd | N/A |

7. Install PowerPath according to the *PowerPath Installation Guide*. Because `rootvg` is not under the control of PowerPath, do not use the section explaining how to install PowerPath with external boot.

Special Recovery Procedures: Path Failure

If one or more paths fail in a mirrored configuration, you actually have two restorations to perform at all times. You must always keep in mind that once PowerPath has been restored to its original functioning condition, you must also make sure that `rootvg` is synchronized as well.

1. One indication that `rootvg` is no longer synchronized will be errors in the error log similar to the following. Also, you can check at anytime by using the `lsvg` command. Notice that under the column LV STATE, the logical volume is opened but us marked stale.

```
(0) (AIX5.1) > errpt | more
IDENTIFIER    TIMESTAMP    T C    RESOURCE_NAME  DESCRIPTION
EAA3D429     0502112902  U S    LVDD           PHYSICAL PARTITION MARKED STALE
EAA3D429     0502112902  U S    LVDD           PHYSICAL PARTITION MARKED STALE
EAA3D429     0502112902  U S    LVDD           PHYSICAL PARTITION MARKED STALE
EAA3D429     0502112902  U S    LVDD           PHYSICAL PARTITION MARKED STALE
EAA3D429     0502112902  U S    LVDD           PHYSICAL PARTITION MARKED STALE
EAA3D429     0502112902  U S    LVDD           PHYSICAL PARTITION MARKED STALE
```

```
(0) (AIX4.3) l82bc102/> lsvg -l rootvg
rootvg:
LV NAME      TYPE      LPs    PPs    PVs    LV STATE      MOUNT POINT
hd5          boot      1      2      2      closed/syncd  N/A
hd6          paging    64     128    2      open/syncd    N/A
hd8          jfslog    1      2      2      open/stale    N/A
hd4          jfs       1      2      2      open/stale    /
hd2          jfs       67     134    2      open/stale    /usr
hd9var       jfs       1      2      2      open/stale    /var
hd3          jfs       2      4      2      open/stale    /tmp
hd1          jfs       1      2      2      open/syncd    /home
```

2. Once the problem with the failed path has been fixed to correct this condition, the `rootvg` volume group can be synchronized using the `varyonvg` command:


```
(0) (AIX5.1) >varyonvg rootvg
```
3. Always make sure that after a period of time, all of your root volume group partitions eventually return to open/syncd status.

Using MKSYSB for Backup and Restore

Prerequisites

For AIX 4.3, minimum level EMC ODM definitions kit is 4.3.3.4. For AIX 5.1 and 5.2, minimum ODM definition is 5.1.0.0.

1. Before backing up your system you are required to run `pprotectdev fix`. If you do not execute this command, your backup process will fail right away.
2. After running `pprotectdev fix`, back up your host using `mksysb`; use any command syntax that you would ordinarily use for this process.

3. Because the AIX OS is not currently designed to deal with duplicate paths to devices without third-party software support, before starting the restore process you must reduce the number of paths to your boot device to one.
4. Boot from your `mksysb` tape.
5. If your configuration has not changed, it is best to restore using defaults. If you do need to select a new for restoration, select it, and then proceed with the restore. As part of the restoration process, the host will reboot automatically.

Note: During the first automated reboot that occurs after the `mksysb restore` is complete, several scripts are run to rebuild devices since by default `recover_devices = yes` is set. This includes the `powerpath0` driver that is necessary for PowerPath to properly configure devices and also boot externally.

6. Log in to the host. Since during this reboot the `powerpath0` driver necessary for operation was just created, it is necessary to reboot the host again. This will replace the exclusive disk reservation on `rootvg` with a group reserve, allowing multiple paths to access this device.
7. Add the remaining cables to the host, and then run `cfgmgr -vl` against each individual adapter as it is added.
8. Run `powermt config` to reintegrate all of your additional devices and paths back into PowerPath.
9. Run `pprootdev fix`.
10. Run `bosboot -ad /dev/ipldevice`.
11. Update the bootlist to include all of the hdisks containing `rootvg`.

SRDF

This section describes requirements and procedures for configuring and booting a pair of AIX hosts from a set of Symmetrix SRDF® disks. **SRDF, when used in conjunction with fibre boot, is currently not supported with AIX 4.3.3 and AIX 5.1, and generally supported with AIX 5.2.** This installation method assumes that all minimum host and HBA firmware revisions, operating-system revisions, and Symmetrix code revisions documented have been followed. This installation method also assumes that you have an existing AIX installation on internal disks at the minimum levels, and that this copy of AIX will be cloned using `alt_disk_install` to the Symmetrix boot device. The administrator can then boot off of either the internal or external disks as needed for maintenance or troubleshooting. It is possible to perform these tasks without using internal disks, but not practical. Several problems can occur when performing upgrades, or troubleshooting without them, in a disaster recovery situation.

For SRDF to function as a disaster recovery solution for a bootable `rootvg` image, the two physical systems that are used (local host and remote host) need to be similar in model (i.e., `chrp` and `mp`), adapter type, and layout to ensure that the servers can both boot off the same `rootvg` volume group image.

Figure 2 is an example of a basic SRDF configuration that will be used as a reference throughout this section.




Figure 2. Basic SRDF Configuration

SRDF Installation Overview

1. Identify the host(s) to be used for fibre boot.
2. Identify the type of host adapter to be used for fibre boot and install it as described in *Adding the Adapters to Your Host*.
3. Update the adapter microcode as described in *Updating the Adapter Microcode*.
4. Update the system/service processor microcode as described in *Updating the System/Service Processor Combined Microcode*.
5. Install the EMC ODM support package.

6. Attach the external disk
7. Install Solutions Enabler SYMCLI any license keys required.
8. Configure SRDF device groups.
9. Perform the clone from the current boot volume to the external disk using `alt_disk_install` (recommended) or perform a new installation.
10. Test cloned `rootvg` on local and remote hosts.

Identify the Host(s) to Be Used for Fibre Boot

For SRDF to function as a disaster recovery solution for a bootable `rootvg` image, the two physical systems that are used (local host and remote host) need to be similar in model (i.e., `chrp` and `mp`), adapter type, and layout to ensure that the servers can both boot off the same `rootvg` volume group image.

Identify the Type of Host Adapter to Be Used for Fibre Boot and Install It

If the hosts do not have the adapter(s) installed, install only the adapter to be used for booting at this time. If the host already has multiple adapters installed, to avoid confusion you may want to remove or remove the cables from all fibre adapters except the one intended to be used as the boot adapter(s).

Install the EMC ODM Support Package

AIX 5.2

1. From the FTP server, `ftp://ftp.emc.com/pub/elab/aix/ODM_DEFINITIONS`, download the most recent ODM fileset for the AIX version you are using. Currently, the fileset for AIX 5.2 is `EMC.AIX.5.2.0.1.tar.Z`. If a more recent fileset is available, use it as a substitute.
2. In the `/tmp` directory, uncompress the fileset. Substitute the appropriate filename revision for the command to work properly:


```
uncompress EMC.AIX.5.2.0.1.tar.Z
```
3. After you uncompress the fileset, you will be left with a file with the `.tar` extension. Untar the resulting file. Substitute the appropriate filename revision for the command to work properly:


```
tar -xvf EMC.AIX.5.2.0.1.tar
```
4. If the SMIT menu interface is preferred, invoke `smit installp` from the `/tmp` directory. Select **Install and Update from LATEST Available Software**.
5. Use the List function to select `/tmp` as the installation directory.
6. Use the **List** function to select **EMC Symmetrix AIX Support Software and EMC Symmetrix Fibre Channel Support Software**.
7. Press **Enter** *after* making all desired changes. You can use the default options.
8. Be sure to scroll all the way down to the bottom of the window to see the Installation Summary, and verify that the `SUCCESS` message is displayed.
9. Reboot the host for all changes to take effect.

Attach the External Disk

1. Attach the external disk and the adapter into the SAN switch.
2. Verify that the adapters have logged in to the switch and the World Wide Names match the new adapters. These may be labeled on the back of the physical adapter, or can be found with the `lscfg -pv | pg` command and search for the `fc#` stanza of the adapter you are using. The Network Address field in the `lscfg` output is equal to the World Wide Name.
3. Once you have attached the disks, run `cfgmgr` on the AIX host to add the newly discovered disks into the servers ODM. It is a good idea to only attach one physical connection at this time, even if you are later going to use a multipathing solution (e.g., PowerPath) due to problems that could occur during installation.
4. Zone the switch or disk array such that the machine being installed can only discover the disk(s) to be installed to. After the installation has completed, you can reopen the zoning so the machine can discover all necessary devices. Zone the R1 devices to the local host; zone the R2 devices to the remote host.
5. Assign PVIDs (physical volume identifiers) to all disks from the internal image of the AIX system that can access the disks. To do this, use the command `chdev -l hdiskX -a pv=yes`, where *X* is the appropriate number.

6. Create a table mapping PVIDs to physical disks.

7. Add the World Wide Name to the table with the following command:

```
lsattr -El hdisk# | grep ww_name
```

8. Add LUN IDs to the table with the following commands:

```
lsattr -El hdisk# | grep lun_id.
```

This information will be very helpful in determining the exact path and disk that you will be installing to from the boot/install menus or troubleshooting any problems. The PVIDs will be visible from the install menus by selecting the alternate attributes option when selecting the disk to install to. The World Wide Name and LUN IDs will be visible on the SMS multiboot bootlist. Do this from both servers because some values will be unique per server.

Example:

```
chdev -l hdisk2 -a pv=yes
chdev -l hdisk3 -a pv=yes
.
.
.
```

Local host disk details:

```

Hdisk Number          hdisk2
Current Status =      Available
Location =            20-58-01
Description =         EMC Symmetrix FCP RDF1 Raid1
World Wide Name =     0x5006048acc82de1
Logical Storage array Number (LUN) ID = 0x0
Volume Group Assigned = None
Physical Volume (PVID) ID = 0002175f3043527f

```

```

Hdisk Number =       hdisk3
Current Status =     Available
Location =           20-58-01
Description =        EMC Symmetrix FCP RDF1 Raid1
World Wide Name =    0x5006048acc82de1
Logical Storage array Number (LUN) ID = 0x1000000000000
Volume Group Assigned = None
Physical Volume (PVID) ID = 0002175f30435da6

```

Remote host disk details:

```

Hdisk Number          hdisk2
Current Status =      Available
Location =            20-58-01
Description =         EMC Symmetrix FCP RDF2 Raid1
World Wide Name =     0x5006048acc82dff
Logical Storage array Number (LUN) ID = 0x100000000000000
Volume Group Assigned = None
Physical Volume (PVID) ID = 0002175f3043527f

```

```

Hdisk Number          hdisk3
Current Status =      Available
Location =            20-58-01
Description =         EMC Symmetrix FCP RDF2 Raid1
World Wide Name =     0x5006048acc82dff
Logical Storage array Number (LUN) ID = 0x110000000000000
Volume Group Assigned = None
Physical Volume (PVID) ID = 0002175f30435da6

```

Verify that the external disks that you have selected to hold the external rootvg image are recognized to be bootable by both the local and remote AIX servers. To do this, run the following command:

```
bootinfo -B hdisk#
```

The command will return a 1 if the disk is capable of booting that particular server. A return code of 0 means that you will not be able to boot off that disk device.

Install Solutions Enabler SYMCLI Software

1. Install the Solutions Enabler SYMCLI software on both servers.
2. Add the `.profile` of the `root` user on local and remote AIX servers:


```
export PATH=$PATH:/usr/symcli/bin
export PATH=$PATH:/usr/lpp/EMC/Symmetrix/bin
```
3. Install license keys by running the `/usr/symcli/bin/symlmf` script and inserting the appropriate license keys.
4. Verify that the license keys are valid by running local commands on both local and remote hosts:

```
emc_cfgmgr — Detects new Symmetrix devices on the AIX servers.
symcfg discover — Builds or rebuilds the database file that stores device details.
symcfg list — Lists the Symmetrix storage arrays that are visible to that server.
symcfg list -RA all — Shows the RDF Directors available to that server.
symrdf list — Shows state of all SRDF devices.
```

Configure SRDF Device Groups

An SRDF configuration has at least one source storage array and one target storage array. SRDF configurations may transfer data in a unidirectional or bidirectional manner. In a unidirectional configuration, all R1 devices reside in the source Symmetrix storage array and all R2 devices in the target Symmetrix storage array. Under normal operating conditions, data flows from the R1 devices to the target R2 devices.

The SRDF volumes should have been associated when your EMC CE/SE configured the Symmetrix storage arrays. Device groups are system-specific, so it is a good idea to create group names that reflect the type activity you would perform on that group.

Create device groups for your boot devices, which will allow you to perform operations to all the devices at the same time. You may have one device group to hold the bootable `rootvg` image, and another device group to hold a database or mission-critical data. By using separate device groups, you can perform actions on them individually if necessary.

1. Create a device group to hold your disks containing the `rootvg` image:

on local server:

```
symdg create prodrootvg -type RDF1
```

on remote server:

```
symdg create prodrootvg -type RDF2
```

2. Add the devices you have selected for that external copy of `rootvg`. This example uses `hdisk2` and `hdisk3`. Choose ones that are deemed bootable with the `bosboot -B` command.

```
symdev list -r1
```

```
Symmetrix ID: 000187900087
```

| | Device Name | Directors | | Device | | | Cap (MB) |
|------|--------------|--------------|----------|-----------|-----|--|-------------|
| Sym | Physical | SA :P DA :IT | Config | Attribute | Sts | | |
| 0000 | /dev/rhdisk2 | 02C:1 01A:C0 | RDF1+Mir | N/Grp'd | RW | | 4155 |
| 0001 | /dev/rhdisk3 | 02C:1 16B:C0 | RDF1+Mir | N/Grp'd | RW | | 4155 |

```
0002 /dev/rhdisk4 02C:1 16A:C1 RDF1+Mir N/Grp'd RW 4155
0003 /dev/rhdisk5 02C:1 01B:C1 RDF1+Mir N/Grp'd RW 4155
```

```
Hdisk Number =          hdisk2
Current Status =        Available
Location =              20-58-01
Description =           EMC Symmetrix FCP RDF1 Raid1
World Wide Name =       0x5006048accc82de1
Logical Storage array Number (LUN) ID = 0x0
Volume Group Assigned = none
Physical Volume (PVID) ID = 0002175f3043527f
```

```
Hdisk Number =          hdisk3
Current Status =        Available
Location =              20-58-01
Description =           EMC Symmetrix FCP RDF1 Raid1
World Wide Name =       0x5006048accc82de1
Logical Storage array Number (LUN) ID = 0x1000000000000
Volume Group Assigned = none
Physical Volume (PVID) ID = 0002175f30435da6
```

Add the disks to the `prodrootvg` device group with the following commands (perform on both local and remote hosts):

On local server (your RANGE value may be different):

```
symld -g prodrootvg -RANGE 000:001 addall dev
```

on remote server (your RANGE value may be different):

```
symld -g prodrootvg -RANGE 010:011 addall dev
```

3. You can now verify that the disks are part of the `prodrootvg` group.

```
symdev list -r1
```

```
Symmetrix ID: 000187900087
```

| Device Name | | Directors | | | Device | | | | |
|-------------|--------------|-----------|--------|----------|--------|--------|--------------|-----|----------|
| Sym | Physical | SA | :P | DA | :IT | Config | Attribute | Sts | Cap (MB) |
| 0000 | /dev/rhdisk2 | 02C:1 | 01A:C0 | RDF1+Mir | | | Grp'd | RW | 4155 |
| 0001 | /dev/rhdisk3 | 02C:1 | 16B:C0 | RDF1+Mir | | | Grp'd | RW | 4155 |
| 0002 | /dev/rhdisk4 | 02C:1 | 16A:C1 | RDF1+Mir | | | N/Grp'd | RW | 4155 |
| 0003 | /dev/rhdisk5 | 02C:1 | 01B:C1 | RDF1+Mir | | | N/Grp'd | RW | 4155 |

You can also view specifics about the RDF devices in the group with the following command:

```
symrdf -g prodrootvg query
```

```
Device Group (DG) Name      : prodrootvg
DG's Type                   : RDF1
DG's Symmetrix ID           : 000187900087
```

| Source (R1) View | | | | | Target (R2) View | | | | | MODES | | |
|------------------|----------|--------|-----|-------|------------------|-------|------|-------|--------|--------|-----|--------------|
| Standard | Logical | Device | Dev | ST | LI | ST | Dev | E | Tracks | Tracks | MDA | STATE |
| | | | | A | N | A | | | | | | |
| | | | | T | K | T | | | R1 Inv | R2 Inv | | RDF Pair |
| | | | | E | S | E | | | Tracks | Tracks | | STATE |
| DEV001 | 0000 | RW | | 0 | 0 | RW | 0010 | WD | 0 | 0 | S.. | Synchronized |
| DEV002 | 0001 | RW | | 0 | 0 | RW | 0011 | WD | 0 | 0 | S.. | Synchronized |
| Total | | ----- | | ----- | | ----- | | ----- | | ----- | | |
| | Track(s) | | | 0 | 0 | | | | 0 | 0 | | |
| | MB(s) | | | 0.0 | 0.0 | | | | 0.0 | 0.0 | | |

Note: This output shows you the R1-to-R2 relationship as well as the State of the pair.

- We can now establish the SRDF mirrors (if not performed previously) that will synchronize and overwrite the R2 side with an exact copy of the data of R1:

```
symrdf -g prodrootvg establish -full
```

- Depending if there is actual data on the R1 devices, this activity may take some time. To verify when the process is complete, perform the following:

```
symrdf -g prodrootvg -synchronized verify
```

If you receive the following message, the establish has completed:

```
All devices in the RDF group 'prodrootvg' are in the 'Synchronized' state.
```

Perform the Clone to the External Disk Using alt_disk_install (Recommended)

You are now ready to perform the `alt_disk_install` procedure to create a clone of the `rootvg` file system currently on the internal disks.

This example uses `hdisk2` and `hdisk3`. Be careful not to select the Volume Logix Database Device that is represented as an `hdisk` on the AIX servers. In our example, it is shown as:

```
hdisk18 Available 20-58-01      EMC Symmetrix FCP RDF1 Raid1
```

You can verify which hdisk representation is your Volume Logix database using the following commands:

```
inq -dev /dev/rhdisk# -page0
```

```
-----
BYTE 109:  VCMState  VCMDevice  GKDevice  MetaDevice  Shared
SA State   1         1         0         0         1
-----
```

Look for the VCMDevice flag. If it is set to 1, then that is your Volume Logix device. Do not select that device by mistake when installing or booting.

If it is not already installed on your system, install the `bos.alt_disk_install` LPP from AIX CDs. Also, install the latest patch from IBM:

```
installp -a -d/dev/cd0 bos.alt_disk_install
```

As an added safety precaution, make a backup `mksysb` of your progress up to this point.

```
smitty mksysb
```

Create the clone image using the `alt_disk_install` command:

```
Alt_disk_install -C hdisk2 hdisk3
```

```
Calling mkszfile to create new /image.data file.
Checking disk sizes.
Creating cloned rootvg volume group and associated logical volumes.
Creating logical volume alt_hd5
Creating logical volume alt_hd6
Creating logical volume alt_hd8
Creating logical volume alt_hd4
Creating logical volume alt_hd2
Creating logical volume alt_hd9var
Creating logical volume alt_hd3
Creating logical volume alt_hd1
Creating logical volume alt_hd10opt
Creating logical volume alt_dumplv
Creating /alt_inst/ file system.
Creating /alt_inst/home file system.
Creating /alt_inst/opt file system.
Creating /alt_inst/tmp file system.
Creating /alt_inst/usr file system.
Creating /alt_inst/var file system.
Generating a list of files
for backup and restore into the alternate file system...
Backing-up the rootvg files and restoring them to the
alternate file system...
Modifying ODM on cloned disk.
Building boot image on cloned disk.
forced unmount of /alt_inst/var
forced unmount of /alt_inst/usr
forced unmount of /alt_inst/tmp
forced unmount of /alt_inst/opt
forced unmount of /alt_inst/home
forced unmount of /alt_inst
forced unmount of /alt_inst
Changing logical volume names in volume group descriptor area.
```

```
Fixing LV control blocks...
Fixing file system superblocks...
Bootlist is set to the boot disk: hdisk2
```

Note: The bootlist has now been updated to hdisk2.

New Installation to R1 Disk (Optional)

If you choose not to use the `alt_disk_install` procedure you can perform a new installation from the AIX installation media (CD, tape, or NIM image) using the below procedure:

1. Make sure that you have gathered the disk information for the specific disks you plan to install to. The installation process will display disk details similar to the following:


```
SCSI 4356 MB FC Harddisk id =@5006048accc82dff,10000000000000 ( slot=1 )
SCSI 4356 MB FC Harddisk id =@5006048accc82dff,11000000000000 ( slot=1 )
```
2. Put the installation media into the host and shut down. After the shutdown, you can remove the internal disk if desired (not recommended).
3. Boot from the installation media. If the installation media device was not part of the bootlist, you can select it by entering the SMS menus and selecting your installation device as you boot device.
4. Select the display to be used as the console during install.
5. Select the language to be used.
6. Select option 2, **Change/Show Installation Settings and Install**.
7. Select option 1, **System Settings**, and ensure that **New and Complete Overwrite** is selected.
8. If you have more than one disk allocated to the adapter, verify your installation disk by selecting option 77, **Display More Disk Information**, to see the PVID of the hdisks. Depending on the method that you used to identify the installation disk, as you continue selecting 77, additional information such as WWPN, SCSI ID, and LUN ID will appear. Also, on the initial disk selection screen, verify that **Yes** appears under the field marked bootable for the hdisk to which you will install. Once you have identified the correct hdisk device(s), select it and then continue. After the installation is complete, the host will automatically reboot from your external boot device.
9. Because you installed a new image, you will need to install the EMC ODM support package to this new image (follow the procedure in *Installing the EMC ODM Support Package* on page 18).

Test SRDF rootvg Boot on Local and Remote Hosts

SRDF Control Operation Relationships

In a SRDF `rootvg` environment, there are two primary control operations to consider.

- **Split / establish** — A split operation stops remote mirroring between the R1 and R2 devices and makes R2 available to the remote host and R1 available to the local host. The establish operation would then **overwrite R2** with an exact copy of R1 and break the remote connection to R2 preventing any test data from interfering with production data.
- **Failover / failback** — A failover operation switches data processing from the R1 side to the R2 side. This disables the local host to R1 and the mirroring between the R1 and R2 side. It makes R2 available to the remote host. The failback operation will **Update the R1 copy** with changes made on the R2 side. This would be used in a true disaster recovery situation where you would not want to lose the changes made to the `rootvg` data.

You will now notice (if you performed an `alt_disk_install`) that you have an `altinst_rootvg` volume group on the external disks. The configuration is set to boot off of those disks during the next reboot of the local host.

Example:

lspv

```
hdisk0          00601910967f88f3          rootvg
hdisk1          00015458bfe2e31f          rootvg
hdisk2          0002175f3043527f          altinst_rootvg
hdisk3          0002175f30435da6          altinst_rootvg
```

Reboot the local host to activate the production `rootvg` on the external disk and verify that it can indeed boot off the cloned `rootvg` on the external devices.

Shutdown -Fr

After booting from the cloned installation, you will notice that the original installation is called `old_rootvg`.

lspv

```
hdisk0          00601910967f88f3          old_rootvg
hdisk1          00015458bfe2e31f          old_rootvg
hdisk2          0002175f3043527f          rootvg
hdisk3          0002175f30435da6          rootvg
```

All that must be performed now to boot from either copy is to change the bootlist to point to whichever copy of `rootvg` we may need. `Bootlist -m normal -o` listd the current boot device. `Bootlist -m normal hdisk#` changes the copy of `rootvg` from which you want to boot. If you have forgotten the original boot device, use the `alt_disk_install -q` command against any disk in an inactive copy of `rootvg` to find which device contains the bootable image.

alt_disk_install -q hdisk1

```
hdisk0 # (This is output from the above command.)
```

At this point, you have three copies of `rootvg` from which to boot: One copy on each server's internal disk, and a "floating" production copy on external disk. Next, boot off the local hosts internal disks again so we can move the cloned copy to the remote host and boot off of it:

bootlist -m normal -o

```
hdisk2 # (This is output from the above command.)
```

bootlist -m normal hdisk0

bootlist -m normal -o

```
hdisk0 # (This is output from the above command.)
```

shutdown -Fr

With both servers booted off their internal copies of `rootvg`, run the following command to fail over the R1 devices to the remote host. This is done to ensure that the shared `rootvg` copy contains the proper

configuration in their ODM files for all devices. This will write-disable the local hosts connection to the R1, suspend RDF link traffic, and read/write-enable R2 to the remote host.

symrdf -g prodrootvg failover

Execute an RDF 'Failover' operation for device group 'prodrootvg' (y/[n]) ? **y**

An RDF 'Failover' operation execution is in progress for device group 'prodrootvg'. Please wait...

```
Write Disable device(s) on SA at source (R1).....Done.
Suspend RDF link(s).....Done.
Read/Write Enable device(s) on RA at target (R2).....Done.
```

The RDF 'Failover' operation successfully executed for device group 'prodrootvg'.

You should now see the state of the devices changed to "Failed over" in both the symrdf list and symrdf -g prodrootvg query commands.

symrdf list

Symmetrix ID: 000187900087

Local Device View

| Sym Dev | RDF RDev | Typ:G | STATUS | | MODES | | R1 Inv Tracks | R2 Inv Tracks | RDF S T A T E S | | |
|---------|----------|-------|--------|----|-------|-----|---------------|---------------|-----------------|------|-------------|
| | | | SA | RA | LNK | MDA | | | Dev | RDev | Pair |
| 0000 | 0010 | R1:2 | WD | RW | NR | S.. | 0 | 0 | WD | RW | Failed Over |
| 0001 | 0011 | R1:2 | WD | RW | NR | S.. | 0 | 0 | WD | RW | Failed Over |

symrdf -g prodrootvg query

```
Device Group (DG) Name      : prodrootvg
DG's Type                   : RDF1
DG's Symmetrix ID          : 000187900087
```

| Source (R1) View | | | | | Target (R2) View | | | | | MODES | |
|------------------|------|--------|--------|--------|------------------|------|--------|--------|----------|-------|-------------|
| Standard | ST | LI | ST | | Standard | ST | LI | ST | | | |
| Logical | T | R1 Inv | R2 Inv | K | Logical | T | R1 Inv | R2 Inv | RDF Pair | | |
| Device | Dev | E | Tracks | Tracks | S | Dev | E | Tracks | Tracks | MDA | STATE |
| DEV001 | 0000 | WD | 0 | 0 | NR | 0010 | RW | 0 | 0 | S.. | Failed Over |
| DEV002 | 0001 | WD | 0 | 0 | NR | 0011 | RW | 0 | 0 | S.. | Failed Over |
| Total | | | | | | | | | | | |
| Track(s) | 0 | | 0 | | 0 | | 0 | | | | |
| MB(s) | 0.0 | | 0.0 | | 0.0 | | 0.0 | | | | |

To avoid confusion, shut down the local host at this time to avoid a conflict on the network with two servers responding to the primary IP_LABEL of the local host. You can also use the `uname -a` output to keep track of the physical server you are on.

On the local host, issue the following commands:

```
>uname -a
AIX local 1 5 000154584C00
>shutdown -F
```

On the remote host, verify the proper boot device by using the PVID and LUN ID. Note that the World Wide Name will be different because of the adapter used on this server. This is the R2 device that mirrors the bootable R1 device on the local host.

```
Hdisk Number =                hdisk2
Current Status =              Available
Location =                    20-58-01
Description =                 EMC Symmetrix FCP RDF2 Raid1
World Wide Name =             0x5006048acc82dff
Logical Storage array Number (LUN) ID = 0x10000000000000
Volume Group Assigned =       None
Physical Volume (PVID) ID =   0002175f3043527f
```

In this example, the external bootable drive is defined as hdisk2 on both servers; however, the internal bootable disk on the remote host is hdisk1. Change the bootlist to boot the external bootable disk and reboot the remote host.

```
bootlist -m normal -o
hdisk1 # (This is output from the above command.)
```

```
bootlist -m normal hdisk2
```

```
bootlist -m normal -o
hdisk2 # (This is output from the above command.)
```

```
uname -a
AIX remote 1 5 006019104C00
```

```
shutdown -Fr
```

Adding PowerPath

As with any procedure involving `rootvg`, an `mksysb` should be performed to back up any critical data. The installation procedure migrates volume groups built on Symmetrix devices from AIX hdisks to PowerPath `hdiskpower` devices. After installing PowerPath, you need to vary on the volume groups and mount any file systems that use those volume groups, but you do not need to reconfigure the volume groups themselves.

1. Note the adapters that will be used for the PowerPath configuration. Remember that only one of them should have any devices configured down it at this time. In this configuration, there are two adapters:

```
> lsdev -Cc adapter | grep fcs

fcs0      Available   20-58      FC Adapter
fcs1      Available   20-60      FC Adapter
```

2. Verify that devices are configured down only one path at a time. You should see your `hdisk` devices with all of the location codes matching only one of the adapters you determined in step 1. The exception to this would be your internal `hdisk` if one were present. All of these `hdisks`, with the exception of the internal disk, are available for the adapter located at 20-58, which corresponds to `fcs0` (this text has been truncated).

```

> lsdev -Cc disk
hdisk2 Available 20-58-01 EMC Symmetrix FCP RDF1 Raid1
hdisk3 Available 20-58-01 EMC Symmetrix FCP RDF1 Raid1
.
.
.

```

3. Install PowerPath as described in the *PowerPath for UNIX Installation Guide*. Refer to the section on AIX. It will guide you through properly installing the PowerPath software and checking your registration.

4. Because `rootvg` is varied on, there is an exclusive disk reserve placed on the device preventing it from being accessed down the second path. The next three steps are necessary to replace that exclusive disk reserve with a group reserve. The group reserve allows the device to be accessed down multiple paths. Start by using the PowerPath configuration command if you have not already done so.

```

> powermt config
powerpath0 created

```

5. Use the `pprootdev` tool to enable PowerPath protection of `rootvg`. The `pprootdev` tool is used to change AIX configuration rules and update the boot image so that the AIX Logical Volume Manager will use `hiskpower` devices to vary on the root volume group the next time the system is booted.

```

> pprootdev on
bosboot: Boot image is 8669 512 byte blocks.
PowerPath boot is enabled for the next system boot

```

6. To actually place the group reserve on this device, the host must now be rebooted.

```

> shutdown -Fr

```

7. Once the system is back up, the `rootvg` device should have the group reservation set. You can now attach and configure the second path to the Symmetrix storage array at this time. Using the `emc_cfgmgr` script will configure automatically all paths and then update the PowerPath configuration.

```

> /usr/lpp/EMC/Symmetrix/bin/emc_cfgmgr

```

8. Verify that all of your `hdisk`s have been configured properly down each path. This process is covered in greater detail in the *PowerPath for UNIX Installation Guide* (in this example, `hdisk2` and `hdisk23` are the same disk through different paths, represented by `hdiskpower0`).

```

> lsdev -Cc disk
...
hdisk2 Available 20-58-01 EMC Symmetrix FCP RDF1 Raid1
hdisk3 Available 20-58-01 EMC Symmetrix FCP RDF1 Raid1
...
hdiskpower0 Available 20-58-01 PowerPath Device
hdiskpower1 Available 20-58-01 PowerPath Device
...
hdisk23 Available 20-60-01 EMC Symmetrix FCP RDF1 Raid1
hdisk24 Available 20-60-01 EMC Symmetrix FCP RDF1 Raid1
...

```

9. Run the `pprootdev fix` command:

```

> pprootdev fix
You may now run bosboot.
PowerPath boot remains enabled for the next system boot.

```

10. Run `bosboot` to reinitialize the boot record on the boot device:

```
> bosboot -ad /dev/ipldevice
```

11. Update the bootlist to include all of the underlying `hdisk` devices of `rootvg`. Be sure to include any additional devices that you may need.

```
> bootlist -m normal hdisk2 hdisk23
```

12. Reboot the host again:

```
> shutdown -Fr
```

13. When the host is back online you can verify that you have booted from the PowerPath-protected `rootvg` by using the `lspv` command. Notice that `rootvg` now belongs to an `hdiskpower` device (this text has been truncated).

```
> lspv
```

```
...
hdiskpower0      00015458eb951002          rootvg
hdiskpower1      00015458eb9524df          rootvg
```

14. Use the `powermt display` command to check the status of the various paths to your disks:

```
>powermt display
```

```
Symmetrix logical device count=20
CLARiiON logical device count=0
=====
----- Host Bus Adapters ----- I/O Paths ----- Stats -----
### HW Path Summary Total Dead IO/Sec Q-IOs Errors
=====
  0 fscsi0      optimal      20         0         -         0         0
  1 fscsi1      optimal      16         0         -         0         0
```

Potential Problems

BOOT Device LEDs

If the server won't boot or hangs at E1F5 or 20EE000B appears on the LED panel or other errors occur, the server can't find the boot device. The issue is related to the way the server uniquely identifies the drives. When you change anything in the path to the drive, the firmware recognizes this as a different `hdisk`. On the shared copy of `rootvg`, there is a definition for `hdisk2`, which is the R1 path to the drive. When you attempt to boot that image, it may conflict with the definition of the original R1 path that no longer exists. This is where the documentation is invaluable.

The external copy of `rootvg` is configured to boot off of `hdisk2` through the World Wide Name path of `0x5006048accc82de1 LUN ID 0 (0x0)` in its normal state.

```
Hdisk Number =          hdisk2
Current Status =        Available
Location =             20-58-01
Description =          EMC Symmetrix FCP RDF1 Raid1
World Wide Name =      0x5006048accc82de1
Logical Storage array Number (LUN) ID = 0x0
Volume Group Assigned = none
Physical Volume (PVID) ID = 0002175f3043527f
```

When you split or fail over the shared copy of `rootvg` and boot the remote host, the path to the device will have changed. The AIX server firmware may decide that the change in the physical path is a different `hdisk`. In our example, it is:

```
SCSI 4356 MB FC Harddisk id=@5006048acc82dff,10000000000000 ( slot=1 )
```

This matches our remote host `hdisk2` definition from earlier.

```

Hdisk Number          hdisk2
Current Status =     Available
Location =           20-58-01
Description =        EMC Symmetrix FCP RDF2 Raid1
World Wide Name =    0x5006048acc82dff
Logical Storage array Number (LUN) ID = 0x10000000000000
Volume Group Assigned = None
Physical Volume (PVID) ID = 0002175f3043527f

```

To resolve, boot to the SMS screen by pressing a **F1** on a GUI (or type **1** on a tty) on the console before the fifth keyword appears on the banner screen:

```

RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000
RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000
RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000
RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000
RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000
RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000
RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000
RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000
RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000
RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000
RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000
RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000
RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000
RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000
RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000
RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000
RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000
RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000
RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000 RS/6000

```

```

memory      keyboard      network      scsi

```

<type "F1" on a graphics Keyboard or "1" on a tty>

```

RS/6000 Firmware
Version WIL03115
(c) Copyright IBM Corp. 2000 All rights reserved.
-----

```

System Management Services

```

1 Display Configuration
2 Multiboot
3 Utilities
4 Select Language

```

<select "2" to enter the Multiboot menu>


```

Hdisk Number =                hdisk19
Current Status =              Available
Location =                    20-58-01
Description =                  EMC Symmetrix FCP RDF2 Raid1
World Wide Name =              0x5006048accc82dff
Logical Storage array Number (LUN) ID = 0x100000000000000
Volume Group Assigned =        rootvg
Physical Volume (PVID) ID =    0002175f3043527f

```

>lspv

```

hdisk0          00601910967f88f3          old_rootvg
hdisk1          00015458bfe2e31f          old_rootvg
hdisk19         0002175f3043527f          rootvg
hdisk20         0002175f30435da6          rootvg

```

You may also get errors running Symmertrix commands due to the path differences.

>symrdf -g prodrootvg query

Error opening the gatekeeper device for communication to the Symmetrix

If this happens, run `symcfg discover` to dynamically rediscover and repair the connections.

symcfg discover

This operation may take up to a few minutes. Please be patient...

>symrdf -g prodrootvg query

```

Device Group (DG) Name      : prodrootvg
DG's Type                   : RDF1
DG's Symmetrix ID          : 000187900087

```

| Source (R1) View | | | | | Target (R2) View | | | | | MODES | |
|------------------|-------|--------|--------|--------|------------------|------|--------|------------|--------|----------|-------------|
| Standard | LI | ST | LI | ST | Standard | LI | ST | LI | ST | | |
| Logical | T | R1 Inv | R2 Inv | K | Logical | T | R1 Inv | R2 Inv | K | RDF Pair | |
| Device | Dev | E | Tracks | Tracks | S | Dev | E | Tracks | Tracks | MDA | STATE |
| DEV001 | 0000 | WD | 0 | 0 | NR | 0010 | RW | 221 | 0 | S.. | Failed Over |
| DEV002 | 0001 | WD | 0 | 0 | NR | 0011 | RW | 63 | 0 | S.. | Failed Over |
| Total | | | | | | | | | | | |
| Track(s) | ----- | | 0 | 0 | ----- | | | 284 | 0 | | |
| MB(s) | 0.0 | 0.0 | | | 8.9 | 0.0 | | | | | |

Note: There are now invalid R1 tracks due to the changes during the boot cycle.

LED 554 Bootable Media Recovery Procedure

1. Boot in Maintenance Mode (from NIM in this example). You will be prompted to select Root Volume Group.

```

Access a Root Volume Group

Type the number for a volume group to display the logical volume information
and press Enter.

1) Volume Group 0052904a00004c00000000fba68bb032 contains these disks:
   hdisk0 17357 10-60-00-0,0
2) Volume Group 0052904a00004c00000000fc5ac13c3a contains these disks:
   hdisk2 21502      20-58-01

Choice: 2

```

2. Select the entry corresponding to your rootvg, (check VGID).

Notice that **hdisk2** is the device configured by the AIX kernel (Maintenance Mode) in this example.

```

Volume Group Information
-----
Volume Group ID 0052904a00004c00000000fc5ac13c3a includes the following
logical volumes:

   hd5      hd6      hd8      hd4      hd2      hd9var
   hd3      hd1      hd10opt  lv00     lv01
-----

Type the number of your choice and press Enter.

1) Access this Volume Group and start a shell
2) Access this Volume Group and start a shell before mounting filesystems

99) Previous Menu

Choice [99]: 1

```

3. Select **Access this Volume Group and start a shell**.

4. Once rootvg is imported, lspv shows the following volumes:

```
# lspv
hdisk0          0052904a0c2d3952      old_rootvg
hdisk1          none                    None
hdisk3          0052904a5ac12d8a      rootvg
hdisk4          none                    None
hdisk5          none                    None
hdisk6          0052904a5ac12d8a      rootvg
hdisk7          none                    None
hdisk8          none                    None
```

Notice that hdisk2 is missing. This is because the `lspv` command gets the info from the ODM on the disk. According to the previous display, `rootvg` belongs to hdisk3 and hdisk6 on ODM.

5. Check the boot logical volume.

```
# lslv -m hd5
hd5:N/A
LP   PP1  PV1          PP2  PV2          PP3  PV3
0001 0001 hdisk3
```

Notice that `hd5` belongs to `hdisk3`.

6. Remove the additional path (hdisk6 in this example):

```
# rmdev -dl hdisk6
hdisk6 deleted
```

7. Execute the `bosboot` command, but be sure to execute it against the correct device.

Remember that the `rootvg` disk configured in maintenance mode is `hdisk2`; the system has also configured a special device file `/dev/ipldevice` pointing to `hdisk2`.

Consequently, the `bosboot` command must be executed on `/dev/ipldevice`.

If you select `hdisk3` (device associated with `hd5`), the `bosboot` command will fail with the following error message :

```
# bosboot -ad /dev/hdisk3
odmget: Could not retrieve object for CuAt, odm errno 5904

bosboot: Boot image is 18040 512 byte blocks.
#
```

Execute the `bosboot` command on `/dev/ipldevice`:

```
# bosboot -ad /dev/ipldevice
bosboot: Boot image is 18040 512 byte blocks.
```

8. Once the boot image is successfully created, reboot the host:

```
# shutdown -Fr
```

Special Considerations when PowerPath Protection of rootvg Is Enabled

As explained before, if you receive a 552/554 LED when trying to boot off the R2 copy of rootvg, it may be caused by the multiple paths to your boot disk.

It may also be caused because `pprootdev on` (enabled) and the PVID of rootvg disk are associated with `hdiskpower` device.

With AIX 5.1, the PVID of rootvg *must* be associated with `hdisk` even if PowerPath protection of rootvg is enabled (`pprootdev on`).

To resolve this issue, the `pprootdev fix` command must be integrated into a script called by `/etc/inittab` at boot time or run after every boot. This will ensure that PVID is always associated with `hdisk(s)`.

Special Considerations when Data Volume Groups Are Configured on hdiskpower Devices

When data volume groups are configured on R1 devices under PowerPath control, these volume groups may not be activated automatically in case of failover to the R2 side.

This is because the `hdiskpower` PVID entry of R1 device remains configured in the ODM when a failover is executed. Furthermore, when booting off the R2 side, new `hdiskpower` devices will be configured for data disks. That results in duplicate PVID entries, as shown next:

```
# odmget -q 'value like *0030379ba4960689* and attribute=pvid' CuAt
```

```
CuAt:
```

```
name = "hdiskpower4"
attribute = "pvid"
value = "0030379ba49606890000000000000000"
type = "R"
generic = "D"
rep = "s"
nls_index = 2
```

```
CuAt:
```

```
name = "hdiskpower5"
attribute = "pvid"
value = "0030379ba49606890000000000000000"
type = "R"
generic = "D"
rep = "s"
nls_index = 2
```

```
# lspv
```

```
...
hdiskpower0      0052904a5ac12d8a      rootvg
hdiskpower3      0030379b63cb7a4c      None
hdisk5           none                  None
hdisk10          none                  None
hdiskpower5      0030379ba4960689      datavg
```

```
# lsdev -Cc disk
```

```
...
hdiskpower4 Defined 3p-08-01 PowerPath Device
hdiskpower5 Available 4k-08-01 PowerPath Device
```

To fix this issue, remove the `hdiskpower` entry in Defined state:

```
# rmdev -dl hdiskpower4
```

```
hdiskpower4 deleted
```

Then, you can varyonvg the data volume group and mount the file systems:

```
# varyonvg datavg

# lsvg -o
datavg
rootvg
```

Failure Recovery Scenarios

Scheduled Shutdown of Production Local Server

During normal procedures in this example we have the local server booting off of the production copy of rootvg. To intentionally move this over to the remote server for maintenance or scheduled outage, perform the following steps:

On the local server:

Stop any applications and perform a shutdown to stop all processes and close files.

```
>shutdown -F now
```

On remote server:

If you are planning to retain the changes made on the rootvg during this outage, perform a “failover” operation. If you are going to discard the changes, perform a “split” operation. This will stop the synchronized traffic from the R1 to the R2 side and make the R2 disk available to the remote host.

```
> symrdf -g prodrootvg failover
```

```
Execute an RDF 'Failover' operation for device group 'prodrootvg' (y/[n]) ? y
```

```
An RDF 'Failover' operation execution is in progress for device group
'prodrootvg'. Please wait...
```

```
Write Disable device(s) on SA at source (R1).....Done.
Suspend RDF link(s).....Done.
Read/Write Enable device(s) on RA at target (R2).....Done.
```

```
The RDF 'Failover' operation successfully executed for device group 'prodrootvg'.
```

Change the bootlist from the active rootvg on the remote server to the “production” R2 disk.

```
>bootlist -m normal hdisk2 hdisk3
```

Reboot the remote server off of the production rootvg.

```
>shutdown -Fr
```

Power Loss or Hang of Production Local Server

On the local server:

Record any LEDs or obtain any information that may help diagnose the cause of the problem.

On the remote server:

If you are planning to retain the changes made on the rootvg during this outage, perform a failover operation. If you are going to discard the changes, perform a split operation. This will stop the synchronized traffic from the R1 to the R2 side and make the R2 disk available to the remote host.

```
> symrdf -g prodrootvg failover
```

```
Execute an RDF 'Failover' operation for device group 'prodrootvg' (y/[n]) ? y
```

```
An RDF 'Failover' operation execution is in progress for device group
'prodrootvg'. Please wait...
```

```
Write Disable device(s) on SA at source (R1).....Done.
Suspend RDF link(s).....Done.
Read/Write Enable device(s) on RA at target (R2).....Done.
```

```
The RDF 'Failover' operation successfully executed for device group
'prodrootvg'.
```

Change the bootlist from the active rootvg on the remote server to the production R2 disk.

```
>bootlist -m normal hdisk2 hdisk3
```

Reboot the remote server off of the production rootvg:

```
>shutdown -Fr
```

Sysdump of Production Local Server

If the local server performs a system dump, you may not be able to diagnose the system dump on the remote host. You may have to fail back to the original server to run analysis tools to debug the error.

On the local server:

Record any LEDs or obtain any information that may help diagnose the cause of the problem.

On the remote server:

If you are planning to retain the changes made on the rootvg during this outage, perform a failover operation. If you are going to discard the changes, perform a split operation. This will stop the synchronized traffic from the R1 to the R2 side and make the R2 disk available to the remote host.

```
> symrdf -g prodrootvg failover
```

```
Execute an RDF 'Failover' operation for device group 'prodrootvg' (y/[n]) ? y
```

```
An RDF 'Failover' operation execution is in progress for device group
'prodrootvg'. Please wait...
```

```
Write Disable device(s) on SA at source (R1).....Done.
Suspend RDF link(s).....Done.
Read/Write Enable device(s) on RA at target (R2).....Done.
```

```
The RDF 'Failover' operation successfully executed for device group
'prodrootvg'.
```

Change the bootlist from the active `rootvg` on the remote server to the production R2 disk:

```
>bootlist -m normal hdisk2 hdisk3
```

Reboot the remote server off of the production `rootvg`:

```
>shutdown -Fr
```

Extended Outages

There may be times when the local server would be unavailable for an extended period of time. Consider using the swap function to reverse the roles of the R1 and R2 sides in this situation. That way, you can keep your local Symmetrix system synchronized during the production activity during the failover, and the delay of synchronizing the R1 copy with all the R2 updates is reduced.

This procedure assumes that the original R1 Symmetrix system is functioning properly and that you want to maintain the production copy of the data that is being updated on the R2 side in a failover situation.

Current state assumes a failover has been previously issued. You should see the state of the devices changed to “Failed over” in both the `symrdf list` and `symrdf -g prodrootvg query` commands.

symrdf list

Symmetrix ID: 000187900087

Local Device View

| ----- | | | | | | | | | | | |
|-------|------|--------|----|----|-------|--------|-----------------|--------|-----|------|-------------|
| | | STATUS | | | MODES | | RDF S T A T E S | | | | |
| Sym | RDF | ----- | | | ----- | R1 Inv | R2 Inv | ----- | | | |
| Dev | RDev | Typ:G | SA | RA | LNK | MDA | Tracks | Tracks | Dev | RDev | Pair |
| ----- | | | | | | | | | | | |
| 0000 | 0010 | R1:2 | WD | RW | NR | S.. | 0 | 0 | WD | RW | Failed Over |
| 0001 | 0011 | R1:2 | WD | RW | NR | S.. | 0 | 0 | WD | RW | Failed Over |
| . | | | | | | | | | | | |
| . | | | | | | | | | | | |
| 0010 | 0000 | R2:3 | RW | RW | NR | S.. | 212 | 0 | RW | WD | Failed Over |
| 0011 | 0001 | R2:3 | RW | RW | NR | S.. | 61 | 0 | RW | WD | Failed Over |

Note that there are invalid R1 tracks as activity continues on the R2 side.

Perform the following on the secondary system:

```
Symrdf -g prodrootvg swap -refresh R1
Symrdf -g prodrootvg establish
```

This command updates the R1 side with the changes made since the failover command, and then reverses the roles of R1/R2 disks. Both copies from this point forward would stay in a synchronized state.

Return the Servers to Production State

To return the servers to their normal state (local host on external `rootvg` and remote host on internal `rootvg`), perform the following:

Reboot the remote host on its original internal `rootvg` copy.

```
>bootlist -m normal -o
hdisk19 # (This is output from the above command.)
```

```
>bootlist -m normal hdisk1
```

```
>bootlist -m normal -o
hdisk1 # (This is output from the above command.)
```

```
>shutdown -Fr
```

When the remote host comes back on its internal disks, you can then set the SRDF pairs back to their normal state and reboot the local host with them. Use the appropriate command to return the R1/R2 pairs to the original state. If you chose to “failover”, then run “failback.” If you chose to “split”, then run “establish.”

```
>symrdf -g prodrootvg failback
```

```
Execute an RDF 'Failback' operation for device
group 'prodrootvg' (y/[n]) ? y
```

```
An RDF 'Failback' operation execution is
in progress for device group 'prodrootvg'. Please wait...
```

```
Write Disable device(s) on RA at target (R2).....Done.
Suspend RDF link(s).....Done.
Merge device track tables between source and target.....Started.
Devices: 0000-0001 ..... Merge device
track tables between source and target.....Done.
Resume RDF link(s).....Done.
Read/Write Enable device(s) on SA at source (R1).....Done.
```

```
The RDF 'Failback' operation successfully executed for
device group 'prodrootvg'.
```

Verify that the R1 and R2 devices are synchronized and there are no invalid tracks.

```
>symrdf -g prodrootvg query
```

```
Device Group (DG) Name      : prodrootvg
DG's Type                   : RDF1
DG's Symmetrix ID          : 000187900087
```

| Source (R1) View | | | | | Target (R2) View | | | | | MODES | |
|------------------|-------|--------|--------|-----|------------------|--------|--------|--------|----------|-------|--------------|
| Standard | ST | | | | LI | ST | | | | | |
| Logical | A | | | | N | A | | | | | |
| Device | T | R1 Inv | R2 Inv | K | T | R1 Inv | R2 Inv | | RDF Pair | | |
| Dev | E | Tracks | Tracks | S | Dev | E | Tracks | Tracks | MDA | STATE | |
| DEV001 | 0000 | RW | 0 | 0 | RW | 0010 | WD | 0 | 0 | S.. | Synchronized |
| DEV002 | 0001 | RW | 0 | 0 | RW | 0011 | WD | 0 | 0 | S.. | Synchronized |
| Total | | | | | | | | | | | |
| Track(s) | ----- | | 0 | 0 | ----- | | 0 | 0 | | | |
| MB(s) | ----- | | 0.0 | 0.0 | ----- | | 0.0 | 0.0 | | | |

You can boot the local host back on its internal disk to verify a problem was fixed, or use the SMS multiboot menus to change the bootlist to its external `rootvg` copy and reboot off the external devices. Use `lspv` to verify which copy is your active one:

>lspv

```
hdisk0          00601910967f88f3          rootvg
hdisk1          00015458bfe2e31f          rootvg
hdisk2          0002175f3043527f          altinst_rootvg
hdisk3          0002175f30435da6          altinst_rootvg
```

>bootlist -m normal -o

hdisk0 # (This is output from the above command.)

>bootlist -m normal hdisk2

>bootlist -m normal -o

hdisk2 # (This is output from the above command.)

>shutdown -Fr

Once your local host has been booted of external disk, you may have to once again run `symcfg discover` to undo the link changes.

>symrdf -g prodrootvg query

Error opening the gatekeeper device for communication to the Symmetrix

If this happens it is due to the paths on the remote host being severed. Run `symcfg discover` to dynamically repair.

>symcfg discover

This operation may take up to a few minutes. Please be patient...

Notice now that you have the R2 devices in a Defined state on the external disk copy of `rootvg`:

>lsdev -Cc disk

```
hdisk0  Available  10-60-00-8,0  16 Bit SCSI Disk Drive
hdisk1  Available  10-60-00-9,0  16 Bit SCSI Disk Drive
hdisk2  Available  20-58-01      EMC Symmetrix FCP RDF1 Raid1
hdisk3  Available  20-58-01      EMC Symmetrix FCP RDF1 Raid1
.
.
.
hdisk19 Defined    20-58-01      EMC Symmetrix FCP RDF2 Raid1
hdisk20 Defined    20-58-01      EMC Symmetrix FCP RDF2 Raid1
```