



**EMC Fibre Channel with  
Emulex Host Bus Adapters in the  
Solaris Environment**

**P/N 30-001-156  
REV A02  
OCTOBER 2003**

**EMC Corporation**  
*Corporate Headquarters:*  
Hopkinton, MA 01748 -9103  
1-508 -435 -1000  
[www.EMC.com](http://www.EMC.com)

**Copyright © 2001–2003 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.

### **Trademark Information**

EMC<sup>2</sup>, EMC, Symmetrix, AutoIS, Celerra, CLARiiON, CLARAlert, DG, E-Infostructure, HighRoad, Navisphere, PowerPath, ResourcePak, SnapView /IP, SRDF, VisualSAN, where information lives, WideSky, and The EMC Effect are registered trademarks and EMC Automated Networked Storage, EMC ControlCenter, EMC Developers Program, EMC Enterprise Storage, EMC Enterprise Storage Network, EMC OnCourse, EMC Proven, EMC Snap, Access Logix, AutoAdvice, Automated Resource Manager, AutoSwap, AVALONidm, C-Clip, CacheStorm, Celerra Replicator, Centera, CentraStar, CLARevent, Connectrix, CopyCross, CopyPoint, CrosStor, Direct Matrix, Direct Matrix Architecture, EDM, E-Lab, Engenuity, FarPoint, FLARE, GeoSpan, InfoMover, MirrorView, NetWin, OnAlert, OpenScale, Powerlink, PowerVolume, RepliCare, SafeLine, SAN Architect, SAN Copy, SAN Manager, SDMS, SnapSure, SnapView, StorageScope, SupportMate, SymmAPI, SymmEnabler, Symmetrix DMX, TimeFinder, Universal Data Tone, and VisualSRM, are trademarks of EMC Corporation. All other trademarks used herein are the property of their respective owners.

<b>Preface</b> .....	v
<b>Chapter 1 Introduction</b>	
Emulex/Solaris Environment .....	1-2
Planning Zoning in a Fabric Environment.....	1-2
<b>Chapter 2 Installing and Configuring the HBA and Driver</b>	
Installing the Hardware .....	2-2
Installing Optical GBICs for CLARiiON .....	2-2
Installing the HBA .....	2-3
Connecting Cables in a CLARiiON Environment .....	2-4
Installing the HBA I/O Driver.....	2-5
Removing CLARiiON Failover Software.....	2-5
Installing a New Driver .....	2-5
Upgrading a 4.x SBus Driver to 5.x.....	2-8
Upgrading a 4.x PCI or 5.x Driver to 5.x .....	2-12
Incorporating EMC Fibre Channel.....	2-16
Fabric Environment Specifics.....	2-16
Editing lpfc.conf.....	2-19
Adding LUNS .....	2-26
Editing /etc/system .....	2-29
Partitioning and Labeling New Devices.....	2-31
Updating the HBA Firmware.....	2-32
Updating the HBA FCode .....	2-33
Creating a Single-HBA Zone in a CLARiiON Environment ...	2-35
Creating a DS-xB or ED-12000B Switch Zone for Each HBA .....	2-35
Replacing HBAs Connected to Shared CLARiiON Systems ...	2-37

Optimizing the HBA Driver in a CLARiiON Environment ....	2-38
FC4500, FC5300, and FC5700 Storage Systems: FC-AL ....	2-38
FC4500, FC4700, and CX-Series Storage Systems: FC-SW	2-39
C3400/3500 Storage Systems .....	2-40
Symmetrix/Fabric Case Study .....	2-41

## Chapter 3    **Configuring a Boot Device on the Storage Array**

Configuring a Symmetrix Boot Device.....	3-2
Requirements .....	3-2
Procedure.....	3-2
Verifying Symmetrix Information.....	3-3
Configuring Partitions.....	3-3
Creating Filesystems .....	3-4
Installing a Bootblk.....	3-4
Copying Required Files .....	3-5
Modifying OpenBoot .....	3-6
Troubleshooting .....	3-9
Useful Monitor Commands .....	3-10
Configuring a CLARiiON Boot Device .....	3-11
Prerequisites .....	3-11
Boot Disk Configuration.....	3-11
Boot Setup.....	3-13
Troubleshooting Suggestions.....	3-16
Helpful ok Prompt Commands.....	3-16
Booting From CLARiiON Over PowerPath .....	3-17

<b>Index</b> .....	i-1
--------------------	-----

---

## Conventions Used in This Guide

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

---

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

---



### **CAUTION**

---

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

---



### **WARNING**

---

*A warning contains information essential to avoid a hazard that can cause severe personal injury, death, or substantial property damage if you ignore the warning.*

---

## Typographical Conventions

EMC uses the following type style conventions in this guide:

AVANT GARDE	Keystrokes
<b>Palatino, bold</b>	<ul style="list-style-type: none"> <li>◆ Dialog box, button, icon, and menu items in text</li> <li>◆ Selections you can make from the user interface, including buttons, icons, options, and field names</li> </ul>
<i>Palatino, italic</i>	<ul style="list-style-type: none"> <li>◆ New terms or unique word usage in text</li> <li>◆ Command line arguments when used in text</li> <li>◆ Book titles</li> </ul>
<i>Courier, italic</i>	Arguments used in examples of command line syntax.
Courier	<p>System prompts and displays and specific filenames or complete paths. For example:</p> <pre>working root directory [/user/emc]: c:\Program Files\EMC\Symapi\db</pre>
<b>Courier, bold</b>	<ul style="list-style-type: none"> <li>◆ User entry. For example: <b>sympoll -p</b></li> <li>◆ Options in command line syntax</li> </ul>

# Introduction

This document describes the procedures for installing one or more EMC-qualified Emulex LightPulse host bus adapters (HBAs) into a Sun or Fujitsu Siemens Solaris host and configuring the host for connection to an EMC storage array over Fibre Channel.

Review the *EMC Support Matrix* for the latest information on approved HBAs and drivers.

- ◆ Emulex/Solaris Environment .....1-2

## Emulex/Solaris Environment

Using the Emulex adapter with the Solaris operating system requires HBA I/O driver software. The driver functions as the host adapter driver in the host's Common SCSI Architecture (CSA), which is a layer below the Solaris SCSI Target Driver (sd), to present the EMC Fibre Channel devices to the operating system as if they were standard SCSI devices.

An Emulex HBA is identified in the Solaris host by the **lpfcX**, where **X** is the driver instance number of the HBA. This information appears in console messages (execute **dmesg**), and can be viewed in the file **/var/adm/messages**.

The instance number of the **lpfc** can exceed the number of the adapters. An administrator can determine the mapping between the physical card with the driver instance **X** by disconnecting the cable from the HBA and watching the console, which will display a message similar to the following example:

```
NOTICE:lpfcX: ...WWPN:10:00:00:00:c9:YY:YY:YY WWNN:10:00:00:00:c9:YY:YY:YY
```

where **lpfcX** is the interface of the specific Emulex adapter.

The EMC<sup>®</sup> Symmetrix<sup>®</sup> Fibre Channel implementation does not currently support the mixing of different types of HBA in a single server.

All EMC-approved Emulex HBAs use the same driver in a SPARC host.

## Planning Zoning in a Fabric Environment

Before setting up the hardware in a fabric switch configuration with Symmetrix, you should plan an effective zone map. Check the switch manufacturer's user documentation for help on defining zones.

---

This chapter describes the procedures for installing and configuring the Emulex HBA and driver.

- ◆ Installing the Hardware .....2-2
- ◆ Installing the HBA I/O Driver.....2-5
- ◆ Incorporating EMC Fibre Channel .....2-16
- ◆ Partitioning and Labeling New Devices.....2-31
- ◆ Updating the HBA Firmware .....2-32
- ◆ Updating the HBA FCode .....2-33
- ◆ Creating a Single-HBA Zone in a CLARiiON Environment .....2-35
- ◆ Replacing HBAs Connected to Shared CLARiiON Systems .....2-37
- ◆ Optimizing the HBA Driver in a CLARiiON Environment .....2-38
- ◆ Symmetrix/Fabric Case Study.....2-41

## Installing the Hardware

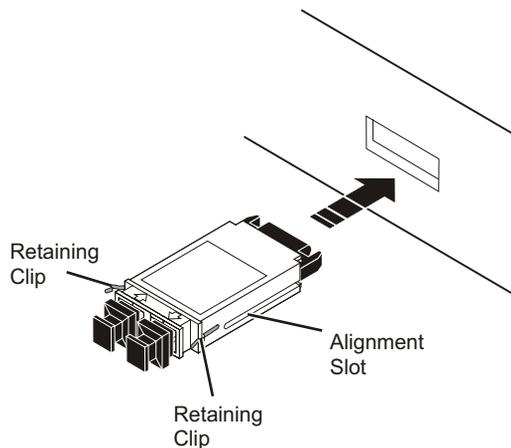
This section describes the procedure for configuring and installing the HBA. Also included are procedures for installing a GBIC and connecting the cables in a CLARiiON® environment.

### Installing Optical GBICs for CLARiiON

In a CLARiiON environment, a 1 Gb HBA requires an optical GBIC (gigabit interface converter). Follow the steps in this section to install an optical GBIC onto an HBA:

For 2 Gb HBAs, go to *Installing the HBA* on page 2-3.

1. Position the optical GBIC connector with the alignment slot toward the right.
2. Push the optical GBIC connector into the mating connector on the HBA until it is firmly seated and the retaining clips are engaged.



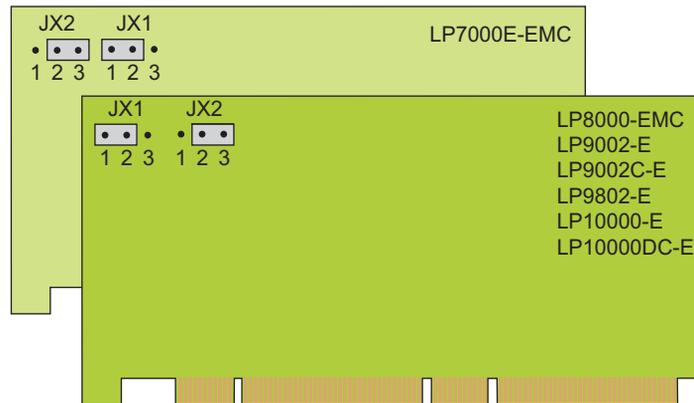
The hardware might be rotated 90 degrees clockwise from the orientation shown.

## Installing the HBA

For CLARiON environments, review your server documentation for slot recommendations that meet or exceed the HBA capabilities: 64-bit, 33 MHz. Also make sure a GBIC is installed on the HBA.

1. Set the HBA jumpers as shown below to enable it for use with the Solaris host.

The LP9002S-E and LP9002DC-E have no user-programmable jumpers.



2. With host system power removed, install the HBA as instructed in the server documentation. The card installs into a single PCI bus slot.

For a CLARiON environment, connect the cables as described under *Connecting Cables in a CLARiON Environment* on page 2-4.

3. Re-apply power and allow the system to boot normally.

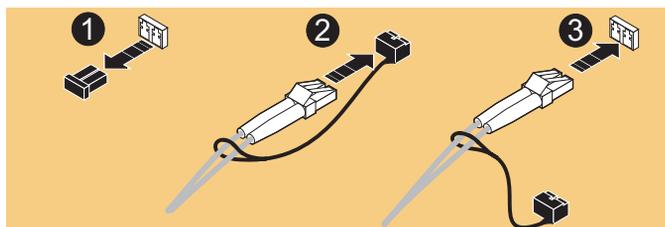
## Connecting Cables in a CLARiiON Environment

This section describes how to connect optical cables to an optical HBA connected to a CLARiiON storage array.

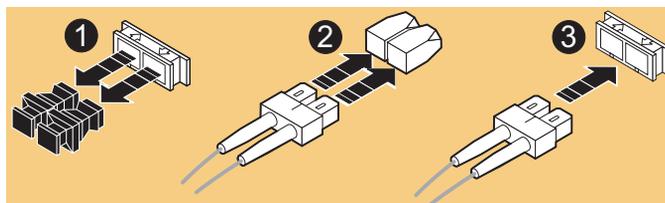
Connect optical cables to the HBA as follows:

1. Remove the protective covers on each fiber-optic cable.
2. Plug one end of the cable into the connector on the HBA as shown in the appropriate figure. (The hardware might be rotated 90 degrees clockwise from the orientation shown.)

- LC Optical Cable:



- SC Optical Cable:



3. Plug the other end of the cable into either an SP connector on the storage system or a switch port.
4. Label each cable to identify the HBA and the switch port to which it connects.
5. After connecting all HBAs in the server, power up the server.

## Installing the HBA I/O Driver

Using the Emulex adapter with the Solaris operating system requires HBA I/O driver software. The driver functions at a layer below the Solaris SCSI driver to present Fibre Channel devices to the operating system as if they were standard SCSI devices.

### Removing CLARiiON Failover Software

If CLARiiON failover software (ATF/CDE) is installed in the host, you must remove it before installing the HBA driver. After you have installed the HBA driver, you may reinstall the failover software.

Be sure you have the latest qualified driver software. If necessary, you can download the driver from the Emulex website (as described later in this procedure). Refer to release notes provided with the driver for information that might be unique to new driver revisions.

The next step depends on the type of installation, and the existing driver if you are upgrading to the latest. The v4.x driver was different for SBus HBAs than for PCI HBAs, but the v5.x driver is the same for both:

Type of Installation	Current Driver	Go to:
New	N/A	<i>Installing a New Driver</i> on page 2-5
Upgrade	v4.x SBus	<i>Upgrading a 4.x SBus Driver to 5.x</i> on page 2-8
	V4.x PCI or v5.x	<i>Upgrading a 4.x PCI or 5.x Driver to 5.x</i> on page 2-12

### Installing a New Driver

Follow these steps to perform a first-time install.

#### Obtain the Latest Driver

#### From the Emulex Website:

1. Use a Web browser to access this link:  
<http://www.emulex.com/ts/fc/docs/framemc.htm>
2. Click the link to your HBA model.
3. Under **Drivers for Solaris**, find the the correct driver version (as shown in the *EMC Support Matrix*) for your HBA, and click **Installation and Configuration** in the **Online Manuals** column.
4. Proceed to *Install the Driver* on page 2-6.

**From the Emulex CD-ROM:**

1. Insert the Emulex CD-ROM.
2. Make an `/emulex` directory:  

```
mkdir /emulex
```
3. Change to the driver directory:  

```
cd /cdrom/emulex/solaris_sparc/driver
```
4. Copy the Driver and Applications kit to the host:  

```
cp Solaris-xxx.tar /emulex
```

  
where `xxx` is the kit revision.
5. Proceed to *Install the Driver*.

**Install the Driver**

Once you have the latest driver:

1. Change to the `/emulex` directory (to which you copied the Driver and Applications kit) and untar the kit:

```
cd /emulex
tar xvf Solaris-xxx.tar
```

When the untar is complete you will find these components:

- `lpfc-yyy-sparc.tar` (where `yyy` is the driver revision) — Emulex Fibre Channel Solaris driver
- `EmlxAppsZZZ.tar` (where `ZZZ` is the HBA anywhere version) — Emulex Applications kit (containing: **lputil**, **HBA API**, and **HBA anywhere**)
- `readme.first.txt`

2. Install the driver:

```
tar xvf lpfc-yyy-sparc.tar
pkgadd -d .
```

3. You are prompted:

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

- To accept the default value, press ENTER
- If you need help, type `?` or `??` and press ENTER.

4. You are prompted:

```
Rebuild manual pages database for section 7d [y,n,?]
```

Read the following note, then press either Y or N.

---

Section 7d of the online **man** pages contains the manual pages for **lpfc**. Normally the **catman** command can be run to create preformatted versions of the online manual from the **nroff(1)** input files. Each manual page is examined and those whose preformatted versions are missing or out of date are recreated. If any changes are made, **catman** recreates the windex database. Depending on your system, this operation can take from 1 to 10 minutes. If you enter N, the **catman** will be skipped, allowing the installation to complete quickly. The skipped **catman** command will be output for the administrator to run later, if desired.

---

5. You are prompted:

```
Use IP networking over Fibre Channel [y,n,?]:
```

Press N to disable IP networking through the Emulex card.

6. You are prompted:

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

Press Y to confirm the installation.

7. The installation package continues the installation. When it finishes, you are prompted:

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

Press Q to quit the installation package. The following appears:

```
***IMPORTANT NOTICE ***
This machine must now be rebooted in order to ensure
safe operation.
```

Do not reboot yet. (You will reboot at the end of this procedure.)

8. Install the Emulex Applications Kit:

```
tar xvf EmlxAppsZZZ.tar
pkgadd -d .
```

9. You are prompted:

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

Select **HBAnywhere** and press ENTER.

10. You are prompted:

```
Do you want to continue with the installation of
<HBAnywhere> [yn,?]
```

Press Y to confirm the installation.

11. When the installation finishes, you are prompted:

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

Press Q to quit the installation package.



### CAUTION

**EMC does not support the Emulex HBAnywhere GUI.**

12. Edit any necessary files as described under *Incorporating EMC Fibre Channel* on page 2-16.
13. Reboot the host with the **-r** (reconfigure) option:

```
reboot -- -r
```

---

## Upgrading a 4.x SBus Driver to 5.x

### Back Up Files and Remove the Old Driver

If you are upgrading from a v4.x SBus driver to the latest v5.x (SBus and PCI) driver, follow these steps.

1. Back up the configuration file:

```
cp /kernel/drv/lpfs.conf /kernel/drv/lpfs.conf.bak
```

2. Back up the `sd.conf` file:

```
cp /kernel/drv/sd.conf /kernel/drv/sd.conf.bak
```

3. Back up the `path_to_inst` file:

```
cp /etc/path_to_inst /etc/path_to_inst.bak
```

4. Remove the 4.x driver:

```
pkgrm lpfs
```

5. Restore the `path_to_inst` file:

```
cp /etc/path_to_inst.bak /etc/path_to_inst
```

6. Proceed to *Obtain the Latest Driver* on page 2-9.

## Obtain the Latest Driver

### From the Emulex Website:

1. Use a Web browser to access this link:  
`http://www.emulex.com/ts/fc/docs/framemc.htm`
2. Click the link to your HBA model.
3. Under **Drivers for Solaris**, find the the correct driver version (as shown in the *EMC Support Matrix*) for your HBA, and click **Installation and Configuration** in the **Online Manuals** column.
4. Proceed to *Install the Driver* on page 2-9.

### From the Emulex CD-ROM:

1. Insert the Emulex CD-ROM.
2. Make an `/emulex` directory:  

```
mkdir /emulex
```
3. Change to the driver directory:  

```
cd /cdrom/emulex/solaris_sparc/driver
```
4. Copy the Driver and Applications kit to the host:  

```
cp Solaris-xxx.tar /emulex
```

where `xxx` is the kit revision.
5. Proceed to *Install the Driver*.

## Install the Driver

Once you have the latest driver:

1. Change to the `/emulex` directory (to which you copied the Driver and Applications kit) and untar the kit:

```
cd /emulex
tar xvf Solaris-xxx.tar
```

When the untar is complete you will find these components:

- `lpfc-yy-sparc.tar` (where `yy` is the driver revision) — Emulex Fibre Channel Solaris driver
- `EmlxAppsZZZ.tar` (where `ZZZ` is the HBA anywhere version) — Emulex Applications kit (containing: **lputil**, **HBA API**, and **HBA anywhere**)
- `readme.first.txt`

2. Install the driver:

```
tar xvf lpfc-yyy-sparc.tar
pkgadd -d .
```

3. You are prompted:

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

- To accept the default value, press ENTER
- If you need help, type ? or ?? and press ENTER.

4. You are prompted:

```
Rebuild manual pages database for section 7d [y,n,?]:
```

Read the following note, then press either Y or N.

---

Section 7d of the online **man** pages contains the manual pages for **lpfc**. Normally the **catman** command can be run to create preformatted versions of the online manual from the **nroff(1)** input files. Each manual page is examined and those whose preformatted versions are missing or out of date are recreated. If any changes are made, **catman** recreates the windex database. Depending on your system, this operation can take from 1 to 10 minutes. If you enter N, the **catman** will be skipped, allowing the installation to complete quickly. The skipped **catman** command will be output for the administrator to run later, if desired.

---

5. You are prompted:

```
Use IP networking over Fibre Channel [y,n,?]:
```

Press N to disable IP networking through the Emulex card.

6. You are prompted:

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

Press Y to confirm the installation.

7. The installation package continues the installation. When it finishes, you are prompted:

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

Press Q to quit the installation package. The following appears:

```
***IMPORTANT NOTICE ***
This machine must now be rebooted in order to ensure
safe operation.
```

Do not reboot yet. (You will reboot at the end of this procedure.)

8. Install the Emulex Applications Kit:

```
tar xvf EmlxAppsZZZ.tar
pkgadd -d .
```

9. You are prompted:

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

Select **HBAnywhere** and press ENTER.

10. You are prompted:

```
Do you want to continue with the installation of
<HBAnywhere> [yn,?]
```

Press Y to confirm the installation.

11. When the installation finishes, you are prompted:

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

Press Q to quit the installation package.



### **CAUTION**

**EMC does not support the Emulex HBAnywhere GUI.**

12. Restore any customized parameter settings in `sd.conf` and `lpfc.conf`:
  - a. Copy `/kernel/drv/sd.conf.bak` to `/kernel/drv/sd.conf`.
  - b. Merge `/kernel/drv/lpfs.conf.bak` with `/kernel/drv/lpfc.conf`.
  - c. Rename all occurrences of **lpfs** to **lpfc** in both `/kernel/drv/sd.conf` and `/kernel/drv/lpfc.conf`.
13. Reboot the host with the **-r** (reconfigure) option:

```
reboot -- -r
```

## Upgrading a 4.x PCI or 5.x Driver to 5.x

### Back Up Files and Remove the Old Driver

If you are upgrading from a v4.x PCI driver or an older v5.x driver to the latest v5.x (SBus and PCI) driver, follow these steps.

1. Back up the configuration file:
 

```
cp /kernel/drv/lpfc.conf /kernel/drv/lpfc.conf.bak
```
2. The next step depends on the current installed driver version:
  - If the version is 5.01b or older, go to step 3.
  - If the version is later than 5.01b, remove the HBAnywhere application:
 

```
pkgrm HBAnywhere
```
3. Remove the old driver:
 

```
pkgrm lpfc
```
4. Proceed to *Obtain the Latest Driver*.

### Obtain the Latest Driver

#### From the Emulex Website:

1. Use a Web browser to access this link:
 

```
http://www.emulex.com/ts/fc/docs/framemc.htm
```
2. Click the link to your HBA model.
3. Under **Drivers for Solaris**, find the the correct driver version (as shown in the *EMC Support Matrix*) for your HBA, and click **Installation and Configuration** in the **Online Manuals** column.
4. Proceed to *Install the Driver*.

#### From the Emulex CD-ROM:

1. Insert the Emulex CD-ROM.
2. Make an /emulex directory:
 

```
mkdir /emulex
```
3. Change to the driver directory:
 

```
cd /cdrom/emulex/solaris_sparc/driver
```
4. Copy the Driver and Applications kit to the host:
 

```
cp Solaris-xxx.tar /emulex
```

where **xxx** is the kit revision.

5. Proceed to *Install the Driver* on page 2-13.

## Install the Driver

Once you have the latest driver:

1. Change to the `/emulex` directory (to which you copied the Driver and Applications kit) and untar the kit:

```
cd /emulex
tar xvf Solaris-xxx.tar
```

When the untar is complete you will find these components:

- `lpfc-yyy-sparc.tar` (where `yyy` is the driver revision) — Emulex Fibre Channel Solaris driver
  - `EmlxAppsZZZ.tar` (where `ZZZ` is the HBAAnywhere version) — Emulex Applications kit (containing: **lputil**, **HBAAPI**, and **HBAAnywhere**)
  - `readme.first.txt`
2. Install the driver:

```
tar xvf lpfc-yyy-sparc.tar
pkgadd -d .
```

3. You are prompted:

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

- To accept the default value, press ENTER
  - If you need help, type `?` or `??` and press ENTER.
4. You are prompted:

```
Rebuild manual pages database for section 7d [y,n,?]:
```

Read the following note, then press either Y or N.

---

Section 7d of the online **man** pages contains the manual pages for **lpfc**. Normally the **catman** command can be run to create preformatted versions of the online manual from the **nroff(1)** input files. Each manual page is examined and those whose preformatted versions are missing or out of date are recreated. If any changes are made, **catman** recreates the windex database. Depending on your system, this operation can take from 1 to 10 minutes. If you enter N, the **catman** will be skipped, allowing the installation to complete quickly. The skipped **catman** command will be output for the administrator to run later, if desired.

---

5. You are prompted:

Use IP networking over Fibre Channel [y,n,?]:

Press N to disable IP networking through the Emulex card.

6. You are prompted:

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

Press Y to confirm the installation.

7. The installation package continues the installation. When it finishes, you are prompted:

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

Press Q to quit the installation package. The following appears:

```
***IMPORTANT NOTICE ***
This machine must now be rebooted in order to ensure
safe operation.
```

Do not reboot yet. (You will reboot at the end of this procedure.)

8. Install the Emulex Applications Kit:

```
tar xvf EmlxAppsZZZ.tar
pkgadd -d .
```

9. You are prompted:

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

Select **HBAnywhere** and press ENTER.

10. You are prompted:

Do you want to continue with the installation of <HBAnywhere> [yn,?]

Press Y to confirm the installation.

---

EMC does not support the Emulex HBAnywhere GUI.

---

11. When the installation finishes, you are prompted:

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

Press **Q** to quit the installation package.

12. Restore any customized parameter settings in  
`/kernel/drv/lpfc.conf.bak` to `/kernel/drv/lpfc.conf`.
13. Reboot the host with the **-r** (reconfigure) option:

```
reboot -- -r
```

## Incorporating EMC Fibre Channel

Once the EMC storage array has devices assigned with device addresses, the host can see all the target EMC devices assigned to that interface. Devices are presented to the host in the same manner as devices accessed through a standard SCSI interface.

To insure smooth integration and maximum performance from the host and the EMC storage array, you must edit these files (as described in detail later):

- ◆ Emulex driver configuration file `lpfc.conf` — Refer to *Editing lpfc.conf* on page 2-19.
- ◆ Host sd driver configuration file `sd.conf` — Refer to *Adding LUNS* on page 2-26.
- ◆ System specification file `etc/system` — Refer to *Editing /etc/system* on page 2-29.

After editing the files, type `reboot -- -r` and press ENTER. This reboots the host with the `-r` (reconfigure) option, which rebuilds the kernel and implements the changes.

---

### Fabric Environment Specifics

The Emulex HBA and the EMC storage array act as separate nodes that participate in the same switched fabric environment. Therefore, prior to the setup of the Emulex HBA driver file (`lpfc.conf`), the administrator must enable the storage array to participate in the fabric. The administrator must also configure the switch to enable the EMC storage array and the hosts to "see" each other. (Refer to the user manual for the switch.)

### Persistent Binding Implementation

Emulex persistent binding is implemented through both `/kernel/drv/lpfc.conf` and `/kernel/drv/sd.conf`. The matching between the EMC storage array port, host adapter instance, and target ID occurs inside `lpfc.conf`. On the other hand, the matching between the target ID, host adapter instance, and LUN occurs inside `sd.conf`. All of these configurations assume that the settings for the persistent binding in `lpfc.conf` is enabled (for example: `automap=2`) and the EMC storage array is being configured in the physical addressing mode (V-bit is off).



### CAUTION

**The modification of the default entries (target 0 to target 15) inside `/kernel/drv/sd.conf` might cause the host machine fail to boot properly because the Solaris operating system uses targets 0 to 15 as the reference or the bind to the internal boot drive.**

Inside `lpfc.conf`, the EMC storage array port (target) is mapped against a target number that is defined as one of the entries inside `sd.conf`. The mapping is done by defining the `lpfcXtY` properties, where the `X` corresponds to the driver instance number, and the `Y` is the target ID. Note that this target ID does not correspond to the loop ID defined by the Fibre Channel `AL_PA` (as in an arbitrated loop configuration); however, it is a logical definition of the relationship between the EMC storage array port ID and the device address.

Although the current Emulex HBA driver enables three different binding methods (World Wide Node Name, World Wide Port Name, Port ID), the recommended setup uses the World Wide Port Name (WWPN).

On a Symmetrix, the WWPN can be obtained by issuing the Inlines command `e2` for the designated Fibre Channel director port.

The format for a bind entry is `NNNNNNNNNNNNNNNNNN:lpfcXtY`, where:

- ◆ `NNNNNNNNNNNNNNNNNN` is a 16-digit BCD representation of the WWPN of the EMC storage array port.

The port's WWPN and WWNN (World Wide Node Name) are the same. On a Symmetrix, this can be displayed using the Inlines `e2` command. The output is similar to the following example for Fibre Channel port 1A, with a WWPN of 50080482b8912b8e:

```
XWWN HI      50080482 / 50080482
WWN  LOW     b8912b8e / b8912b8e
              WWNN      WWPN
```

- ◆ `X` is the driver instance number.

You can determine the driver instance number `X` or `lpfcX` by using a Fibre Channel loopback connector to identify the `lpfcX` number that corresponds to the path to the Fibre Channel port.

- ◆ `Y` is the target assignment.

Multiple entries must be separated by a comma (,) with the last entry terminated by a semi-colon (;). Target assignments, with all supported LUNs must also be configured in `sd.conf`, `st.conf`, or `cmdk.conf`.

*Sample entry:*

```
fcp-bind-WWPN= "50080482b8912b8e:lpfc0t0",
               "50080482b8912b9e:lpfc0t1";
```

The `class="lpfc"` and `class="scsi"` entries inside `/kernel/drv/sd.conf` cause the SCSI layer to probe multiple adapters, even multiple adapters across multiple drivers (for example the **glm**, **isp**, and **fas**). Emulex implements a method to enable the administrator to individually bind each LUN (which corresponds to a single device in the EMC storage array) to a specific HBA. This is accomplished by adding `hba="lpfcX"` (where **X** is the driver instance number) to the `/kernel/drv/sd.conf` properties.

*LUN binding example:*

```
NOTICE: lpfc0: ...WWPN:10:00:00:00:c9:YY:YY:YY WWNN:10:00:00:00:c9:YY:YY:YY
        name="sd" parent="lpfc" target=16 lun=0 hba="lpfc0";
```



### CAUTION

**This entry will not cancel the effect of any other `class="scsi"` entries for `target=16 lun=0`. If the user wants the SCSI layer to only probe for `target=16` and `lun=0` on device `lpfc0`, the `class="scsi"` entries for `target=16` and `lun=0` must be deleted.**

### Using `lputil` to Set Up Persistent Binding

The Emulex driver comes with a utility called `lputil` to set up persistent binding (among other functions). To set up binding, follow these steps, repeating them for each adapter:

1. Type the following and press ENTER:

```
/usr/sbin/lpfc/lputil
```

This displays a menu.

2. Select **5. - Persistent Bindings**.
3. Select **5. - Bind Automapped Targets**.

4. Select an adapter, **0–n**, where **n** is the number of adapters installed in your system.
5. For Bind all auto-mapped targets? (Y/N) = >, select **Y** (yes).
6. Select **2. - Bind by port name**.

## Editing lpfc.conf

Configuration file `lpfc.conf` is created automatically by `pkgadd`. The file contains important information about how the driver should behave. If necessary, edit the file as shown in the following table. Descriptions of the variables follow the table.

Changes to the configuration file will not take effect until the host is rebooted.

Variable	Default Value	Recommended Setting		
		If No PowerPath or Veritas DMP	PowerPath or Veritas DMP <sup>a</sup>	If CLARiiON with ATF <sup>b</sup>
log-verbose	0	0x10	0x10	0
log-only	1	1	1	1
fcplib-WWNN	none	16-digit hex BCD with leading zeroes <sup>c</sup>		none
fcplib-WWPN	none	16-digit hex BCD with leading zeroes <sup>c</sup>		
fcplib-DID	none	6-digit hex BCD with leading zeroes <sup>c</sup>		none
automap	2	2	2	2
fcplib-on	1	1	1	1
lun-queue-depth	30	20	20	20
tgt-queue-depth	0	Symmetrix: 512 CLARiiON FC4500: 256 <sup>d</sup> CLARiiON FC5300: 256 <sup>d</sup> CLARiiON FC4700: 512 <sup>d</sup> CLARiiON CX: 512 <sup>d</sup>	Symmetrix: 512 CLARiiON FC4500: 256 <sup>d</sup> CLARiiON FC5300: 256 <sup>d</sup> CLARiiON FC4700: 512 <sup>d</sup> CLARiiON CX: 512 <sup>d</sup>	CLARiiON FC4500: 256 <sup>d</sup> CLARiiON FC5300: 256 <sup>d</sup> CLARiiON FC4700: 512 <sup>d</sup> CLARiiON CX: 512 <sup>d</sup>
lpfcNtM-lun-throttle	N/A	N/A	N/A	N/A
lpfcNtM-target-throttle	N/A	N/A	N/A	N/A
no-device-delay	1	1	0	0
network-on	0 or 1 <sup>e</sup>	0 or 1 <sup>e</sup>	0 or 1 <sup>e</sup>	0 or 1 <sup>e</sup>

Variable	Default Value	Recommended Setting		
		If No PowerPath or Veritas DMP	PowerPath or Veritas DMP <sup>a</sup>	If CLARiiON with ATF <sup>b</sup>
xmt-que-size	256	256	256	N/A
scan-down	2	FC-AL: <ul style="list-style-type: none"> <li>0 = FCP persistent binding</li> <li>2 = No FCP persistent binding</li> </ul> FC-SW: 0	FC-AL: <ul style="list-style-type: none"> <li>0 = FCP persistent binding</li> <li>2 = No FCP persistent binding</li> </ul> FC-SW: 0	FC-AL: <ul style="list-style-type: none"> <li>0 = FCP persistent binding</li> <li>2 = No FCP persistent binding</li> </ul> FC-SW: 0
linkdown-tmo	30	0	60	30
nodev-holdio	0	1	0	0
nodev-tmo	0	0	60	60
delay-rsp-err	0	0	0	0
num-iocbs	1024	2048	2048	2048
num-bufs	1024	1024	1024	1024
topology	0	FC-AL: 4 FC-SW: 2	FC-AL: 4 FC-SW: 2	FC-AL: 4 FC-SW: 2
ip-class	3	3	3	N/A
fcp-class	3	3	3	3
use-adisc	0	0	0	0
fcpfabric-tmo	0	0	0	0
post-ip-buf	128	128	128	N/A
dqfull-throttle	1	1	1	1
dqfull-throttle-up-time	1	1	1	1
dqfull-throttle-up-inc	1	1	1	1
ack0	0	0	0	0
cr-delay	0	0	0	0
cr-count	0	0	0	0

Variable	Default Value	Recommended Setting		
		If No PowerPath or Veritas DMP	PowerPath or Veritas DMP <sup>a</sup>	If CLARiiON with ATF <sup>b</sup>
link speed	0	0	0	0
check-cond-err	0	0	0	0
fdmi-on	0	0	0	0

- PowerPath v1.x and 2.x.x are supported for Symmetrix only.
- ATF and PowerPath can not coexist on the same system.
- FC-SW only. Only one binding method can be used. With CLARiiON arrays, use only WWPN binding. Refer to the descriptions of these variables for examples.
- Refer to *Optimizing the HBA Driver in a CLARiiON Environment* on page 2-38.
- Enabled (1) or disabled (0) setting depends on the setup during driver installation.

### Descriptions of lpfc.conf Variables:

**log-verbose** — Sets a bit mask to record specific types of verbose messages:

- 0x1 — ELS events
- 0x2 — Device Discovery events
- 0x4 — Mailbox Command events
- 0x8 — Miscellaneous events
- 0x10 — Link Attention events
- 0x20 — IP events
- 0x40 — FCP events
- 0x80 — Node table events

**log-only** — Setting the value to **0** causes log messages print on the console and to be logged to **syslog** (which may send them to the console again if it is configured to do so). Setting **log-only** to **1** causes log messages to go to **syslog** only.

**fcv-bind-WWNN** — Binds a specific World Wide Node Name to a target ID.

*Example:*

**WWNN:SCSI ID**

```
fcv-bind-WWNN = "5008048256789abc:lpfc0t0",
               "50080482370c27f7:lpfc0t2",
               "5008048222222222:lpfc2t2";
```

**fcp-bind-WWPN** — Binds the World Wide Port Name of a storage port to a target ID. Note that multiple entries of the WWPN can be configured for the **lpfcX**; they must be separated by commas, and the list must end with a semicolon.

*Example:*                    **WWPN:SCSI ID**

```
fcp-bind-WWPN = "5008048256789abc:lpfc0t0",
                "50080482370c2855:lpfc0t1",
                "5008048222222222:lpfc2t2";
```

**fcp-bind-DID** — Binds a specific DID to a target ID. The SCSI ID to bind to consists of two parts, the lpfc interface to bind to, and the target number for that interface. Thus lpfc0t2 specifies target 2 on interface lpfc0.

Target IDs, with all LUNs supported, must also be in `sd.conf`.

*Example:*                    **DID:SCSI ID**

```
fcp-bind-DID = "0000ef:lpfc0t3";
```

**automap:**

- **0** = Only devices with persistent binding will be recognized by the system.
- **1** = Force WWNN binding.
- **2** = Force WWPN binding.
- **3** = Force DID binding.

**fcp-on**— Is always **1** for SCSI devices in a Fibre Channel fabric or arbitrated loop environment.

**lun-queue-depth** — Number of outstanding commands per FCP LUN. This value is global, affecting each LUN recognized by the driver, but it may be overridden on a per-LUN basis. (See **lpfcNtM-lun-throttle**.) It might be a good idea to configure RAID arrays using the per-LUN tunable throttles.

**tgt-queue-depth** — Number of outstanding commands per FCP target. This value is global, affecting each target recognized by the driver, but it may be overridden on a per-target basis (see **lpfcNtM-tgt-throttle** below). For example, you might want to configure RAID arrays using the per-target tunable throttles.

**lpfcNtM-lun-throttle** — Maximum number of outstanding commands to permit for each LUN of an FCP target that supports multiple LUNs. The default throttle for the number of commands

outstanding to a single LUN of a multiple-LUN target is the **lun-queue-depth**. For a target that supports multiple LUNs, it may be useful to specify a LUN throttle that differs from the default.

*Example* — **lpfc0t17-lun-throttle=31** means that each LUN on target 17, interface lpfc0 should be allowed up to 31 simultaneous outstanding commands.

**lpfcNtM-tgt-throttle** — Maximum number of outstanding commands to permit for a FCP target. By default, target throttle is disabled.

*Example* — **lpfc0t17-tgt-throttle=256** means that target 17, interface lpfc0 should be allowed up to 256 simultaneous outstanding commands.

**no-device-delay** — Interval in seconds (0 to 30) between deciding to fail back an I/O because there is no way to communicate with its particular device (for example, due to device failure) and the actual fail back. A value of zero implies no delay.

*Cautions:*

- The system generally rounds values smaller than 10,000 microseconds (10 milliseconds) up to roughly 10 milliseconds, depending on the system clock rate.
- Setting a long delay value might permit I/O to build up, each with a pending timeout, possibly resulting in exhaustion of critical Solaris kernel resources. In this case, you might see a fatal message, such as `PANIC - Timeout table overflow`. Note that this value can have an impact on the speed with which a system can shut down with I/Os pending and with the HBA unable to communicate with the loop (if a cable is pulled, for example).

**network-on** — Is **1** if networking is enabled, **0** if disabled. This variable is set during installation of the driver via **pkgadd**.

**xmt-que-size** — Size of the transmit queue (128 to 10240) for mbufs.

**scan-down:**

- ◆ If the value is 0, the lpfc driver assigns target IDs by scanning its AL\_PA map from low AL\_PA to high AL\_PA.
- ◆ If the value is 1, the lpfc driver uses an inverted AL\_PA map, effectively scanning AL\_PAs from high to low, as specified in the FC-AL annex.
- ◆ If the value is 2, the lpfc driver uses an inverted AL\_PA map (the same as value 1), and also bases target assignment in the private loop environment on the AL\_PA.

**linkdown-tmo** — Time (1 to 255 seconds) the driver waits to return the I/O packet as an error for retry when a cable has been pulled or the link has otherwise become inaccessible. Linkdown processing includes failing back commands to the target driver that have been waiting around for the link to come back up.

---

Small values of the link-down timer cause the link to appear to "bounce", while large values of the timer can delay failover in a fault-tolerant environment. Units are in seconds. A value of 0 means never failback commands until the link comes up.

---

**nodev-holdio** — If enabled (1), the lpfc driver holds all I/O errors on devices that disappear until they come back. If disabled (0), the lpfc driver return errors with no-device-delay.

**nodev-tmo** — Time (0 to 255 seconds) I/O errors will be held by the device driver if a device on the loop disappears.

**delay-rsp-err** — Treats FCP RSP errors like **no-device-delay**.

**num-iocbs** — Number of iocb buffers to allocate (128 to 10240).

**num-bufs** — Number of ELS buffers to allocate (128 to 4096). ELS buffers are needed to support Fibre Channel Extended Link Services. ELS buffers are also used for SLI-2 FCP buffers (one per FCP command) and Mailbox commands.

**topology** — Link topology for initialized link.

**ip-class** — Fibre Channel class (2 or 3) to use for the IP protocol.

**fcp-class** — Fibre Channel class (2 or 3) to use for the FCP protocol.

**use-adisc** — If 1, ADISC is used for FCP rediscovery; if 2, PLOGI is used.

**fcpfabric-tmo** — Has no function in an arbitrated loop environment.

**post-ip-buf** — Number of 4K STREAMS buffers to post to an IP ring.

**dqfull-throttle** — Decrements (if set to 1) the LUN throttle on a queue-full condition.

**dqfull-throttle-up-time** — Time (in seconds) until the current queue depth will be incremented (by the value specified by **dqfull-throttle-up-inc**).

**dqfull-throttle-up-inc** — Size of increment of queue depth when incremented after **dqfull-throttle-up-time**.

**ack0** — Determines Class 2 acknowledgement: ACK0 (1) or ACK1 (0).

**cr-delay** — Specifies a count of milliseconds after which an interrupt response is generated if **cr-count** has not been satisfied.

**cr-count** — Specifies a count of I/O completions after which an interrupt response is generated.

**link-speed:**

0 = auto

1 = 1 GB/sec

2 = 2 GB/sec

**fdmi-on:**

0 = FCP access is disabled.

1 = FCP access is enabled.

### Determining the Driver Instance Number (x) in a CLARiiON Environment

If the HBA driver is managing HBAs for other heterogeneous storage devices, you may wish to have the variables apply to a specific Emulex adapter. This can be done by prefixing the configuration file variables with HBA adapter instance number.

To determine the driver instance number, as represented by *x* when appending configuration file variables, do the following:

1. List the devices in the `/dev/rdisk` directory — select only slice `s2`. Refer to the following example:

```
atf-sun-2 [5.7]: ls /dev/rdisk/*s2
/dev/rdisk/c0t0d0s2 /dev/rdisk/c1t0d1s2 /dev/rdisk/c2t0d0s2 /dev/rdisk/c2t0d3s2
/dev/rdisk/c0t6d0s2 /dev/rdisk/c1t0d2s2 /dev/rdisk/c2t0d1s2
/dev/rdisk/c1t0d0s2 /dev/rdisk/c1t0d3s2 /dev/rdisk/c2t0d2s2
```

2. Select one LUN per controller to find the path for that controller (or adapter), as shown in the following example:

```
atf-sun-2 [5.7]: ls -al /dev/rdisk/c1t0d3s2

lrwxrwxrwx  1 root  root          54 Sep 27 14:58
/dev/rdisk/c1t0d3s2 ->
../../../../devices/pci@1f,4000/fibre-channel@2/sd@0,3:c,raw
```

---

The path could contain **fibre-channel** or **lpfc**.

---

3. Use the underlined path in the example above to find the related instance number in the `/etc/path_to_inst` directory for this adapter.

```
atf-sun-2 [5.7]: grep "pci@1f,4000/fibre-channel@2"
/etc/path_to_inst | grep "\"lpfc\""

"/pci@1f,4000/fibre-channel@2" 2  "lpfc"
```

This finds the instance number (underlined in the example above) used in `/kernel/drv/lpfc.conf` and used to set the variable attributes for that adapter:

**lpfc2-scan-down=2;**

---

On occasion multiple entries exist for a particular path in the `path_to_inst` file, making it difficult to find the actual instance number being used. However, executing the command `prtconf > /tmp/output` will show all the adapters configured by the system. Looking through the output for the driver name **fibre-channel** (or **lpfc**) will show if the instance number found in step 3 is actually configured on the system.

---



---

## Adding LUNS

EMC storage array devices are addressed using target IDs and LUNs. The Solaris configuration file `/kernel/drv/sd.conf` lists all disks (LUNs) that a user may want to access. The default file contains entries for LUN 0 only; therefore, you must edit the file to add entries for any LUNs with numbers greater than 0 that you will bind in the storage systems connected to the server.

Devices accessed via Fibre Channel using the Emulex adapter and Emulex driver are addressed behind a single target ID with multiple LUNS (up to 32 with CLARiiON and up to 256 with Symmetrix).

Before editing `/kernel/drv/sd.conf`, make a copy of it, for backup if the file is lost or errors are made. When editing the file, follow these precautions:

- ◆ Do not delete any original data from this file; the Solaris operating system might not boot properly.
- ◆ Do not insert hidden format characters. The file is very sensitive to hidden format characters, and a misplaced tab or other formatting character can make all subsequent entries unreadable.
- ◆ Enter all LUN numbers in decimal format.

Changes to `sd.conf` do not take effect until the host is rebooted.

### FC-AL Specifics

The target ID needed to address a device is determined by the Fibre Channel Arbitrated Loop Physical Address (AL\_PA). A loop ID of 00 (AL\_PA of EF) will have LUNS behind target ID 0, loop ID 01 has LUNs on target ID 1, and so on.

Since each Fibre Channel port can have only one loop ID (or AL\_PA), all devices accessed through a port will have the same target ID. If you need to access existing devices that previously had multiple target IDs, you must be aware of changes required for raw device addressing.

In an arbitrated loop configuration, with `scandown=2`, SCSI IDs are hard-coded based on the Loop ID. This parameter maintains backward compatibility with EMC storage Loop ID assignment. For example, if a Symmetrix FA has a Loop ID of 32 (0x20), the `sd.conf` target is also 32.

### FC-SW Specifics

Modify `/kernel/drv/sd.conf` to have the following definitions for each target ID that needs to support multiple LUNs:

```
name="sd" parent="lpfc" target="x" lun="y" hba="lpfcz";
```

where:

- **x** is the target number that corresponds to the `fc_bind_WWN ID lpfcZtX`.
- **y** is the LUN number that corresponds to Symmetrix volume mapping on the Symmetrix port WWN.
- **z** is the lpfc driver instance number that corresponds to the `fc_bind_WWN ID lpfcZtX`.

## Sample `sd.conf` for PCI HBA

Your actual target/LUN assignment will vary depending on unique system configuration requirements.

Make sure that only the **target=x lun=0** line is specified for any target ID that does not need LUN support. This decreases the time needed to reboot the system.

The following sample `sd.conf` file illustrates how to add LUNs, connected to a PCI Fibre Channel HBA, to `sd.conf`. The lines in bold are added to the file. This example adds LUNs to SP A (target 0) and SP B (target 1) of a CLARiiON storage system. You will add LUNs under the entries for the **lpfc** parent.

```
name="sd" parent="lpfc" target=0 lun=0; hba="lpfcX"
name="sd" parent="lpfc" target=0 lun=1; hba="lpfcX"
name="sd" parent="lpfc" target=0 lun=2; hba="lpfcX"
name="sd" parent="lpfc" target=0 lun=3; hba="lpfcX"
name="sd" parent="lpfc" target=0 lun=4; hba="lpfcX:"
name="sd" parent="lpfc" target=0 lun=5; hba="lpfcX"

name="sd" parent="lpfc" target=1 lun=0; hba="lpfcX"
name="sd" parent="lpfc" target=1 lun=1; hba="lpfcX"
name="sd" parent="lpfc" target=1 lun=2; hba="lpfcX"
name="sd" parent="lpfc" target=1 lun=3; hba="lpfcX"
name="sd" parent="lpfc" target=1 lun=4; hba="lpfcX"

name="sd" parent="lpfc" target=125 lun=0; hba="lpfcX"
```

The entries for Emulex use the **parent="lpfc"** keyword instead of **class="scsi"** to ensure that Solaris specifically probes all adapters controlled by the `lpfc` driver for the specified targets and LUNs (instead of probing all SCSI adapters in the entire system).

The **hba="lpfcX"** parameter affects whether a specific target associated with the `lpfc` instance in the `lpfc.conf` file will be used in association with a Emulex adapter. Using **"lpfc0"** means that the first instance of the Emulex HBA will used a particular target, while any other instance of a Emulex HBA will not.

After you add all of the LUNs that you think you will need, save the file and reboot the operating system with the **-r** (reconfigure) option:

```
reboot -- -r
```

## Editing /etc/system

Operation in a Solaris environment requires that some parameters in the file `/etc/system` first be modified, as described below.

Changes to `/etc/system` will not take effect until the host is rebooted.

The parameter `sd_max_throttle` must be modified specifically for Fibre Channel. Other settings apply to both standard SCSI and Fibre Channel.

1. Add this line: `set sd:sd_max_throttle=20`
2. Add this line: `set scsi_options = 0x7F8`

The bit meaning for `scsi_options` is as follows:

Bit	Mask	Meaning
3	0x8	Disconnect enable
4	0x10	Link enable
5	0x20	Sync xfer enable
6	0x40	Parity support enable
7	0x80	Tag commands enable
8	0x100	Fast SCSI enable
9	0x200	Wide SCSI enable
10	0x400	Ultra SCSI enable

Since `/etc/system` can also affect other disks in the system, consider the effects of implementing the `scsi_options=0x7f8` mask on the other disks. If it is not set, storage array operations will not benefit from these options.

3. Add a line to set the I/O timeout. This setting prevents the host from issuing warning messages when non-disruptive operations are performed on the EMC storage array:
  - If the host will run PowerPath, set the I/O timeout to 60 seconds:  
`set sd:sd_io_time = 0x3C`
  - If the host will not run PowerPath, EMC recommends setting the I/O timeout to 120 seconds:  
`set sd:sd_io_time = 0x78`

## Partitioning and Labeling New Devices

Use the **format** command to partition and label new Fibre Channel devices. The devices will appear under `/dev/dsk`.

To partition and label new devices:

1. At the `root` prompt, type **format** and press ENTER.

The host searches for all disks and displays a list, placing all unlabeled disks at the beginning of the list.

2. At the `Specify Disk` prompt, enter the number of the first EMC drive.
3. The `FORMAT MENU` appears. At the `format` prompt, type **label** and press ENTER.
4. At the `Disk not labeled. Label it now?` prompt, type **y** and press ENTER.
5. Type **disk** and press ENTER at the prompt to display a list of the disks.
6. Repeat steps 2 through 5 for each EMC disk.

## Updating the HBA Firmware

To update your firmware:

1. Copy the appropriate firmware file from the CD-ROM or the Emulex website to the appropriate directory, `/usr/sbin/lpfc`:

HBA Model	File
LP7000E-EMC	<b>sfXXX.awc</b> (where <b>XXX</b> is the version; for example: <b>sf330a7.awc</b> )
LP8000E-EMC	<b>ddXXX.awc</b> (where <b>XXX</b> is the version; for example: <b>dd390a7.awc</b> )
LP9002-E LP9002C-E LP9002DC-E	<b>cdXXX.awc</b> (where <b>XXX</b> is the version; for example: <b>cd390a7.awc</b> )
LP9002S-E	<b>rfXXX.awc</b> (where <b>XXX</b> is the version; for example: <b>rf390a7.awc</b> )
LP9802-E	<b>hdXXX.all</b> (where <b>XXX</b> is the version; for example: <b>hd100a2.all</b> )
LP10000-E LP10000DC-E	<b>tdXXX.all</b> (where <b>XXX</b> is the version; for example: <b>td180a2.all</b> )

2. Start the firmware update utility:
 

```
./lputil
```
3. Select option **3 – Firmware Maintenance**.
4. Select the adapter number to update.
5. Select option **1 – Load Firmware Image**.
6. Type the full file name of the image file (as shown in the table in step 1; **cd390a7.awc**, for example) and press ENTER.
7. Select option **0 – Return to Main Menu**.
8. Repeat steps 3 through 7 for any other adapters.
9. If you are upgrading the firmware on the system that will boot from the EMC storage array, update the FCode as described under *Updating the HBA FCode* on page 2-33.



### CAUTION

**If you are using the HBA for booting, update the FCode immediately after updating the firmware, because the firmware upgrade overwrites parts of the FCode.**

## Updating the HBA FCode

To update your FCode support, follow these steps:

1. Make sure you have updated the firmware (as described under *Updating the HBA Firmware* on page 2-32).
2. Copy the appropriate file for Open Boot support from the CD-ROM or the Emulex website to the appropriate directory, `/usr/sbin/lpfc:`

HBA Model	File
LP7000E-EMC	<b>SOXXX.PRG</b> (where <b>XXX</b> is the version; for example: <b>SO132A0.PRG</b> )
LP8000E-EMC	<b>DOXXX.PRG</b> (where <b>XXX</b> is the version; for example: <b>DO140A0.PRG</b> )
LP9002-E LP9002C-E LP9002DC-E	<b>COXXX.PRG</b> (where <b>XXX</b> is the version; for example: <b>CO140A0.PRG</b> )
LP9002S-E	<b>TSCXXX.PRG</b> (where <b>XXX</b> is the version; for example: <b>TSC240A0.PRG</b> )
LP9802-E	<b>HOXXX.PRG</b> (where <b>XXX</b> is the version; for example: <b>HO140A0.PRG</b> )
LP10000-E LP10000DC-E	<b>TOXXX.PRG</b> (where <b>XXX</b> is the version; for example: <b>TO140A0.PRG</b> )

3. Start the FCode update utility:

```
./lputil
```

4. Select option **3 – Firmware Maintenance**.
5. Select the adapter number to update.
6. Select the appropriate option:

HBA Bus	Select:
PCI	<b>1 – Load Firmware Image</b>
SBus	<b>6 – Load FCode Image</b>

7. Enter the full file name of the image file.
8. Select option **0 – Return to Main Menu**.
9. Repeat steps 4 through 8 for any other adapters.

10. Type `reboot -- -r` and press ENTER to reboot the host with the `-r` (reconfigure) option.



#### **CAUTION**

**Make sure your storage array device path is `/pci/lpfc`. If you see that the path is `/pci/fibre_channel`, you must download the Fcode again and reboot the system with the command `reboot -- -rv`.**

## Creating a Single-HBA Zone in a CLARiiON Environment

---

This section applies only to fabric configurations with CLARiiON storage arrays.

---

EMC requires single-initiator zoning. This means that each HBA has a separate zone that contains it and the SP with which it communicates. The following zoning configuration rules apply:

- Fan-out of 32 initiators registered to each SP port on an FC4700 storage system and 15 hosts to one non-FC4700 storage system.
- Fan-in of four storage systems to one host.

How you set up a zone on a switch depends on the type of switches you are using.

- ◆ DS-8B, DS-16B, DS-16B2, DS-32B2, ED-12000B — Proceed to *Creating a DS-xB or ED-12000B Switch Zone for Each HBA*.
- ◆ DS-16M, DS-16M2, DS-24M2, DS-32M, DS-32M2, ED-64, ED-140M, or ED-1032 — Refer to the documentation for the switch.

---

### Creating a DS-xB or ED-12000B Switch Zone for Each HBA

You use the Web Tools to zone a DS-8B, DS-16B, DS-16B2, DS-32B2, or ED-12000B. If you are unfamiliar with Web Tools or with how zoning operates on these switches, refer to the Web Tools user guides, zoning reference manuals, and Quickloop reference manuals, available at <http://powerlink.emc.com>, under Connectrix in the Document Library.

To create a switch zone for each HBA, follow these steps:

1. Start **Web Tools** and click **Zone Admin**.

The Zone Administrator window opens.

---

You may have to enter your username and password to access the Zone Administrator window.

---

2. In Web Tools, click **Zone Settings**.

The **Zone Settings** tab opens.

3. Click **Create Zone** to open the Create Zone dialog box.
4. Enter the name of the new zone.

5. For each member you want to add to the zone, select the member (HBA and SPs) by World Wide Port Name and click **Add Member**.
6. After adding all members to the zone, click **Apply**.
7. Click **Config Settings**.
8. Either create a new configuration as follows or select an existing configuration:

---

Only one zone configuration can be active at a time.

---

- a. Click **Create Cfg** to open the **Create Cfg** dialog box.
  - b. Enter the name of the new configuration.
9. For each zone you created, select the zone, then click **Add Member**.
  10. Click **Enable Config**, and then click **Apply**.
  11. Repeat steps 1 through 10 to set up the zone settings and zone configuration settings for the second HBA.
  12. Reboot the host with the **-r** (reconfigure) option:

**reboot -- -r**

---

*Important:* You *must* reboot now so the HBA can see the targets (SPs) that you configured through zoning.

---

## Replacing HBAs Connected to Shared CLARiiON Systems

---

A shared CLARiiON storage system is a storage system with Access Logix software.

---

Follow these steps to replace an HBA connected to a shared storage system:

1. Stop all I/O to the storage systems connected to the HBA you are replacing.
2. Remove the connection path to the old HBA from each storage system's persistent memory, using these CLI commands:

- a. `navicli port -list`

- b. `navicli -removeHBA`

Refer to the CLI reference manual for information on this command.

3. Shut down Solaris on the server whose HBA you are replacing.
4. Power down the server.
5. Remove the HBA you want to replace and install the new HBA.
6. Power up the server.
7. After the HBA is registered with the Host Agent running on the server, reconnect any Storage Groups that you disconnected.

The HBA is registered with the Host Agent if the Connectivity Status dialog box in Manager shows **Yes** in the **Registered** column for the HBA.

8. In any switch zones that include the HBA, change the World Wide Port Name (WWPN) of the old HBA to the WWPN of the new HBA.
9. Reboot the host with the `-r` (reconfigure) option:

```
reboot -- -r
```

## Optimizing the HBA Driver in a CLARiiON Environment

In a CLARiiON environment, if neither the CDE nor ATF option will be installed, optimize your configuration by following the appropriate procedure:

### FC4500, FC5300, and FC5700 Storage Systems: FC-AL

Follow these steps:

1. In `/kernel/drv/lpfc.conf`, set parameters as follows:
  - In a single-server configuration:
    - Set **tgt-queue-depth** to **256**.
    - Calculate **lun-queue-depth** by dividing the **tgt-queue-depth** (in this case 256) by the maximum number of LUNs bound per SP.  
*Example:* If eight LUNs are bound to each SP, you would set this variable to **32** ( $256 / 8 = 32$ ).
    - Set **nodev-tmo** to **20**.
    - Set **dqfull-throttle** to **0**.
  - In a multi-server configuration:
    - Calculate **tgt-queue-depth** by dividing 256 by the maximum number of HBAs (initiators) connected to each SP.  
*Example:* If three HBAs were connected through Fibre Channel Hubs to each SP in a storage system, you would set the value to **85** ( $256 / 3 \text{ HBAs} = 85$ ).
    - Calculate **lun-queue-depth** by dividing the **tgt-queue-depth** (in this case 85) by the maximum number of LUNs bound per SP.  
*Example:* If 10 LUNs are bound to each SP, you would set this variable to **8** ( $85 / 10 = 8$ ).
    - Set **nodev-tmo** to **20**.
    - Set **dqfull-throttle** to **0**.
2. Reboot the server (with a standard reboot) to make your changes take effect:

**Reboot**

---

## FC4500, FC4700, and CX-Series Storage Systems: FC-SW

Follow these steps:

1. In `/kernel/drv/lpfc.conf`, set parameters as follows:
  - Calculate **tgt-queue-depth** by dividing 2000 (256 for the FC4500) by the maximum number of HBAs (initiators) that could have access to Storage Groups on each storage system.

*Example:* If the fabric consists of four servers (each with two HBAs), each connected to one SP in each storage system through switches, you would set the value to **256** (2000/8).

- Calculate **lun-queue-depth** by dividing the **tgt-queue-depth** (in this case 256) by the number of LUNs contained in the largest Storage Group to be accessed from the host being configured.

*Example:* If this host has access to two storage systems on the fabric with a Storage Group on each storage system, you divide **tgt-queue-depth** (this case 256) by the number of LUNs in the largest Storage Group. If the largest Storage Group has eight LUNs, you would set this variable to **32** (256/8).

---

CDE or ATF sets other parameter values after you add persistent bindings. Refer to the *EMC Navisphere Server Software for Solaris Administrator's Guide* (P/N 069001068) for more information on these parameters and their values.

---

2. Reboot the server (with a standard reboot) to make your changes take effect:

**Reboot**

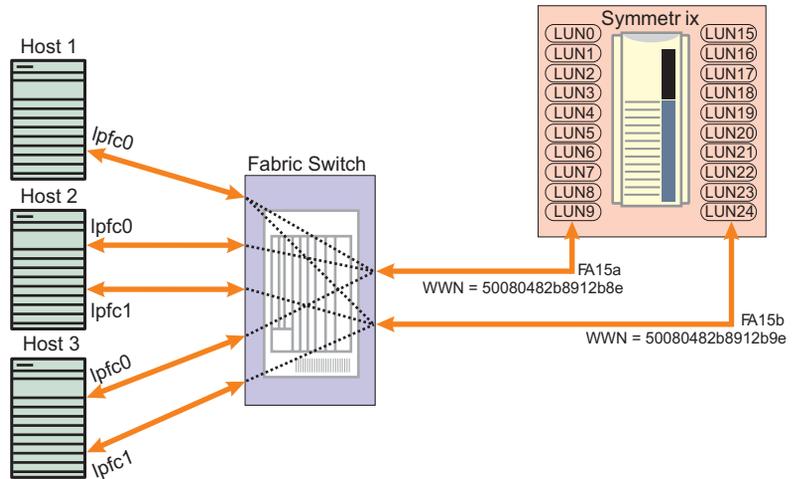
## C3400/3500 Storage Systems

1. In `/kernel/drv/lpfc.conf`, set parameters as follows:
  - In a single-server configuration:
    - Set **tgt-queue-depth** to **128**.
    - Calculate **lun-queue-depth** by dividing the **tgt-queue-depth** (in this case 128) by the maximum number of LUNs bound per SP.  
*Example:* If four LUNs are bound to each SP, you would set this variable to **32** ( $128 / 4 = 32$ ).
    - Set **nodev-tmo** to **20**.
    - Set **dqfull-throttle** to **0**.
  - In a dual-server configuration:
    - Set **tgt-queue-depth** to **64**, allowing each initiator 64 outstanding I/Os per SP.
    - Calculate **lun-queue-depth** by dividing the **tgt-queue-depth** (in this case 64) by the maximum number of LUNs bound per SP.  
*Example:* If four LUNs are bound to each SP, you would set this variable to **16** ( $64 / 4 = 16$ ).
    - Set **nodev-tmo** to **20**.
    - Set **dqfull-throttle** to **0**.
2. Reboot the server (with a standard reboot) to make your changes take effect:

### Reboot

## Symmetrix/Fabric Case Study

This case study incorporates all of the information for the persistent binding setup and the host configuration files in a Solaris fabric environment.



Host	HBA	FA	Target	LUNs
Host 1	lpfc0	15a	0	0-3
		15b	1	15-17
Host 2	lpfc0	15a	0	4-6
	lpfc1	15b	1	18, 20
Host 3	lpfc0	15a	0	7-9
	lpfc1	15b	1	19, 21-24

**Host 1** In `/kernel/drv/lpfc.conf`:

```
fcplib-bind-WWPN= "50080482B8912B8E:lpfc0t0",  
"50080482B8912B9E:lpfc0t1";
```

In `/kernel/drv/sd.conf`:

```
name="sd" parent="lpfc" target=0 lun=0 hba="lpfc0";  
name="sd" parent="lpfc" target=0 lun=1 hba="lpfc0";  
name="sd" parent="lpfc" target=0 lun=2 hba="lpfc0";  
name="sd" parent="lpfc" target=0 lun=3 hba="lpfc0";  
name="sd" parent="lpfc" target=1 lun=15 hba="lpfc0";  
name="sd" parent="lpfc" target=1 lun=16 hba="lpfc0";  
name="sd" parent="lpfc" target=1 lun=17 hba="lpfc0";
```

**Sun2** In `/kernel/drv/lpfc.conf`:

```
fcplib-bind-WWPN= "50080482B8912B8E:lpfc0t0",  
"50080482B8912B9E:lpfc1t1";
```

In `/kernel/drv/sd.conf`:

```
name="sd" parent="lpfc" target=0 lun=4 hba="lpfc0";  
name="sd" parent="lpfc" target=0 lun=5 hba="lpfc0";  
name="sd" parent="lpfc" target=0 lun=6 hba="lpfc0";  
name="sd" parent="lpfc" target=1 lun=18 hba="lpfc1";  
name="sd" parent="lpfc" target=1 lun=20 hba="lpfc1";
```

**Sun3** In `/kernel/drv/lpfc.conf`:

```
fcplib-bind-WWPN= "50080482B8912B8E:lpfc0t0",  
"50080482B8912B9E:lpfc1t1";
```

In `/kernel/drv/sd.conf`:

```
name="sd" parent="lpfc" target=0 lun=7 hba="lpfc0";  
name="sd" parent="lpfc" target=0 lun=8 hba="lpfc0";  
name="sd" parent="lpfc" target=0 lun=9 hba="lpfc0";  
name="sd" parent="lpfc" target=1 lun=19 hba="lpfc1";  
name="sd" parent="lpfc" target=1 lun=21 hba="lpfc1";  
name="sd" parent="lpfc" target=1 lun=22 hba="lpfc1";  
name="sd" parent="lpfc" target=1 lun=23 hba="lpfc1";  
name="sd" parent="lpfc" target=1 lun=24 hba="lpfc1";
```

## Configuring a Boot Device on the Storage Array

Solaris hosts have been qualified for booting from EMC storage array devices interfaced through Fibre Channel as described in the *EMC Support Matrix*. This chapter describes the process to configure a storage array device as a boot device.

- ◆ Configuring a Symmetrix Boot Device .....3-2
- ◆ Configuring a CLARiiON Boot Device..... 3-11

---

## Configuring a Symmetrix Boot Device

Boot support for Symmetrix is available to Emulex HBAs in Solaris hosts as described in the *EMC Support Matrix*.

---

### Requirements

The designated Symmetrix boot device must meet the following requirements:

- ◆ The boot device must have enough disk space to hold the Solaris operating system.
- ◆ For better performance, the boot device should not share the spindle with any other devices. For example, assume devices 304, 203, 405, and 103 occupy physical device C1M1. If the device 203 is being used for booting, no hosts should use device 304, 405, or 103.
- ◆ The boot device must have the similar partition structures, and each partition must be at least as large as the source partition (the partition from where the image of the partition is being copied).
- ◆ The installation requires a host drive that will act as the source of the image for the Symmetrix boot drive.
- ◆ The installation requires administrator understanding of the data backup from the host internal boot drive to the Symmetrix device.

The procedure described in this section assumes that you have installed the Emulex HBA(s), driver(s), and Symmetrix device (over arbitrated loop) onto the Solaris host, and that the Solaris operating system sees the Fibre Channel drives.

If necessary, update the HBA firmware and Open Boot support to at least the minimum levels shown in the previous table.

---

### Procedure

Follow these steps (described in detail later) to configure a boot device:

- ◆ Verify Symmetrix information
- ◆ Configure partitions
- ◆ Create a filesystem
- ◆ Install a Bootblk
- ◆ Copy required files
- ◆ Modify OpenBoot

## Verifying Symmetrix Information

Identify the controller instance and Symmetrix device number for the installation of Solaris. Refer to *Adding LUNS* on page 2-26 for information on Loop ID assignment.

## Configuring Partitions

1. Use the **format** command to verify that the sizes of the partitions on the Symmetrix device chosen for the Solaris installation are large enough to copy the current Solaris 7 operating system partitions. Examine the partitions of the host source drive where the current operating system resides:
  - a. At the %> prompt, type **format** and press ENTER.
  - b. When prompted choose `disk 0`, press ENTER.
  - c. At the `format>` prompt, type **partition** and press ENTER.
  - d. At the `partition>` prompt, type **print** and press ENTER.

Sample output:

Part	Tag	Flag	Cylinders	Size
0	root	wm	0-335	124.69 MB
1	swap	wm	336-792	169.59 MB
2	backup	um	0-5846	2.12GB
3	unassigned	wm	0	0
4	unassigned	wm	0	0
5	unassigned	wm	0	0
6	usr	wm	739-2305	561.40 MB
7	home	wm	2306-5846	1.28GB

2. Select the Symmetrix device on which to install the Solaris operating system as follows:
  - a. At the `partition>` prompt, type **quit** and press ENTER to return to the **format** menu.
  - b. At the `format>` prompt, type **disk** and press ENTER.
  - c. Type **1** and press ENTER when information similar to the following appears:

```
AVAILABLE SELECTIONS
...
1. c1t0d0 <EMC-SYMMETRIX-5265 cyl 4088 alt 2 hd
   15
   __sec 64> /sbus@1f,0/fce@1,400000/sd@0,0
...
specify disk (enter its number):
```

3. Repartition the Symmetrix device (if necessary) to model the partition information from the host source drive in order to match the drive configuration.
  - a. At the `format` prompt, type **partition** and press ENTER.
  - b. At the `partition` prompt, type **modify** and press ENTER.
  - c. After the partition in the Symmetrix device has been created, type **modify** and press ENTER at the `format` prompt.

---

## Creating Filesystems

Create filesystems on the required partitions in the designated Symmetrix boot device.

Below is an example in which the operating system is being copied to a Symmetrix Device of Target 3 LUN 0 on controller 1:

```
newfs /dev/rdisk/c1t3d0s0
newfs /dev/rdisk/c1t3d0s6
newfs /dev/rdisk/c1t3d0s7
```

To use a different Symmetrix Device, change **t3** to **t<Symmetrix Device target number>**. To use a different controller, change **c1** to **c<controller number>**. You must create both the root partition (slice 0) and the `/usr` partition (slice 6). You should also create other partitions (`/export/home`, `/var`, `/opt`) to mirror the current layout. Use the `newfs` command to create the filesystems.

---

## Installing a Bootblk

Install a *bootblk* (the information required by the host to boot the OS) on the root partition of the Symmetrix boot device.

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

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

```
SunOS patriot 5.7 Generic_106541-07 sun4u Sparc SUNW,Ultra -4
```

2. At the `%>` prompt, enter the following to install a bootblk onto the Symmetrix boot device:

```
/usr/sbin/installboot /usr/platform/sun4u/lib/fs/ufs/bootblk /dev/rdisk/  
c<controller number>t<Symmetrix device target number>d0s0
```

For example, entering the following installs a bootblk onto a boot device of Target 3 LUN 0 on controller 1:

```
/usr/sbin/installboot /usr/platform/sun4u/lib/fs/ufs/bootblk /dev/rdisk/
c1t3d0s0
```

## Copying Required Files

Use the **ufsdump** and **ufsrestore** commands to copy the required files and directories from the current host source drive to the Symmetrix boot device.

1. At the %> prompt, type **mount /dev/dsk/c1t3d0s0 /mnt** and press ENTER.

This mounts the root directory of the Symmetrix device to **/mnt**.

2. Use the **ufsdump** command to create the required directory structure on the new Symmetrix device and copy the files.

Below is an example where the current operating system is located on **/dev/dsk/c0t0d0**.

*Example:* At the %> prompt, type the following and press and press ENTER:

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

When the above command completes, the Symmetrix device (**c1t3d0s0**) will have the complete image of the root partition.

3. In the **/mnt/etc** directory update the **vfstab** file to indicate the Symmetrix boot device address to be mounted during boot. Modify all partitions that will be located on the Symmetrix boot device.

a. At the %> prompt, type **cd /mnt/etc** and press ENTER.

b. At the %> prompt, type **vi vfstab** and press ENTER.

*Example* — Before modification:

#device	device	mount	FS	fscck	mount	mount
#to mount to	fscck	point	type	pass	at boot	options
#						
/dev/dsk/c0t0d0s1	-	-	swap	-	no	-
/dev/dsk/c0t0d0s0	/dev/rdisk/c0t0d0s0	/	ufs	1	no	-
/dev/dsk/c0t0d0s6	/dev/rdisk/c0t0d0s6	/usr	ufs	1	no	-
swap - /tmp tmpfs	- yes -					

*Example* — After modification:

#device	device	mount	FS	fsck	mount	mount
#to mount to	fsck	point	type	pass	at boot	options
#						
/dev/dsk/c1t3d0s1	-	-	swap	-	no	-
/dev/dsk/c1t3d0s0	/dev/rdisk/c1t2d0s0	/	ufs	1	no	-
/dev/dsk/c1t3d0s6	/dev/rdisk/c1t2d0s6	/usr	ufs	1	no	-
swap - /tmp tmpfs	- yes -					

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

4. At the `%>` prompt, type **umount /mnt** and press ENTER to unmount the root partition.
5. Repeat steps 1, 2, and 4 for the `/usr` partition. In the current root directory:
  - a. At the `%>` prompt, type **mount /dev/dsk/c1t3d0s6 /mnt** and press ENTER.
  - b. At the `%>` prompt, type the following and press ENTER:  
**ufsdump 0f - /dev/dsk/c0t0d0s6 | ( cd /mnt; ufsrestore rf -)**
  - c. At the `%>` prompt, type **umount /mnt** and press ENTER.
6. Repeat steps 1, 2, and 4 for any other partitions to be mounted from the Symmetrix boot device.

## Modifying OpenBoot

Follow this procedure to modify OpenBoot:

1. Halt the system and issue a reset to get to the OpenBoot environment:
  - a. At the `%>` prompt, type **halt** and press ENTER.
  - b. At the `OK>` prompt, type **setenv auto-boot? false** and press ENTER, so the system will display the `OK>` prompt after reset.
  - c. At the `OK>` prompt, type **reset-all** and press ENTER.
2. Specify the topology for the target Fibre Channel boot disk:
  - a. At the `OK>` prompt, type the following and press ENTER:  
**" <device path name>" select-dev**  
 (Example: **" /pci@1f,4000/lpfc@4" select-dev**)

- b. At the `OK>` prompt, type the appropriate command and press ENTER:
 

```
FC-AL: set-fc-al
FC-SW: set-ptp
```
  - c. At the `OK>` prompt, type `unselect-dev` and press ENTER.
3. At the `OK>` prompt, type `reset-all` and press ENTER to get to the OpenBoot environment.
4. At the `OK>` prompt, type `printenv boot-device` and press ENTER to verify the default boot device. The resulting display should be similar to the following example:
 

```
boot-device disk net
```
5. Set up the new alias for your local boot disk:
  - a. At the `OK>` prompt, type `devalias` and press ENTER.
  - b. Note the first appearance of the long path name associated with your local boot disk. For example:
 

```
disk /pci@1f,4000/scsi@3/disk@0,0
```
  - c. At the `OK>` prompt, type `nvalias localdisk <path name>` and press ENTER. This assigns the alias `localdisk` to the local boot disk.
6. Set up the boot ID for the Symmetrix device boot disk:
  - a. At the `OK>` prompt, type `reset` and press ENTER to reset the host.
  - b. After the system has reset, at the `OK>` prompt, type `probe-scsi-all` and press ENTER to find the AL-PA and WWPN of the Symmetrix boot device.

#### Example of output:

```
probe-scsi-all
/pci@1f,4000/lpfc@5
Target none ALPA e2 WWPN 5006.0482.bbff.4e1f
Unit 0 Disk EMC SYMMETRIX 5265
Unit 3 Disk EMC SYMMETRIX 5265
Unit 4 Disk EMC SYMMETRIX 5265

/pci@1f,4000/lpfc@4
Target none ALPA e4 WWPN 5006.0482.bbff.4e0f
Unit 0 Disk EMC SYMMETRIX 5265
Unit 1 Disk EMC SYMMETRIX 5265
Unit 2 Disk EMC SYMMETRIX 5265
```

- c. Set the boot device ID, using the below as an example.  
*Example* — Assume the Symmetrix boot device has this configuration:

- Path: `/pci@1f,4000/lpfc@4`
- Alpha: `e4`
- WWPN: `5006.0482.bbff.4e0f`
- Target: `2`
- LUN: `0`

Follow these steps to set the boot device ID:

- i. At the `OK>` prompt, type "`/pci@1f,4000/lpfc@4`" `select-dev` and press ENTER.
- ii. At the `OK>` prompt, type the appropriate command and press ENTER:  
`FC-AL: did e4 0 2 set-boot-id`  
`FC-SW: wwpn 5006.0482.bbff.4e0f 0 2 set-boot-id`
- iii. At the `OK>` prompt, type `unselect-dev` and press ENTER.

---

To remove a boot ID just assigned:

- i. At the `OK>` prompt, type "`/pci@1f,4000/lpfc@4`" `select-dev` and press ENTER.
  - ii. At the `OK>` prompt, type `remove-boot-id #` and press ENTER.
  - iii. At the `OK>` prompt, type `unselect-dev` and press ENTER.
- 

7. Set up the alias and boot device for your new Symmetrix boot disk:
  - a. At the `OK>` prompt, type `nvalias symmdisk <pathname>` and press ENTER.  
*Example:* `nvalias symmdisk /pci@1f,4000/lpfc@4/sd@2,0`
  - b. At the `OK>` prompt, type `nvstore` and press ENTER to save your alias.
  - c. At the `OK>` prompt, type `setenv boot-device symmetrix` and press ENTER.
8. Reboot the system with your new boot disk:
  - a. At the `OK>` prompt, type `setenv auto-boot? true` and press ENTER, so your system will not stop at the `OK>` prompt after reset.
  - b. At the `OK>` prompt, type `boot symmdisk` and press ENTER to boot the host with the Symmetrix boot device.

## Troubleshooting

This section lists some possible problems and describes their solutions:

Problem	During boot process cannot mount or <b>fsck</b> .
Cause	<b>/etc/vfstab</b> file on the Symmetrix boot device did not updated correctly.
Solution	At the <b>OK&gt;</b> prompt, type <b>boot localdisk</b> ENTER to boot the system using internal disk. Correct the file <b>/etc/vfstab</b> on the Symmetrix boot device.

Problem	After entering <b>boot disk</b> , messages say that the loaded file is not executable.
Cause	There might not be a boot device on the target disk, or the copy from local to Symmetrix device was unsuccessful.
Solution	At the <b>OK&gt;</b> prompt, type <b>boot localdisk</b> ENTER to boot from the internal disk. Recreate the Symmetrix boot disk.

Problem	System hangs for a long time after displaying driver information, or reboots just after displaying the driver information.
Cause	Possible incorrect topology in the <b>/kernel/drv/lpfc.conf</b> file on the Symmetrix boot device.
Solution	At the <b>OK&gt;</b> prompt, type <b>boot localdisk</b> ENTER to boot from the internal disk. Correct the <b>/kernel/drv/lpfc.conf</b> on the Symmetrix device.

---

## Useful Monitor Commands

Here are some system monitor commands that can be useful after selecting a boot ID:

Command	Function
<code>.boot-id</code>	Prints the current boot device ID
<code>.version</code>	Prints the current version of Open Boot
<code>.topology</code>	Prints the current topology
<code>.nvram</code>	Prints the current flags for Open Boot

## Configuring a CLARiiON Boot Device

The current version of Open Boot supports FC-AL, private loop, public loop, and fabric point-to-point from a single host. It also supports multi-initiators and multi-LUNs (0 through 255).

The following is a procedure for installing the Solaris operating system onto a storage system LUN (drive). When the installation procedure is complete, the host can boot and run from the LUN (drive).

---

Be sure to back up any critical files on the local boot disk.

---

### Prerequisites

Before starting this procedure make sure that:

- ◆ the host has a local boot disk (internal hard drive).
- ◆ the capacity of the intended (target) Fibre Channel disk is large enough to hold all the file systems and swap slice of the local current boot disk.
- ◆ you use persistent bindings (assigning a specific SCSI ID to a specific WWPN or DID) if the Fibre Channel adapter is connected to more than one logical disk drive.
- ◆ you correlate the intended boot disk displayed by the **probe-scsi-all** command to the devices displayed by the operating system.

### Boot Disk Configuration

Use this procedure to make a bootable system disk from any size logical disk. This procedure also requires that you use a local boot disk to build the Fibre Channel boot disk.

1. At the `ok>` prompt, boot the server's internal disk to set up the new disk by entering the following command:

```
boot
```

2. If you do not see any Fibre Channel disks after booting the system, you may need to do one or both of the following
  - Enter the **touch /reconfigure** command and then reboot.
  - Edit the **/kernel/drv/lpfc.conf** file on the boot disk to change the topology from 4 to 2 or from 2 to 4.

- When Solaris has successfully rebooted, use the **format** command to do the following:

---

The following example assumes that you have selected a logical drive designated as **c1t0d0** (controller 1, target 0, lun 0).

---

- Select the logical drive to become boot disk (for example, **c1t0d0**).
  - Select the partition option and partition disk as desired.
  - Select the label option and write volume label to target disk.
- To install the boot block on the target disk, enter the following command:

```
installboot /usr/platform'uname-i' /lib/fs/ufs/bootblk /dev/rdisk/c1t0d0s0
```

- To create a file system for each partition that will contain a mounted file system, enter the following commands:

```
newfs -v /dev/rdisk/c1t0d0s0 (will become root)
newfs -v /dev/rdisk/c1t0d0s6 (will become usr)
newfs -v /dev/rdisk/c1t0d0s7 (will become export/home)
```

- To create temporary mount points for the new partitions so they can be mounted for copying, enter the following commands:

```
mkdir /root2
mkdir /usr2
mkdir /export2
```

- Mount the new partitions to the appropriate directories.
- Copy the applicable file systems. An example command sequence for the root file system is:

```
cd /root2
ufsdump 0f - /dev/dsk/c0t0d0s0 | ufsrestore rf -
```

- Edit **/root2/etc/vfstab** to point to the new boot disk as the physical location of mount points for **root**, **swap**, **usr**, and so on. Change the current entries as follows:
  - **/dev/dsk/c0t0d0s0** and **/dev/rdisk/c0t0d0s0** to show **c1t0d0s0** (root)
  - **/dev/dsk/c0t0d0s6** and **/dev/rdisk/c0t0d0s6** to show **c1t0d0s6** (usr)

- `/dev/dsk/c0t0d0s7` and `/dev/rdisk/c0t0s0s7` to show `c1t0d0s7` (export)
  - Change entry for swap to show `c1t0d0s3`
10. If multiple targets exist on a Fibre Channel adapter, you must edit `/root2/kernel/drv/lpfc.conf` to set up a persistent binding that will ensure that the desired logical disk is always assigned the correct target ID. Refer to the instructions in the `lpfc.conf` file to set persistent binding.

The disk is now ready to use as a boot disk.

11. Shut down the system.
12. At the `ok>` prompt enter the following command:

```
reset-all
```

## Boot Setup

1. To display the open boot driver version, enter the following command sequence at the `ok>` prompt:

```
select /lpfc  
.version  
unselect-dev
```

2. Change the alias for the current local boot disk from `disk` to another name (for example, `localdisk`) as follows:

- a. Enter the following command:

```
Devalias
```

- b. Record the long pathname after the first mention of `disk`, and use that pathname in the following command:

```
nvalias localdisk <path name>
```

This assigns the alias `localdisk` to the local boot disk.

3. Determine the device WWPN, DID or AL-PA by entering the following command:

```
probe-scsi-all
```

4. Record the actual WWPN, DID or AL-PA and the corresponding boot entry.

5. Select the appropriate topology (point-to-point or FC-AL) for the intended (target) Fibre Channel boot disk:

```
select /xxxx
  where xxxx is the full pathname to the boot device you
  recorded from the probe-scsi-all command. For example:
  /pci@2,2000/lpfc@2
set-fc-al (for arbitrated loop) or set-ptp (for fabric)
unselect-dev
reset-all
```

6. Set the boot device ID.

- To use WWPN, enter the following commands:

```
select /xxxx
  where xxxx is the full pathname to the boot device you
  recorded from the probe -scsi -all command. For example:
  /pci@2,2000/lpfc@2
wwpn <wwpn> <lun> <target id> set-boot-id
  where:
    <wwpn> must be in hex
    <lun> is in hex; to change to decimal type d# <lun>
    <target id> is in hex; to change to decimal,
    type d# <target id>
  For example, a valid boot device ID using the WWPN is:
  22000020.37997f70 d# 140 d# 10 set-boot-id
unselect-dev
reset-all
```

- To use DID (public loop), enter the following commands:

```
select /xxxx
  where xxxx is the full pathname to the boot device you
  recorded from the probe -scsi -all command. For example:
  /pci@2,2000/lpfc@2
did <did> <lun> <target id> set-boot-id
  where:
    <did> must be in hex
    <lun> is in hex; to change to decimal, type d# <lun>
    <target id> is in hex; to change to decimal,
    type d# <target id>
  For example, a valid boot device ID using the DID is:
```

```
2009e1 d# 100 d# 10 set-boot-id
unselect-dev
reset-all
```

- To use AL-PA (private loop), enter the following commands:

```
select /xxxx
```

where **xxxx** is the full pathname to the boot device you recorded from the **probe -scsi -all** command. For example:

```
/pci@2,2000/lpfc@2
```

```
did <alpa> <lun> <target id> set-boot-id
```

where:

<alpa> must be in hex

<lun> is in hex; to change to decimal, type **d# <lun>**

<target id> is in hex; to change to decimal,

type **d# <target id>**

For example, a valid boot device ID using ALPA is:

```
e1 d# 100 d# 10 set-boot-id
```

```
unselect-dev
reset-all
```

7. To remove the boot ID you just assigned, enter the following commands:

```
select /xxxx
```

where **xxxx** is the full pathname to the boot device you recorded from the **probe -scsi -all** command. For example:

```
/pci@2,2000/lpfc@2
```

```
remove-boot-id (to clear boot ID settings)
```

```
unselect-dev
reset-all
```

## Troubleshooting Suggestions

The following table lists symptoms and possible causes for problems you may experience when using a storage-system LUN as a boot device.

Symptom	Possible Cause
During system initialization, the boot sequence cannot mount or filecheck <b>/usr vfstab</b> on the boot disk, or the boot disk is not updated correctly.	Persistent binding is not implemented or not implemented correctly.
After entering boot disk, a message displays stating that the file that was loaded is not executable.	No boot on target disk because copy from local disk to target disk was unsuccessful.
System hangs for a long time after displaying driver information, or System reboots right after displaying driver information.	The topology may be incorrect. Check the topology value in the <b>/kernel/drv/lpfc.conf</b> file for the target disk.

## Helpful ok Prompt Commands

<b>select /lpfc</b>	Selects the lpfc device.
<b>.boot-id</b>	Prints current boot device ID.
<b>.version</b>	Print current version of open boot.
<b>.topology</b>	Prints current topology.
<b>.nvram</b>	Prints out current flags for open boot.
<b>unselect-dev</b>	Deselects the lpfc device.

## Booting From CLARiON Over PowerPath

You can use either a native-named device or *emcpower* device as a boot device. If using an *emcpower* device, the host adapter on the *emcpower* device must support booting in its FCode.

Using a PowerPath device as the boot device provides load balancing and path failover for the boot device. Native-named devices, however, do not provide boot time boot path failover.

This procedure assumes PowerPath is already installed.

### Boot From Native Solaris Device

Follow these steps to use a native device:

1. Select the LUN that will be used for booting by running the **format** command.
2. On the target LUN selected in step 1, create `/`, `/usr`, `/var`, `/export/home`, `/var`, `/opt`, `swap`, and any other required partitions that match the current boot disk's partitions. You must make all of the partitions greater than that of the current boot disk's partitions.
3. Label the disk to save the updated partitions by issuing the **label** command while still in the **format** utility.
4. Create a new filesystem for each partition, excluding `swap` and `backup`.

*Root partition example:* **newfs /dev/rdisk/c2t0d0s0**

5. Create a directory and mount each partition as in the following example:

```
mkdir /bootarray
mount /dev/dsk/c2t0d0s0 /bootarray
mkdir /bootarray/usr
mount /dev/dsk/c2t0d0s6 /bootarray/usr
```

6. Change to the root directory of the LUN you have set up and dump the O/S:

*Example:*

```
cd /bootarray
ufsdump 0f - /dev/dsk/c0t0d0s0 | ufsrestore rf -
```

Repeat this for each partition you created.

- Record the complete path of the boot LUN by running the **format** command and taking note of the path listed directly below the device.

*Example:* `/ssm@0,0/pci@1c,600000/lpfc@1/sd@0,0`

- Once the dump is done, edit `/bootarray/etc/vfstab` to change the original partitions to reflect the partitions on the new boot disk.

*Example:* **vi /bootarray/etc/vfstab**

Change the lines shown in bold below as follows:

*Before:*

```
#device    device    mount    FS        fck    mount    mount
#to mount  to fck    point    type      pass   at boot  options
#
#/dev/dsk/c1d0s2 /dev/rdisk/c1d0s2 /usr    ufs       1       yes     -
fd - /dev/fd fd - no -
/proc - /proc proc - no -
/dev/dsk/c0t0d0s1 - - swap - no -
/dev/dsk/c0t0d0s0 /dev/rdisk/c0t0d0s0 / ufs 1 no -
/dev/dsk/c0t0d0s6 /dev/rdisk/c0t0d0s6 /usr ufs 1 no -
swap - /tmp tmpfs - yes -
```

*After:*

```
#device    device    mount    FS        fck    mount    mount
#to mount  to fck    point    type      pass   at boot  options
#
#/dev/dsk/c1d0s2 /dev/rdisk/c1d0s2 /usr    ufs       1       yes     -
fd - /dev/fd fd - no -
/proc - /proc proc - no -
/dev/dsk/c2t0d0s1 - - swap - no -
/dev/dsk/c2t0d0s0 /dev/rdisk/c2t0d0s0 / ufs 1 no -
/dev/dsk/c2t0d0s6 /dev/rdisk/c2t0d0s6 /usr ufs 1 no -
swap - /tmp tmpfs - yes -
```

- Add the following lines to `/bootarray/etc/system`:

```
rootfs:ufs
rootdev:/ssm@0,0/pci@1c,600000/lpfc@1/sd@0,0:a
```

- Edit `/bootarray/etc/dumpadm.conf` and change the original swap device to the new swap device.

---

You do not have to edit `dumpadm.conf` for Solaris 2.6 or earlier.

---

11. Install the boot block on the new boot device by running the following:

```
installboot /usr/platform/`uname -i`/lib/fs/ufs/bootblk /dev/rdisk/c2t0d0s0
```

12. Bring the host down to the OBP by issuing the **halt** command.
13. While in OBP, issue the **reset-all** command.
14. Issue the **show-devs** command, and select the correct path to the device to be used as the new boot device.

*Example:*

```
" /ssm@0,0/pci@1c,600000/lpfc@1" select-dev
```

15. To set the correct topology, type the following and press ENTER:

- ◆ Fabric: **set-ptp**
- ◆ Loop: **set-fc-al**

16. To display all Fibre Channel devices along with WWPN of the target SP from which you are booting, issue the **show-children** command.

17. Set the boot ID of the new boot device:

- a. Make note of the actual WWPN, DID, or ALPA returned from the **show-children** command, and copy down the corresponding boot entry.

- b. Set the boot ID:

- For point-to-point (WWPN):  
`wwpn <WWPN> <LUN> <target ID> set-boot-id`
- For public loop (DID):  
`did <DID> <LUN> <target ID> set-boot-id`
- For private loop (ALPA):  
`did <ALPA> <LUN> <target ID> set-boot-id`

18. Create an alias for the new boot device:

```
nvalias emcboot /ssm@0,0/pci@1c,600000/lpfc@1/sd@0,0:a  
nvstore
```

19. You are now ready to boot:

```
unselect-dev  
reset-all  
boot emcboot
```

### Boot From emcpower Device

Follow these steps to use an emcpower device:

1. Select the LUN that will be used for booting by running the **powermt display** command.

*Example:* **powermt display dev=0**

*Output from example:*

```
Pseudo name=emcpower0a
CLARiiON ID=WRE00021500561
Logical device ID=6006016AF9080000D6577E33A769D611
state=alive; policy=CLAROpt; priority=0; queued-I/Os=0
=====
----- Host ----- - Stor - -- I/O Path - -- Stats ---
### HW Path          I/O Paths   Interf.   Mode    State  Q-I/Os  Errors
=====
  0 ssm@0,0/pci@1c,700000/lpfc@2 c1t0d8s0  SP A0   active  alive  0  0
  0 ssm@0,0/pci@1c,700000/lpfc@2 c1t1d8s0  SP B0   active  alive  0  0
  0 ssm@0,0/pci@1c,700000/lpfc@2 c1t2d8s0  SP A1   active  alive  0  0
  0 ssm@0,0/pci@1c,700000/lpfc@2 c1t3d8s0  SP B1   active  alive  0  0
  1 ssm@0,0/pci@1c,600000/lpfc@1 c2t0d8s0  SP A0   active  alive  0  0
  1 ssm@0,0/pci@1c,600000/lpfc@1 c2t1d8s0  SP B0   active  alive  0  0
  1 ssm@0,0/pci@1c,600000/lpfc@1 c2t2d8s0  SP A1   active  alive  0  0
  1 ssm@0,0/pci@1c,600000/lpfc@1 c2t3d8s0  SP B1   active  alive  0  0
  2 ssm@0,0/pci@d,700000/lpfc@1 c3t0d8s0  SP A0   active  alive  0  0
  2 ssm@0,0/pci@d,700000/lpfc@1 c3t1d8s0  SP B0   active  alive  0  0
  2 ssm@0,0/pci@d,700000/lpfc@1 c3t2d8s0  SP A1   active  alive  0  0
  2 ssm@0,0/pci@d,700000/lpfc@1 c3t3d8s0  SP B1   active  alive  0  0
  3 ssm@0,0/pci@d,600000/lpfc@1 c4t0d8s0  SP A0   active  alive  0  0
  3 ssm@0,0/pci@d,600000/lpfc@1 c4t1d8s0  SP B0   active  alive  0  0
  3 ssm@0,0/pci@d,600000/lpfc@1 c4t2d8s0  SP A1   active  alive  0  0
  3 ssm@0,0/pci@d,600000/lpfc@1 c4t3d8s0  SP B1   active  alive  0  0
```

Now run the **format** command.

2. On the target LUN selected in step 1, create `/`, `/usr`, `/var`, `/export/home`, `/var`, `/opt`, `swap`, and any other required partitions that match the current boot disk's partitions. You must make all of the partitions greater than that of the current boot disk's partitions.
3. Label the disk to save the updated partitions by issuing the **label** command while still in the **format** utility.
4. Create a new filesystem for each partition, excluding `swap` and `backup`.

*Root partition example:* **newfs /dev/rdisk/emcpower0a**

5. Create a directory and mount each partition as in the following example:

```
mkdir /bootarray
mount /dev/dsk/emcpower0a /bootarray
mkdir /bootarray/usr
mount /dev/dsk/emcpower0g /bootarray/usr
```

6. Change to the root directory of the LUN you have set up and dump the O/S:

*Example:*

```
cd /bootarray
ufsdump 0f - /dev/dsk/c0t0d0s0 | ufsrestore rf -
```

Repeat this for each partition you created.

7. Record the complete path of the boot LUN by running the following command:

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

Output from command:

```
lrwxrwxrwx 1 root  other    33 Jun  4 09:06 /dev/dsk/emcpower0a ->
../../devices/pseudo/emcp@0:a,blk
```

Record the `/pseudo/emcp@0:a,blk`.

8. Once the dump is done, edit `/bootarray/etc/vfstab` to change the original partitions to reflect the partitions on the new boot disk.

Example: **vi /bootarray/etc/vfstab**

Change the lines shown in bold below as follows:

*Before:*

```
#device    device    mount    FS        fskck    mount    mount
#to mount  to fsck   point    type      pass     at boot  options
#
#/dev/dsk/c1d0s2 /dev/rdsk/c1d0s2 /usr    ufs      1        yes     -
fd - /dev/fd fd - no -
/proc - /proc proc - no -
/dev/dsk/c0t0d0s1 - - swap - no -
/dev/dsk/c0t0d0s0 /dev/rdsk/c0t0d0s0 / ufs 1 no -
/dev/dsk/c0t0d0s6 /dev/rdsk/c0t0d0s6 /usr ufs 1 no -
swap - /tmp tmpfs - yes -
```

*After:*

```
#device    device    mount    FS        fskck    mount    mount
#to mount  to fsck   point    type      pass     at boot  options
#
#/dev/dsk/c1d0s2 /dev/rdsk/c1d0s2 /usr    ufs      1        yes     -
fd - /dev/fd fd - no -
/proc - /proc proc - no -
/dev/dsk/emcpower0b - - swap - no -
/dev/dsk/emcpower0a /dev/rdsk/emcpower0a / ufs 1 no -
/dev/dsk/emcpower0g /dev/rdsk/emcpower0g /usr ufs 1 no -
swap - /tmp tmpfs - yes -
```

9. Add the following lines to the end of `/bootarray/etc/system`:

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

10. Edit `/bootarray/etc/dumpadm.conf` and change the original swap device to the new swap device.

---

You do not have to edit `dumpadm.conf` for Solaris 2.6 or earlier.

---

11. Install the boot block on the new boot device by running the following:

```
installboot /usr/platform/`uname -i`/lib/fs/ufs/bootblk /dev/rdsk/emcpower0a
```

12. To find the boot device information to reconfigure the boot disk EEPROM setting, issue the **powermt display** command as you did in step 1, and issue **ls -l** on any one of the devices that appears in the output of **powermt display**.

*Example: **ls -l c1t0d0s0***

Output from command:

```
lrwxrwxrwx 1 root  other   33 Jun  4 09:06 /dev/dsk/c1t0d0s0 ->../
/devices/ssm@0,0/pci@1c,700000/lpfc@2/sd@0,8:a
```

Record the `ssm@0,0/pci@1c,700000/lpfc@2/sd@0,8:a`.

13. Bring the host down to the OBP by issuing the **halt** command.

14. While in OBP, issue the **reset-all** command.

15. Issue the **show-devs** command, and select the correct path to the device to be used as the new boot device (recorded in step 11).

*Example: " **/ssm@0,0/pci@1c,700000/lpfc@2**" **select-dev***

16. To set the correct topology, type the following and press ENTER:

- ◆ Fabric: **set-ptp**
- ◆ Loop: **set-fc-al**

17. To display all Fibre Channel devices along with WWPN of the target SP from which you are booting, issue the **show-children** command.

18. Set the boot ID of the new boot device:

- a. Make note of the actual WWPN, DID, or ALPA returned from the **show-children** command, and copy down the corresponding boot entry.

- b. Set the boot ID:

- For point-to-point (WWPN):  
**wwpn <WWPN> <LUN> <target ID> set-boot-id**
- For public loop (DID):  
**did <DID> <LUN> <target ID> set-boot-id**
- For private loop (ALPA):  
**did <ALPA> <LUN> <target ID> set-boot-id**

19. Create an alias for the new boot device:

```
nvalias emcpower0a /ssm@0,0/pci@1c,600000/lpfc@1/sd@0,0:a
```

```
nvstore
```

20. You are now ready to boot:

```
unselect-dev  
reset-all  
boot emcpower0a
```

## Symbols

/etc/system, editing 2-29

## B

boot device, CLARiiON, configuring 3-11

boot device, Symmetrix, configuring 3-2

## C

cables, connecting (CLARiiON) 2-4

case study, Symmetrix/fabric 2-41

## D

driver instance number, determining  
(CLARiiON) 2-25

driver, HBA

installing 2-5

obtaining 2-5

optimizing in CLARiiON environment 2-38

upgrading

PCI 2-12

SBus 2-8

## E

EMC Fibre Channel, incorporating 2-16

Emulex/Solaris environment 1-2

etc/system, editing 2-29

## F

fabric environment, preparing for 2-16

failover software, CLARiiON, removing 2-5

FCode, HBA, updating 2-33

firmware, HBA, updating 2-32

## H

HBA driver. *See* driver, HBA

HBA types, mixing 1-2

HBA, replacing (CLARiiON) 2-37

## I

instance number, driver, determining  
(CLARiiON) 2-25

## L

labeling devices 2-31

lpfc.conf, editing 2-19

LUNS, adding 2-26

## P

partitioning devices 2-31

persistent binding 2-16

## S

Solaris/Emulex environment 1-2

## **U**

upgrading the driver

PCI 2-12

SBus 2-8

## **Z**

zone, single-HBA, creating (CLARiiON) 2-35

zoning 1-2