



Removing and Replacing Hard Disk Drives on Sun Server Platforms

User Guide

Please Read

Important

Please read this entire guide. If this guide provides installation or operation instructions, give particular attention to all safety statements included in this guide.

Notices

Trademark Acknowledgements

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at www.cisco.com/go/trademarks.

Third party trademarks mentioned are the property of their respective owners.

The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1009R)

Publication Disclaimer

Cisco Systems, Inc. assumes no responsibility for errors or omissions that may appear in this publication. We reserve the right to change this publication at any time without notice. This document is not to be construed as conferring by implication, estoppel, or otherwise any license or right under any copyright or patent, whether or not the use of any information in this document employs an invention claimed in any existing or later issued patent.

Copyright

© 2007, 2010, 2012 Cisco and/or its affiliates. All rights reserved. Printed in the United States of America.

Information in this publication is subject to change without notice. No part of this publication may be reproduced or transmitted in any form, by photocopy, microfilm, xerography, or any other means, or incorporated into any information retrieval system, electronic or mechanical, for any purpose, without the express permission of Cisco Systems, Inc.

Contents

About This Guide	v
Replace a Disk Drive on a Sun Fire V880 and V890 Server	1
6-Disk Configuration.....	2
12-Disk Configuration.....	24
Replace a Disk Drive on a Sun Fire V445 Server	49
Hard Disk Layout	50
Verify Disk Drive Status	51
Prepare to Replace the Hard Drive	62
Avoid Electrostatic Discharge.....	65
Remove the Hard Drive.....	66
Install the Hard Drive.....	68
Reconnect the New Hard Drive.....	69
Customer Information	74

About This Guide

Introduction

This guide provides instructions for removing and replacing hard disk drives on various Sun Microsystems (Sun) servers.

Purpose

The purpose of this guide is to provide support engineers with instructions for removing and replacing hard disk drives on the Sun servers.

Audience

This guide is written for system operators who use a Sun server as their Digital Network Control System (DNCS) or Application Server platform. Cisco field service engineers and internal support technicians who assist system operators in managing their systems will also find this guide useful.

The individual responsible for completing the procedures in this guide must have the following capabilities:

- Familiarity with system administration using the Solaris operating system (OS)
- Knowledge of how to use a text editor, such as UNIX vi
- Privileges as a root user

Document Version

This is the third formal release of this document. In addition to minor text and graphic changes, the following table provides the technical changes to this document.

Description	See Topic
Instructions for removing and replacing hard drives on the Sun Fire V445 were added to the guide.	See <i>Replace a Disk Drive on a Sun Fire V445 Server</i> (on page 49).

1

Replace a Disk Drive on a Sun Fire V880 and V890 Server

Introduction

This chapter provides instructions for removing and replacing a hard disk drive on the Sun Fire V880 and V890 servers.

Important: Cisco engineers do not recommend shutting down the DNCS when replacing a hard disk drive. Your metadevices may become corrupted.

In This Chapter

- 6-Disk Configuration..... 2
- 12-Disk Configuration..... 24


6-Disk Configuration

V880/V890 Hardware (6-Disk Configuration)

Hard Disk Layout

Depending on your system, the Sun Fire V880 or V890 server is configured with either 6 or 12 disk drives.

The following illustration shows the front view of the 6-disk V880/V890 system layout.



Verify Disk Drive Status

If a disk problem occurs on your V880 or V890 server, an orange banner with white text (banner window) appears on your console to alert you to the problem. Do not

close this window; the window disappears when the problem is resolved. The disk problem may be an event where one or all portions of a disk drive become inaccessible.

Should one or all portions of a disk drive become inaccessible, you must use the **luxadm** utility to accomplish the following steps:

- Verify the enclosure name assigned to the server's internal storage array
- Determine the status of the disk drive

Note: Cisco's default enclosure name is FCloop.

After using the **luxadm** utility, use the **metadb** utility to determine if a meta database is on the disk drive. If a meta database exists on the drive, you must delete the meta database *prior* to removing the disk drive.

Finding the Enclosure Name

Complete these steps to verify the enclosure name for the internal storage array using the **luxadm probe** command.

- 1 If necessary, open an xterm window on the V880 or V890 server.
- 2 Complete the following steps to log on to the xterm window as **root** user.
 - a Type **su -** and press **Enter**. The password prompt appears.
 - b Type the root password and press **Enter**.
- 3 Type the following command and then press **Enter**. The name of the enclosure appears.

```
luxadm probe
```

Example: Sample output follows:

```
Found Enclosure:
```

```
SUNWGS INT FCBPL   Name:FCloop   Node WWN:508002000011d7d8
```

```
    Logical Path:/dev/es/ses0
```

```
    Logical Path:/dev/es/ses1
```

Note: **FCloop** is the enclosure name in this example.

Determining the Status of the Disk Drive

Complete these steps to determine the disk drive status using the **luxadm display** command.

Note: If you have been following the procedures in this chapter in order, you should still be logged in to an xterm window as root user.

Type the following command and press **Enter** to determine the disk drive status:

```
luxadm display [enclosure name]
```

Example: luxadm display FCloop

Sample output showing the disk status follows:

```
SUNWGS INT FCBPL
          DISK STATUS
SLOT   DISKS                (Node WWN)
0      On (O.K.)            2000002037f87099
1      On (O.K.)            2000002037f39796
2      On (Login failed)
3      On (O.K.)            2000002037ffab19
4      On (O.K.)            2000002037ff9ca5
5      On (O.K.)            2000002037ffa3c8
6      Not Installed
7      Not Installed
8      Not Installed
9      Not Installed
10     Not Installed
11     Not Installed

          SUBSYSTEM STATUS
FW Revision:9218   Box ID:0
Node WWN:508002000011d7d8   Enclosure Name:FCloop
SSC100's - 0=Base Bkpln, 1=Base LoopB, 2=Exp Bkpln, 3=Exp
LoopB
          SSC100 #0:      O.K.(11.A)
          SSC100 #1:      O.K.(11.A)
```

```
SSC100 #2:    Not Installed
SSC100 #3:    Not Installed
```

```
Temperature Sensors - 0 Base, 1 Expansion
0:21°C
1:Not Installed
```

Note: In this example, notice the status of slot 2 is **Login failed**. This indicates that Slot 2 is the disk drive that needs to be replaced.

Important: The World Wide Name (WWN) is unique for each disk drive. Whenever you remove a drive and replace it, you get a new WWN value for the new drive.

Examine Mirrored Devices (6-Disk Setup)

The V880 and V890 servers support disk drive mirroring. Through disk mirroring, the DNCS stores identical information across sets of hard disks.

After using the **luxadm display** command to determine the drive status, use the **check_metadevices** command to determine if a single file system or multiple file systems have been affected.

The V880/V890 6-disk configuration has a total of three mirrored disk drives. The hard disks in drive bay 0 and drive bay 3 are mirrored, as are the hard disks in drive bay 1 and drive bay 4. The hard disks in drive bay 2 and drive bay 5 are also mirrored.

Before you disable the disk mirroring functions of the V880/V890 in preparation for replacing a disk drive, you should examine the status of the mirrored disks in the system. Examining the status of the mirrored devices also tells you if a specific disk slice needs maintenance.

Examining the Mirrored Devices

Complete these steps to examine the status of the mirrored disk drives on your DNCS using the **check_metadevices** command.

Note: If you have been following the procedures in this chapter in order, you should still be logged in to an xterm window as root user.

1 Type the following command and press **Enter**:

```
su - dncs
```

- 2 Type the following command and press **Enter**. The metadvice tool examines the mirrored disk drives on your DNCS and indicates if a drive or a specific disk slice needs maintenance.

```
/dvs/dnccs/bin/check_metadevices
```

Example: Sample output pertaining to when a disk slice needs maintenance follows:

```
WARNING: d520 metadvice Needs maintenance
```

```
Found Disk Problems at: Mon Jul 12 11:08:20 EDT 2004
```

Important: Write the name of the disk drive or disk slice that needs maintenance on a separate sheet of paper.

- 3 Type **exit** and then press **Enter** to exit from the DNCS shell.

What Is the Status of the Mirrored Disk Drives?

After using the `check_metadevices` command to examine the status of mirrored disk drives, the following scenarios could occur:

- If one slice of the disk needs maintenance, resync the disk slice. Go to *Resync the Disk Drive* (on page 7).
- If you have errors on multiple slices per disk, complete the following steps to verify the disk drive status.
 - a Go back to *Verify Disk Drive Status* (on page 2) and run the `luxadm probe` command again.
 - b Choose one of the following options based on the `luxadm probe` command results:
 - If this command does *not* show a disk drive problem, contact Cisco Services for assistance.
 - If this command *shows* **login failed [drive name]**, you must replace the disk drive. Go to *Remove a Hard Disk Drive (6-Disk Setup)* (on page 9).

Important: If maintenance is required on drives that are mirrors of each other, such as `c1t2d0s7` and `c2t5d0s0`, contact Cisco Services.

Resync the Disk Drive

After examining the mirrored devices, when one disk slice needs maintenance, you must resync the disk drive. The resync process could take from a half-hour to 7 hours to complete, based on your system.

Complete these steps to resync the disk drive.

Note: If you have been following the procedures in this chapter in order, you should still be logged in to an xterm window as root user.

- 1 Using the disk slice you recorded in *Examining the Mirrored Devices* (on page 5), type the following command and press **Enter**.

```
metastat d[disk slice]
```

Example: metastat d520

Result: The submirrors of d520 are listed.

```
d520: Mirror
  Submirror 0: d420
    State: Needs maintenance
  Submirror 1: d720
    State: Okay
  Pass: 1
  Read option: roundrobin (default)
  Write option: parallel (default)
  Size: 125256384 blocks

d420: Submirror of d520
  State: Needs maintenance
  Invoke: metareplace d520 c1t2d0s0 <new device>
  Size: 125256384 blocks
  Stripe 0:
    Device      Start Block  Dbase State      Hot Spare
    c1t2d0s0      0           No   Maintenance

d720: Submirror of d520
  State: Okay
  Size: 125256384 blocks
  Stripe 0:
    Device      Start Block  Dbase State      Hot Spare
    c2t5d0s0      0           No   Okay
```

Notes:

- In this example, notice the state of Submirror 1: d420 is **Needs Maintenance**. This indicates that d420 is the submirror of d520 that requires maintenance.

- Also in this example, notice that the Device assignment for d420 is **c1t2d0s0**. You will use this value when you perform the meta replace command in the next step.
 - Refer to the *V880/V890 6-Disk Device Assignments* (on page 22) table for more information.
- 2 Locate the disk slice in the *V880/V890 Meta Replace Commands* (on page 23) table and type the meta replace command for the disk slice.
Example: In the previous example, the device assignment for d420 is c1t2d0s0. Type the following to resync the disk drive:

```
metareplace -e d520 c1t2d0s0
```


Important: Type the meta replace command that affects the specific disk slice, only.
Result: You return to a prompt, and the orange banner window shows the disk drive is resyncing.
 - 3 After you resync the drive, go back to the *Examining the Mirrored Devices* (on page 5) procedure and re-examine the mirrored devices using the check metadevices tool.
 - 4 Did the metadevice tool indicate that the disk slice needs maintenance?
 - If **yes**, you must replace the disk drive. Go to *Remove a Hard Disk Drive (6-Disk Setup)* (on page 9).
 - If **no**, you do not need to replace the disk drive. However, keep a written record of the disk slice on which you performed the meta replace command. If a problem continues to occur with the same disk slice, you may consider replacing the disk drive. No further action is required. You may stop reading this guide.

Remove a Hard Disk Drive (6-Disk Setup)

Finding a Meta Database

After examining the mirrored disk drives and determining that multiple portions of a disk drive are inaccessible, you must remove the hard disk drive. Before you can remove the drive, use the **metadb** utility to determine if the drive in question has a metadvice.

Important: The V880/V890 6-disk setup has meta databases on all six disk drives.

Note: For more information about metadvice assignments, refer to the *V880/V890 6-Disk Device Assignments* (on page 22) table.

Complete these steps using the metadb utility to determine if a meta database is located on the disk drive.

- 1 Type the following command and press **Enter** to display the status of the disk drive:

```
metadb -i
```

Example: Sample output follows:

```

      flags          first blk      block count
      a m p luo      16          1034
/dev/dsk/clt0d0s7
      a   p luo      16          1034
/dev/dsk/clt1d0s7
      a R p luo      16          1034
/dev/dsk/clt2d0s7
      a   p luo      16          1034
/dev/dsk/c2t3d0s7
      a   p luo      16          1034
/dev/dsk/c2t4d0s7
      a   p luo      16          1034
/dev/dsk/c2t5d0s7
o - replica active prior to last mddb configuration change
u - replica is up to date
l - locator for this replica was read successfully
c - replica's location was in /etc/lvm/mddb.cf
p - replica's location was patched in kernel
m - replica is master, this is replica selected as input
W - replica has device write errors
a - replica is active, commits are occurring to this
replica
M - replica had problem with master blocks
```

Chapter 1 Replace a Disk Drive on a Sun Fire V880 and V890 Server

D - replica had problem with data blocks
F - replica had format problems
S - replica is too small to hold current data base
R - replica had device read errors

- 2 Notice the drive flagged with a capital letter. This indicates which meta database has an error.

Example: In this example, **c1t2d0s7** is the drive with the meta database error.

- 3 Refer to the *V880/V890 6-Disk Diagram and Configuration Table* (on page 22) to verify the slot number matches the drive with the meta database error.

Example: Using the configuration table, the drive `/dev/disk/c1t2d0s7` is located in physical slot 2.

- 4 Based upon the information in the **V880/V890 6-Disk Diagram and Configuration Table**, did the meta database match the slot number of the hard disk drive?

- If **yes**, you must delete the meta database from the hard disk drive prior to removing the drive. Go to *Deleting a Meta Database* (on page 11).
- If **no**, contact Cisco Services for assistance.

Deleting a Meta Database

Complete these steps to delete a meta database on the hard disk drive.

Note: If you have been following the procedures in this chapter in order, you should still be logged in to an xterm window as root user.

- 1 Type the following command and press **Enter**:

```
metadb -d [device name]
```

Example: `metadb -d c1t2d0s7`

Result: The meta database is deleted from the disk drive.

Note: Write the device name that you delete on a separate piece of paper. You will use this name when you add a meta database to the disk drive, later in this chapter.

- 2 Before you physically remove the hard disk drive, you must put the drive into an offline state. Go to *Putting a Hard Disk Drive Into an Offline State* (on page 11).

Putting a Hard Disk Drive Into an Offline State

When the status of the disk drive is anything other than *not installed*, you must put the drive into the offline state.

Complete these steps to put the disk drive into the offline state.

Note: If you have been following the procedures in this chapter in order, you should still be logged in to an xterm window as **root** user.

- 1 Before you remove the drive, Cisco highly recommends that you have a recent backup of the drive data. Refer to your backup and restore guide for backup procedures.

Note: Refer to the *V880/V890 6-Disk Diagram and Configuration Table* (on page 22) to determine the disk role for the drive you are removing.

Examples:

- An error was found on drive **c1t0d0s7**. Based on the information in the 6-Disk configuration table, the role of drive **c1t0d0s7** is / disk1, meaning it has a file system role. Because a file system contains the Cisco software, you would perform a file system backup.
- If an error is found on a drive with a role of DB, this means that a drive with an Informix database role has an error. You would perform a database backup on the drive.

- 2 After backing up the drive data, type the following command and press **Enter**:

```
luxadm remove_device [enclosure name],[slot]
```

Example: `luxadm remove_device FCloop,s2`

Result: A warning message appears prompting you to be sure all data on the disk drive is backed up.

```
WARNING!!! Please ensure that no filesystems are mounted on
these device(s). All data on these devices should have been
backed up.
```

```
The list of devices being used (either busy or reserved) by
the host:
```

```
1: Box Name:      "FCloop"  slot 2
```

```
Please enter 's' or <CR> to Skip the "busy/reserved"
device(s) or
```

```
'q' to Quit and run the subcommand with
```

```
-F (force) option. [Default: s]: q
```

- 3 Type **q** to quit and then press **Enter**.

- 4 Type the following command and press **Enter**:

```
luxadm remove_device -F [enclosure name],[slot]
```

Example: `luxadm remove_device -F FCloop,s2`

Result: The WWN entries and device links are removed for the disk drive and you are prompted to verify the disk drive to remove.

WARNING!!! Please ensure that no filesystems are mounted on these device(s). All data on these devices should have been backed up.

The list of devices which will be removed is:

```
1: Box Name:      "FCloop" slot 2
   Node WWN:      20000002037ff9bee
   Device Type:   Disk device
   Device Paths:
     /dev/rdisk/c1t2d0s7
     /dev/rdisk/c2t2d0s2
```

- 5 Press **Enter**. You are prompted to remove the disk drive.

```
Please verify the above list of devices and
then enter 'c' or <CR> to Continue or 'q' to Quit.
```

```
[Default: c]:
```

```
stopping: Drive in "FCloop" slot 2....Done
```

```
offlining: Drive in "FCloop" slot 2....Done
```

```
Hit <Return> after removing the device(s).
```

- 6 Press **Enter** again to quit.
7 Go to *Replace a Hard Disk Drive (6-Disk Setup)* (on page 13).

Replace a Hard Disk Drive (6-Disk Setup)

Installing a New Hard Disk


If you have any questions about the equipment, refer to your copy of the *Owner's Guide* before proceeding with these instructions.



CAUTION:

Static sensitive device. Always wear a properly grounded wrist strap when working on this equipment. Failure to observe this caution may result in equipment damage or premature equipment failure.

Use the instructions in this section to install a new hard disk into your V880 or V890. Refer to the following illustration for guidance as you complete this procedure.



T11323

Note: Use only Sun Microsystems Inc. (Sun) internal 72 GB, 10000 RPM, 1-in. high, UltraSCSI disk drives (Sun part number 5234A).

- 1 Swing open the disk drive bay door on the front of the V880/V890. Three columns of internal drive bays are exposed.
- 2 Remove the new hard disk from its protective packaging.
- 3 Slide the thumb latch UP on the old hard disk. The silver drive handle on the hard disk swings down.



CAUTION:

Allow approximately 30 seconds for the hard disk drive to spin down completely before removing the drive from the drive bay. Failure to observe this caution may result in equipment damage.

- 4 Holding the silver drive handle, remove the old hard disk from the drive bay.
- 5 Complete the following steps to install the new hard disk.
 - a Orient the new hard disk so that the silver drive handle is on the bottom.
 - b Using the drive bay guide rails, slide the new disk drive *halfway* into the drive bay.

Important: Do not slide the new disk drive completely into the drive bay at this time. You must put the new drive in the online state first before you can complete the disk installation.

Note: In the event that the new disk drive arrives without guide rails, attach the guide rails from the old disk drive to the new disk drive.

- 6 Go to *Putting the New Disk Drive Into an Online State* (on page 16).

Putting the New Disk Drive Into an Online State

Before you can complete the new disk drive installation, you have to put the new disk drive into an online state.



CAUTION:

Static sensitive device. Always wear a properly grounded wrist strap when working on this equipment. Failure to observe this caution may result in equipment damage or premature equipment failure.

Complete these steps to put the new disk drive in the online state.

Note: If you have been following the procedures in this chapter in order, you should still be logged in to an xterm window as root user.

- 1 Type the following command and press **Enter**:
`luxadm insert_device`
- 2 Press **Enter** again.
- 3 Complete the following steps to install the new hard disk.
 - a Slide the hard disk completely into the drive bay and seat it securely against the back plane of the server.
 - b Swing the silver drive handle closed until it latches.

Note: When you latch the silver drive handle, it engages the hard disk with its back plane connector.
- 4 Is the green LED for the newly installed hard disk illuminated?

- If **yes**, you have installed the hard disk correctly. Go to step 5.
 - If **no**, contact Cisco Services for help with troubleshooting the hard disk installation.
- 5 Close the drive bay door on the front of the server.
 - 6 Press **Enter**. The computer incorporates the new disk drive into the operating system.
- Note:** This process takes a couple of minutes to complete.

- 7 Type the following command and press **Enter**:

```
luxadm display [enclosure name]
```

Example: `luxadm display FCloop`

Result: The disk status appears.

```

SUNWGS INT FCBPL
          DISK STATUS
SLOT   DISKS                (Node WWN)
0      On (O.K.)            2000002037f87099
1      On (O.K.)            2000002037f39796
2      On (Login ok)
3      On (O.K.)            2000002037ffab19
4      On (O.K.)            2000002037ff9ca5
5      On (O.K.)            2000002037ffa3c8
6      Not Installed
7      Not Installed
8      Not Installed
9      Not Installed
10     Not Installed
11     Not Installed

          SUBSYSTEM STATUS
FW Revision:9218   Box ID:0
Node WWN:508002000011d7d8   Enclosure Name:FCloop
SSC100's - 0=Base Bkpln, 1=Base LoopB, 2=Exp Bkpln, 3=Exp
LoopB
SSC100 #0:   O.K.(11.A)
SSC100 #1:   O.K.(11.A)
SSC100 #2:   Not Installed
SSC100 #3:   Not Installed

          Temperature Sensors - 0 Base, 1 Expansion
0:21°C
          1:Not Installed

```

Note: In this example, notice the status of slot 2 changed from **Login failed** to **Login ok**, indicating that Slot 2 is the disk drive that was replaced.

- 8 Does the system list slots 0 through 5?
- If **yes**, the system recognizes the newly installed hard disk; go to step 9.
 - If **no**, contact Cisco Services for assistance.

- 9 Type the following command and press **Enter**:

```
prtvtoc /dev/rdisk/[device of drive mirror] > /tmp/[device
of drive mirror].vtoc
```

Example: `prtvtoc /dev/rdisk/c2t5d0s2 > /tmp/c2t5d0.vtoc`

Note: Refer to the *V880/V890 6-Disk Diagram and Configuration Table* (on page 22) to determine the drive mirror. For example, c2t5 is mirrored to c1t2.

Result: The mirrored drive format is copied to use on the new drive.

- 10 Type the following command and press **Enter**:

```
fmthard -s /tmp/[mirrored device name].vtoc  
/dev/rdisk/[device name of replaced drive]
```

Example: `fmthard -s /tmp/c2t5d0.vtoc /dev/rdisk/c1t2d0s2`

Result: The replacement hard disk is formatted. The replacement hard disk is c1t2d0s2.

Note: If formatting the hard drive is unsuccessful, contact Cisco Services.

- 11 Go to *Add a Meta Database to the New Drive (6-Disk Setup)* (on page 20).

Add a Meta Database to the New Drive (6-Disk Setup)

Adding a Meta Database to the Disk Drive

A V880 or a V890 server with a 6-disk setup utilizes meta databases on each disk drive. Because the defective drive has a meta database, you must add a meta database to the new drive. Use the following instructions to add a meta database to the new disk drive.

Note: If you have been following the procedures in this chapter in order, you should still be logged in to an xterm window as root user.

- 1 Type the following command and press **Enter**:

```
metadb -a [new disk drive]
```

Example: `metadb -a c1t2d0s7`

Note: The new disk drive name is the device name that you recorded on a separate piece of paper during the *Deleting a Meta Database* (on page 11) procedure.

Result: The meta database is added to the new disk drive.

- 2 Go to *Using the Meta Replace Utility to Resynchronize Mirrors* (on page 20).

Using the Meta Replace Utility to Resynchronize Mirrors

Use the following instructions to resynchronize the mirrors on the new disk drive.

Note: This step could take from 1 to 7 hours to complete, based on your system.

- 1 Type the following command and press **Enter**:

```
metareplace -e [mirror device] [disk device]
```

Example: `metareplace -e d520 c1t2d0s0`

Result: The mirrors on the new disk drive are synchronized.

Important: You must run the `metareplace -e` command for every disk slice on the disk drive that you replaced. Refer to the *V880/V890 Meta Replace Commands* (on page 23) table to determine the `metareplace -e` commands you

must run.

- 2 Go to *Using the Meta Check Utility to Monitor Disk Activity* (on page 21).

Using the Meta Check Utility to Monitor Disk Activity

Performing a meta replace results in extensive disk activity. After you resynchronize the mirrors, run the **check_metadevices** command to monitor the disk activity.

Note: The check_metadevices utility automatically runs upon reload of the DBDS.

Go to *Install a Boot Block on a Mirrored Disk Drive (6-Disk Setup)* (on page 21).

Install a Boot Block on a Mirrored Disk Drive (6-Disk Setup)

Installing a Boot Block on a Mirrored Disk Drive

If the replaced disk is an OS disk, you must install a boot block on the hard drive to make the hard drive bootable.

Note: A boot block is a portion of computer code that makes the hard drive on which it is installed able to start itself.

Use the following instructions to install a boot block on a mirrored disk drive of the V880 or V890 server. You should still be logged in to an xterm window as root user.

- 1 Type the following command and press **Enter**:

```
installboot /usr/platform/sun4u/lib/fs/ufs/bootblk
/dev/rdisk/[replaced drive device name]
```

Example: `installboot /usr/platform/sun4u/lib/fs/ufs/bootblk /dev/rdisk/clt0d0s0`

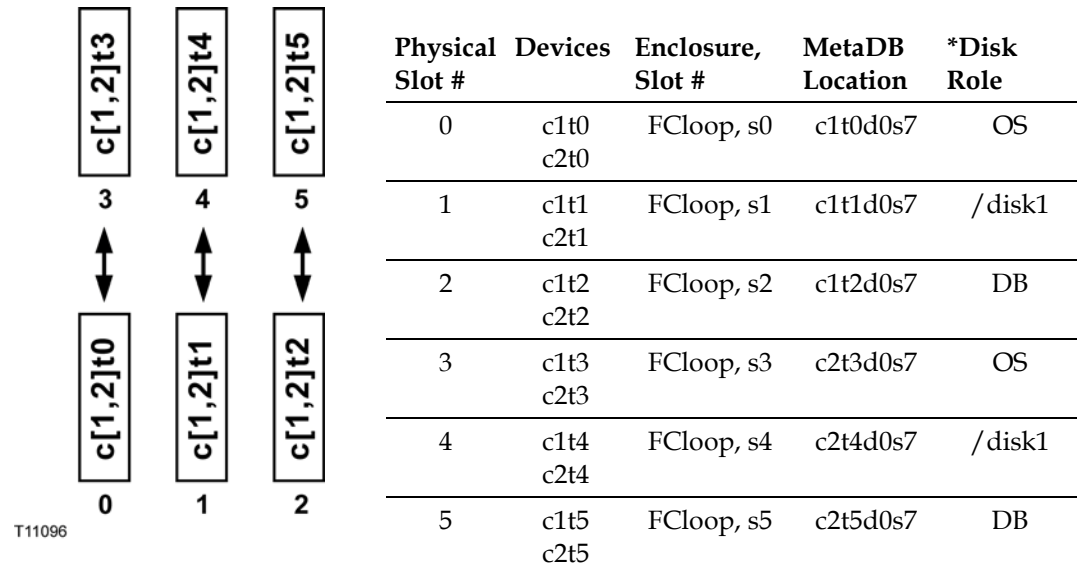
Result: The system installs the boot block on the replaced disk drive.

- 2 Type **exit** and then press **Enter** to log out the root user.

V880/V890 6-Disk Configuration Details

V880/V890 6-Disk Diagram and Configuration Table

The following diagram and table describe the configuration of a 6-disk V880/V890.



Notes:

- In the diagram, arrows indicate the mirrored disks (0 to 3, 1 to 4, and 2 to 5).
- *Disk Role:
 - **OS** is the Solaris operating system drive
 - **/disk 1** is file system containing the Cisco software
 - **DB** is the Informix database drive

V880/V890 6-Disk Device Assignments

The following table describes the metadevice assignments to the physical disk devices for the 6-disk V880 or V890 server. This information is provided as a reference tool for understanding the disk slice assignments.

Purpose	Metadevice	Submirrors	Physical Device
/	d500	d400 d700	c1t0d0s0 c2t3d0s0
swap	d501	d401 d701	c1t0d0s1 c2t3d0s1
/var	d503	d403 d703	c1t0d0s3 c2t3d0s3

Purpose	Metadevice	Submirrors	Physical Device
/export/home	d507	d407 d707	c1t0d0s5 c2t3d0s5
/disk1	d510	d410 d710	c1t1d0s0 c2t4d0s0
soft partitions	d520	d420 d720	c1t2d0s0 c2t5d0s0

V880/V890 Meta Replace Commands

The following table contains the `metareplace -e` commands for each disk slice on the 6-disk V880 or V890 server. Use the information in this table when performing the procedure to resynchronize mirrors (see *Using the Meta Replace Utility to Resynchronize Mirrors* (on page 20)).

Important: You must run the `metareplace -e` command for *every* disk slice on the disk drive that you replaced.

Disk Replaced	Meta Replace Command
Slot 0	metareplace -e d500 c1t0d0s0 metareplace -e d501 c1t0d0s1 metareplace -e d503 c1t0d0s3 metareplace -e d507 c1t0d0s5
Slot 1	metareplace -e d510 c1t1d0s0
Slot 2	metareplace -e d520 c1t2d0s0
Slot 3	metareplace -e d500 c2t3d0s0 metareplace -e d501 c2t3d0s1 metareplace -e d503 c2t3d0s3 metareplace -e d507 c2t3d0s5
Slot 4	metareplace -e d510 c2t4d0s0
Slot 5	metareplace -e d520 c2t5d0s0


12-Disk Configuration

V880/V890 Hardware (12-Disk Configuration)

Hard Disk Layout

Depending on your system, the V880/V890 is configured with either 6 or 12 disk drives.

The following illustration shows the front view of the 12-disk V880 and V890 server layout.



Verify Disk Drive Status (12-Disk Setup)

If a disk problem occurs on your V880 or V890 server, an orange banner with white text (banner window) appears on your console to alert you to the problem. Do not

close this window; the window disappears when the problem is resolved. The disk problem may be an event where one or all portions of a disk drive become inaccessible.

Should one or all portions of a disk drive become inaccessible, you must use the **luxadm** utility to accomplish the following steps:

- Verify the enclosure name assigned to the V880/V890 internal storage array
- Determine the status of the disk drive

Note: Cisco's default enclosure name is FCloop.

After using the **luxadm** utility, use the **metadb** utility to determine if a meta database is on the disk drive. If a meta database exists on the drive, you must delete the meta database *prior* to removing the disk drive.

Finding the Enclosure Name

Complete these steps to verify the enclosure name for the internal storage array using the **luxadm probe** command.

- 1 If necessary, open an xterm window on the V880 or V890 server.
- 2 Complete the following steps to log on to the xterm window as **root** user.
 - a Type **su -** and press **Enter**. The password prompt appears.
 - b Type the root password and press **Enter**.
- 3 Type the following command and then press **Enter**. The name of the enclosure appears.

```
luxadm probe
```

Example: Sample output follows:

```
Found Enclosure:
```

```
SUNWGS INT FCBPL   Name:FCloop   Node WWN:508002000011d7d8
```

```
    Logical Path:/dev/es/ses0
```

```
    Logical Path:/dev/es/ses1
```

Note: **FCloop** is the enclosure name in this example.

Determining the Status of the Disk Drive

Complete these steps to determine the disk drive status using the **luxadm display** command.

Note: If you have been following the procedures in this chapter in order, you should still be logged in to an xterm window as root user.

- 1 Type the following command and press **Enter**:

```
luxadm display [enclosure name]
```

Example: `luxadm display FCloop`

Result: The disk status appears.

```
SUNWGS INT FCBPL
                                DISK STATUS
SLOT   DISKS                    (Node WWN)
0      On (O.K.)                2000002037ffa51d
1      On (O.K.)                2000002037ff9c4a
2      On (Login failed)
3      On (O.K.)                2000002037ffab19
4      On (O.K.)                2000002037ff9ca5
5      On (O.K.)                2000002037ffa3c8
6      On (O.K.)                20000000870ebeda
7      On (O.K.)                20000000870f9e8f
8      On (O.K.)                20000000870f9d47
9      On (O.K.)                20000000870f9eb6
10     On (O.K.)                20000000870e3397
11     On (O.K.)                20000000870ea2ad
                                SUBSYSTEM STATUS
FW Revision:9226   Box ID:0
  Node WWN:50800200001ac508   Enclosure Name:FCloop
SSC100's - 0=Base Bkpln, 1=Base LoopB, 2=Exp Bkpln, 3=Exp
LoopB
  SSC100 #0:   O.K.(9226/ FD99)
  SSC100 #1:   O.K.(9226/ FD99)
  SSC100 #2:   O.K.(9226/ FD99)
  SSC100 #3:   O.K.(9226/ FD99)
          Temperature Sensors - 0 Base, 1 Expansion
          0:26°C 1:24°C (All temperatures are NORMAL.)
Default Language is USA English, ASCII
```

Note: In this example, notice the status of slot 2 is Login failed. This indicates that Slot 2 is the disk drive that needs to be replaced.

Important: The World Wide Name (WWN) is unique for each disk drive. Whenever you remove a drive and replace it, you get a new WWN value for the

new drive.

- 2 Go to *Examine Mirrored Devices (12-Disk Setup)* (on page 23).

Examine Mirrored Devices (12-Disk Setup)

The V880 and V890 servers support disk drive mirroring. Through disk mirroring, the DNCS stores identical information across sets of hard disks.

After using the **luxadm display** command to determine the drive status, use the **check_metadevices** command to determine if a single file system or multiple file systems have been affected.

The V880/V890 12-disk configuration has a total of six mirrored disk drives. The hard disks in drive bay 0 and drive bay 8 are mirrored, as are the hard disks in drive bay 1 and drive bay 9. The hard disks in drive bays 2, 3, 4, 5 and drive bays 10, 11, 12, 13, respectively, are also mirrored.

Before you disable the disk mirroring functions of the V880 or V890 server in preparation of replacing a disk drive, you should examine the status of the mirrored disks in the system. Examining the status of the mirrored devices also tells you if a specific disk slice needs maintenance.

Examining the Mirrored Devices

Complete these steps to examine the status of the mirrored disk drives on your DNCS using the **check_metadevices** command.

Note: If you have been following the procedures in this chapter in order, you should still be logged in to an xterm window as root user.

- 1 Type the following command and press **Enter**:

```
su - dnscs
```

- 2 Type the following command and press **Enter**. The metadvice tool examines the mirrored disk drives on your DNCS and indicates if a drive or a specific disk slice needs maintenance.

```
/dvs/dnscs/bin/check_metadevices
```

Example: Sample output pertaining to when a disk slice needs maintenance follows:

```
WARNING: d520 metadvice Needs maintenance
```

```
Found Disk Problems at: Mon Jul 12 11:08:20 EDT 2004
```

Important: Write the name of the disk drive or disk slice that needs maintenance on a separate sheet of paper.

- 3 Type **exit** and then press **Enter** to exit from the DNCS shell.

What Is the Status of the Mirrored Disk Drives?

After using the `check_metadevices` command to examine the status of mirrored disk drives, the following scenarios could occur:

- If one slice of the disk needs maintenance, resync the disk slice. Go to *Resync the Disk Drive* (on page 25).
- If you have errors on multiple slices per disk, complete the following steps to verify the disk drive status.
 - a Go back to *Verify Disk Drive Status (12-Disk Setup)* (on page 24) and run the `luxadm probe` command again.
 - b Choose one of the following options based on the results of the **luxadm probe** command:
 - If this command does not show a disk drive problem, contact Cisco Services for assistance.
 - If this command shows login failed <drive name>, you must replace the disk drive. Go to *Remove a Hard Disk Drive (12-Disk Setup)* (on page 27).

Important: If maintenance is required on drives that are mirrors of each other, such as `c1t2d0s7` and `c2t5d0s0`, contact Cisco Services.

Resync the Disk Drive

After examining the mirrored devices, when one disk slice needs maintenance, you must resync the disk drive.

Note: The resync process could take from a half-hour to 7 hours to complete, based on your system.

Complete these steps to resync the disk drive.

- 1 Using the disk slice you recorded in the *Examining the Mirrored Devices* (on page 28) procedure, type the following command and press **Enter**:

```
metastat d[disk slice]
```

Example: `metastat d520`

Result: The submirrors of d520 are listed.

```
d520: Mirror
  Submirror 0: d420
    State: Needs maintenance
  Submirror 1: d720
    State: Okay
  Pass: 1
  Read option: roundrobin (default)
  Write option: parallel (default)
  Size: 125256384 blocks

d420: Submirror of d520
  State: Needs maintenance
  Invoke: metareplace d520 c1t2d0s0 <new device>
  Size: 125256384 blocks
  Stripe 0:
    Device      Start Block  Dbase State      Hot Spare
    c1t2d0s0      0           No   Maintenance

d720: Submirror of d520
  State: Okay
  Size: 125256384 blocks
  Stripe 0:
    Device      Start Block  Dbase State      Hot Spare
    c2t10d0s0    0           No   Okay
```

Notes:

- In this example, notice the state of Submirror 1: d420 is **Needs Maintenance**. This indicates that d420 is the submirror of d520 that requires maintenance.
- Also in this example, notice that the Device assignment for d420 is **c1t2d0s0**.

You will use this value when you perform the meta replace command in the next step.

- Refer to the *V880/V890 12-Disk Device Assignments* (on page 46) table for more information.
- 2 Locate the disk slice in the *V880/V890 Meta Replace Commands* (on page 47) table and type the meta replace command for the disk slice.
Example: In the previous example, the device assignment for d420 is c1t2d0s0. Type the following to resync the disk drive:

```
metareplace -e d520 c1t2d0s0
```


Important: Type the meta replace command that affects the specific disk slice, only.
Result: You return to a prompt, and the orange banner window shows the disk drive is resyncing.
 - 3 After you resync the drive, go back to the *Examining Mirrored Devices* (on page 28) procedure and re-examine the mirrored devices using the check metadevices tool.
 - 4 Did the metadvice tool indicate that the disk slice needs maintenance?
 - If **yes**, you must replace the disk drive. Go to *Remove a Hard Disk Drive (12-Disk Setup)* (on page 27).
 - If **no**, you do not need to replace the disk drive. However, keep a written record of the disk slice on which you performed the meta replace command. If a problem continues to occur with the same disk slice, you may consider replacing the disk drive. No further action is required. You may stop reading this guide.

Remove a Hard Disk Drive (12-Disk Setup)

Finding a Meta Database

After examining the mirrored disk drives and determining that multiple portions of a disk drive are inaccessible, you must remove the hard disk drive. Before you can remove the drive, use the **metadb** utility to determine if the drive in question has a metadvice.

Note: For more information about metadvice assignments, refer to the *V880/V890 12-Disk Device Assignments* (on page 46) table.

Complete these steps using the metadb utility to determine if a meta database is located on the disk drive. If you have been following the procedures in this chapter in order, you should still be logged in to an xterm window as root user.

- 1 Type the following command and press **Enter** to display the status of the disk drive:

```
metadb -i
      flags          first blk      block count
      a m p luo      16          1034
/dev/dsk/c1t0d0s7
      a   p luo      16          1034
/dev/dsk/c1t1d0s7
      a R p luo      16          1034
/dev/dsk/c1t2d0s7
      a   p luo      16          1034
/dev/dsk/c2t8d0s7
      a   p luo      16          1034
/dev/dsk/c2t9d0s7
      a   p luo      16          1034
/dev/dsk/c2t10d0s7
o - replica active prior to last mddb configuration change
u - replica is up to date
l - locator for this replica was read successfully
c - replica's location was in /etc/lvm/mddb.cf
p - replica's location was patched in kernel
m - replica is master, this is replica selected as input
W - replica has device write errors
a - replica is active, commits are occurring to this
replica
M - replica had problem with master blocks
D - replica had problem with data blocks
F - replica had format problems
```


S - replica is too small to hold current data base

R - replica had device read errors

- 2 The drive flag with a capital letter indicates which meta database file has an error. In this example, **c1t2d0s7** is the drive with the meta database error.
- 3 Refer to the *V880/V890 12-Disk Diagram and Configuration Table* (on page 45) to verify the slot number matches the drive with the meta database error.

Example: Using the configuration table, the drive `/dev/disk/c1t2d0s7` is located in physical slot 2.

- 4 Based upon the information in the **V880/V890 12-Disk Diagram and Configuration Table**, did the meta database match the slot number of the hard disk drive?
 - If **yes**, you must delete the meta database from the hard disk drive prior to removing the drive. Go to *Deleting a Meta Database* (on page 33).
 - If **no**, contact Cisco Services for assistance.

Deleting a Meta Database

Complete these steps to delete a meta database on the hard disk drive.

Note: If you have been following the procedures in this chapter in order, you should still be logged in to an xterm window as root user.

- 1 Type the following command and press **Enter** to delete the meta database from the disk drive:

```
metadb -d [device name]
```

Example: `metadb -d c1t2d0s7`

Note: Write the device name that you delete on a separate piece of paper. You will use this name when you add a meta database to the disk drive, later in this chapter.

- 2 Before you physically remove the hard disk drive, you must put the drive into an offline state. Go to *Putting a Hard Disk Drive Into an Offline State* (on page 29).

Putting the Hard Disk Drive Into an Offline State

When the status of the disk drive is anything other than not installed, you must put the drive into the offline state.

Complete these steps to put the disk drive into the offline state.

Note: If you have been following the procedures in this chapter in order, you should still be logged in to an xterm window as root user.

- 1 Before you remove the drive, Cisco highly recommends that you have a recent backup of the drive data. Refer to your current backup and restore guide for backup procedures.

Note: Refer to the *V880/V890 12-Disk Diagram and Configuration Table* (on page 45) to determine the disk role for the drive you are removing.

Examples:

- An error was found on drive **c1t0d0s7**. Based on the information in the 12-Disk configuration table, the role of drive **c1t0d0s7** is **/disk1**, meaning it has a file system role. Because a file system contains the Cisco software, you would perform a file system backup for this drive.
- If an error is found on a drive with a role of DB Stripe, this means that a drive with an Informix database role has an error. You would perform a database backup on the drive.

- 2 After backing up the drive data, type the following command and press **Enter**:

```
luxadm remove_device [enclosure name],[slot]
```

Example: `luxadm remove_device FCloop,s2`

Result: A warning message appears prompting you to be sure all data on the disk drive is backed up.

```
WARNING!!! Please ensure that no filesystems are mounted on
these device(s). All data on these devices should have been
backed up.
```

The list of devices being used (either busy or reserved) by the host:

```
1: Box Name:      "FCloop"  slot 2
Please enter 's' or <CR> to Skip the "busy/reserved"
device(s) or
'q' to Quit and run the subcommand with
-F (force) option. [Default: s]: q
```

- 3 Type **q** to quit and then press **Enter**.

- 4 Type the following command and press **Enter**:

```
luxadm remove_device -F [enclosure name],[slot]
```

Example: `luxadm remove_device -F FCloop,s2`

Result: The WWN entries and device links are removed for the disk drive and you are prompted to verify the disk drive to remove.

WARNING!!! Please ensure that no filesystems are mounted on these device(s). All data on these devices should have been backed up.

The list of devices which will be removed is:

```
1: Box Name:      "FCloop" slot 2
   Node WWN:      20000002037ff9bee
   Device Type:   Disk device
   Device Paths:
     /dev/rdisk/c1t2d0s7
     /dev/rdisk/c2t2d0s2
```

- 5 Press **Enter**. The system prompts you to remove the disk drive. Please verify the above list of devices and then enter 'c' or <CR> to Continue or 'q' to Quit. [Default: c]:
stopping: Drive in "FCloop" slot 2....Done
offlining: Drive in "FCloop" slot 2....Done
Hit <Return> after removing the device(s).
- 6 Press **Enter** again to quit.
- 7 Go to *Replace a Hard Disk Drive (12-Disk Setup)* (on page 31).

Replace a Hard Disk Drive (12-Disk Setup)

Installing a New Hard Disk


If you have any questions about the equipment, refer to your copy of the *Sun Fire 880 Server Owner's Guide* or *Sun Fire 890 Server Owner's Guide* before proceeding with these instructions.



CAUTION:

Static sensitive device. Always wear a properly grounded wrist strap when working on this equipment. Failure to observe this caution may result in equipment damage or premature equipment failure.

Use the instructions in this section to install a new hard disk into your V880/V890. Refer to the following illustration for guidance as you complete this procedure.



T11323

Note: Use only Sun Microsystems Inc. (Sun) internal 72 GB, 10000 RPM, 1-in. high, UltraSCSI disk drives (Sun part number 5234A).

- 1 Swing open the disk drive bay door on the front of the server. Three columns of internal drive bays are exposed.
- 2 Remove the new hard disk from its protective packaging.
- 3 Slide the thumb latch UP on the old hard disk. The silver drive handle on the hard disk swings down.



CAUTION:

Allow approximately 30 seconds for the hard disk drive to spin down completely before removing the drive from the drive bay. Failure to observe this caution may result in equipment damage.

- 4 Holding the silver drive handle, remove the old hard disk from the drive bay.
- 5 Complete the following steps to install the new hard disk.
 - a Orient the new hard disk so that the silver drive handle is on the bottom.
 - b Using the drive bay guide rails, slide the new disk drive *halfway* into the drive bay.

Important: Do not slide the new disk drive completely into the drive bay at this time. You must put the new drive in the online state first before you can complete the disk installation.

Note: In the event that the new disk drive arrives without guide rails, attach the guide rails from the old disk drive to the new disk drive.

- 6 Go to *Putting the New Disk Drive Into an Online State* (on page 38).

Putting the New Disk Drive Into an Online State

Before you can complete the new disk drive installation, you have to put the new disk drive into an online state.



CAUTION:

Static sensitive device. Always wear a properly grounded wrist strap when working on this equipment. Failure to observe this caution may result in equipment damage or premature equipment failure.

Complete these steps to put the new disk drive in the online state.

Note: If you have been following the procedures in this chapter in order, you should still be logged in to an xterm window as root user.

- 1 Type the following command and press **Enter**:
`luxadm insert_device`
- 2 Press **Enter** again.
- 3 Complete the following steps to install the new hard disk.
 - a Slide the hard disk completely into the drive bay and seat it securely against the back plane of the server.
 - b Swing the silver drive handle closed until it latches.

Note: When you latch the silver drive handle, it engages the hard disk with its back plane connector.
- 4 Is the green LED for the newly installed hard disk illuminated?

- If **yes**, you have installed the hard disk correctly. Go to step 5.
 - If **no**, contact Cisco Services for help with troubleshooting the hard disk installation.
- 5 Close the drive bay door on the front of the server.
 - 6 Press **Enter**. The computer incorporates the new disk drive into the operating system.

Note: This process takes a couple of minutes to complete.

- 7 Type the following command and press **Enter**:

```
luxadm display [enclosure name]
```

Example: `luxadm display FCloop`

Result: The disk status appears.

```
SUNWGS INT FCBPL
          DISK STATUS
SLOT    DISKS                (Node WWN)
0       On (O.K.)            2000002037f87099
1       On (O.K.)            2000002037f39796
2       On (Login ok)
3       On (O.K.)            2000002037ffab19
4       On (O.K.)            2000002037ff9ca5
5       On (O.K.)            2000002037ffa3c8
6       On (O.K.)            20000000870ebeda
7       On (O.K.)            20000000870f9e8f
8       On (O.K.)            20000000870f9d47
9       On (O.K.)            20000000870f9eb6
10      On (O.K.)            20000000870e3397
11      On (O.K.)            20000000870ea2ad
          SUBSYSTEM STATUS
FW Revision:9218   Box ID:0
          Node WWN:508002000011d7d8   Enclosure Name:FCloop
SSC100's - 0=Base Bkpln, 1=Base LoopB, 2=Exp Bkpln, 3=Exp
LoopB
          SSC100 #0:    O.K.(11.A)
          SSC100 #1:    O.K.(11.A)
          SSC100 #2:    Not Installed
          SSC100 #3:    Not Installed

          Temperature Sensors - 0 Base, 1 Expansion
          0:21°C
          1:Not Installed
```

Note: In this example, notice the status of slot 2 changed from Login failed to **Login ok**, indicating that Slot 2 is the disk drive that was replaced.

- 8 Does the system list disks 0 through 11?
- If **yes**, the system recognizes the newly installed hard disk; go to step 9.
 - If **no**, contact Cisco Services for assistance.

- 9 Type the following command and press **Enter**:

```
prtvtoc /dev/rdisk/<device of drive mirror> > /tmp/[device  
of drive mirror].vtoc
```

Example: `prtvtoc /dev/rdisk/c2t10d0s2 > /tmp/c2t10d0.vtoc`

Note: Refer to the *V880/V890 12-Disk Diagram and Configuration Table* (on page 45) to determine the drive mirror. For example, c2t9 is mirrored to c1t1.

Result: The mirrored drive format is copied to use on the new drive.

- 10 Type the following command and press **Enter**:

```
fmthard -s /tmp/[mirrored device name].vtoc  
/dev/rdisk/[device name of replaced drive]
```

Example: `fmthard -s /tmp/c2t10d0.vtoc /dev/rdisk/c1t2d0s2`

Result: The replacement hard disk is formatted. The replacement hard disk is **c1t2d0s2**.

Note: If formatting the hard drive is unsuccessful, contact Cisco Services.

- 11 Did the hard disk drive have a meta database?

- If **yes**, go to *Add a Meta Database to the New Drive (12-Disk Setup)* (on page 42).
- If **no**, go to the *Using the Meta Replace Utility to Resynchronize Mirrors* (on page 35).

Add a Meta Database to the New Drive (12-Disk Setup)

A V880 or a V890 server with a 12-disk setup has meta databases in slot numbers 0, 1, 2, 6, 7, and 8, only. If the defective drive has a meta database, you must add a meta database to the new drive.

Note: If the new disk drive did not have a meta database, go to *Using the Meta Replace Utility to Resynchronize Mirrors* (on page 35).

Adding a Meta Database to the Disk Drive

Use the following instructions to add a meta database to the new disk drive.

Note: If you have been following the procedures in this chapter in order, you should still be logged in to an xterm window as root user.

- 1 Type the following command and press **Enter**:

```
metadb -a [new disk drive]
```

Example: `metadb -a c1t2d0s7`

Note: The new disk drive name is the device name that you recorded on a separate piece of paper during the *Deleting a Meta Database* (on page 33)

procedure.

Result: The meta database is added to the new disk drive.

- 2 Go to *Using the Meta Replace Utility to Resynchronize Mirrors* (on page 35).

Using the Meta Replace Utility to Resynchronize Mirrors

Use the following instructions to resynchronize the mirrors on the new disk drive.

Note: This step could take from 1 to 7 hours to complete, based on your system.

- 1 Type the following command and press **Enter**:

```
metareplace -e [mirror device] [disk device]
```

Example: `metareplace -e d520 c1t2d0s0`

Result: The mirrors on the new disk drive are synchronized.

Important: You must run the `metareplace -e` command for every disk slice on the disk drive that you replaced. Refer to the *V880/V890 Meta Replace Commands* (on page 47) table to determine the `metareplace -e` commands you must run.

- 2 Go to *Using the Meta Check Utility to Monitor Disk Activity* (on page 44).

Using the Meta Check Utility to Monitor Disk Activity

Performing a meta replace results in extensive disk activity. After you resynchronize the mirrors, perform `check_metadevices` to monitor the disk activity.

Note: The `check_metadevices` utility automatically runs upon reload of the DBDS.

Go to *Install a Boot Block on a Mirrored Disk Drive (12-Disk Setup)* (on page 44).

Install a Boot Block on a Mirrored Disk Drive (12-Disk Setup)

Installing a Boot Block on a Mirrored Disk Drive

If the replaced disk is an OS disk, you must install a boot block on the hard drive to make the hard drive bootable.

Note: A boot block is a portion of computer code that makes the hard drive on which it is installed able to start itself.

Use the following instructions to install a boot block on a mirrored disk drive of the V880 or V890 server. You should still be logged in to an xterm window as root user.

- 1 Type the following command and press **Enter**:

```
installboot /usr/platform/sun4u/lib/fs/ufs/bootblk  
/dev/rdisk/[replaced drive device name]
```

Example: `installboot /usr/platform/sun4u/lib/fs/ufs/bootblk
/dev/rdisk/c1t0d0s0`


Result: The system installs the boot block on the replaced disk drive.

- 2 Type `exit` and then press **Enter** to log out the root user.

V880/V890 12-Disk Configuration Details

V880/V890 12-Disk Diagram and Configuration Table

The following diagram and table describe the configuration of a 12-disk V880 and V890 server.



Physical Slot #	Devices	Enclosure, Slot #	MetaDB Location	* Disk Role
0	c1t0 c2t0	FCloop, s0	c1t0d0s7	OS
1	c1t1 c2t1	FCloop, s1	c1t1d0s7	/disk1
2	c1t2 c2t2	FCloop, s2	c1t2d0s7	DB Stripe
3	c1t3 c2t3	FCloop, s3	N/A	DB Stripe
4	c1t4 c2t4	FCloop, s4	N/A	DB Stripe
5	c1t5 c2t5	FCloop, s5	N/A	DB Stripe
6	c1t8 c2t8	FCloop, s6	c2t8d0s7	OS
7	c1t9 c2t9	FCloop, s7	c2t9d0s7	/disk1
8	c1t10 c2t10	FCloop, s8	c2t10d0s7	DB Stripe
9	c1t11 c2t11	FCloop, s9	N/A	DB Stripe
10	c1t12 c2t12	FCloop, s10	N/A	DB Stripe
11	c1t13 c2t13	FCloop, s11	N/A	DB Stripe

Notes:

- In the diagram, arrows indicate the mirrored disks (0 to 8, 1 to 9, and 2, 3, 4, 5 to 10, 11, 12, 13).
- *Disk Role:
 - **OS** is the Solaris operating system drive
 - **/disk 1** is file system containing the Cisco software

- **DB Stripe** is one of four disks that makeup the Informix database

V880/V890 12-Disk Device Assignments

The following table describes the metadevice assignments to the physical disk devices for the 12-disk V880/V890. This information is provided as a reference tool for understanding the disk slice assignments.

Purpose	Metadevice	Submirrors	Physical Device
/	d500	d400	c1t0d0s0
		d700	c2t8d0s0
swap	d501	d401	c1t0d0s1
		d701	c2t8d0s1
/var	d503	d403	c1t0d0s3
		d703	c2t8d0s3
/export/home	d507	d407	c1t0d0s5
		d707	c2t8d0s5
/disk1	d510	d410	c1t1d0s0
		d710	c2t9d0s0
soft partitions	d520	d420	c1t2d0s0
			c1t3d0s0
			c1t4d0s0
			c1t5d0s0
			d720
			c2t11d0s0
			c2t12d0s0
			c2t13d0s0

V880/V890 Meta Replace Commands

The following table contains the *metareplace -e* commands for each disk slice on the 12-disk V880 or V890 server. Use the information in this table when performing the resynchronize mirrors procedure. See *Using the Meta Replace Utility to Resynchronize Mirrors* (on page 35) for details.

Important: You must run the *metareplace -e* command for *every* disk slice on the disk drive that you replaced.

Disk Replaced	Meta Replace Command
Slot 0	metareplace -e d500 c1t0d0s0 metareplace -e d501 c1t0d0s1 metareplace -e d503 c1t0d0s3 metareplace -e d507 c1t0d0s5
Slot 1	metareplace -e d510 c1t1d0s0
Slot 2	metareplace -e d520 c1t2d0s0
Slot 3	metareplace -e d520 c1t3d0s0
Slot 4	metareplace -e d520 c1t4d0s0
Slot 5	metareplace -e d520 c1t5d0s0
Slot 6	metareplace -e d500 c2t8d0s0 metareplace -e d501 c2t8d0s1 metareplace -e d503 c2t8d0s3 metareplace -e d507 c2t8d0s5
Slot 7	metareplace -e d510 c2t9d0s0
Slot 8	metareplace -e d520 c2t10d0s0
Slot 9	metareplace -e d520 c2t11d0s0
Slot 10	metareplace -e d520 c2t12d0s0
Slot 11	metareplace -e d520 c2t13d0s0

2

Replace a Disk Drive on a Sun Fire V445 Server

Introduction

This chapter provides instructions for removing and replacing a hard disk drive on the Sun Fire V445 server.

In This Chapter

■ Hard Disk Layout	50
■ Verify Disk Drive Status	51
■ Prepare to Replace the Hard Drive	62
■ Avoid Electrostatic Discharge	65
■ Remove the Hard Drive	66
■ Install the Hard Drive	68
■ Reconnect the New Hard Drive	69

Hard Disk Layout

Depending on your system, the Sun Fire V445 server is configured with either 4 or 8 disk drives. For the purpose of this document, the logical drive layout on the 4-disk and 8-disk systems are the same with the following exception: slots 2, 3, 6, and 7 on the 8-disk server are used as extra storage for database slices.

Physical Slot	Device	metadb	Disk Role	Comment
0	c1t0	c1t0d0s4	OS	
1	c1t1	c1t1d0s4	DB	
2	c1t2	c1t2d0s4	DB	8-disk configuration only
3	c1t3	c1t3d0s4	DB	8-disk configuration only
4	c1t4	c1t4d0s4	OS	
5	c1t5	c1t5d0s4	DB	
6	c1t6	c1t6d0s4	DB	8-disk configuration only
7	c1t7	c1t7d0s4	DB	8-disk configuration only

Verify Disk Drive Status

Follow these instructions to determine whether one or more metadevices have failed on the system.

- 1 If necessary, open an xterm window on the DNCS.
- 2 Type the following command and press **Enter**.

```
metastat
```

Result: Lengthy output, similar to the following, appears. As you examine the output, notice devices that have a **State** of **Needs maintenance**.

```
d510: Mirror
  Submirror 0: d410
    State: Okay
  Submirror 1: d710
    State: Needs maintenance
  Pass: 1
  Read option: roundrobin (default)
  Write option: parallel (default)
  Size: 50340672 blocks (24 GB)
d410: Submirror of d510
  State: Okay
  Size: 50340672 blocks (24 GB)
  Stripe 0:
    Device      Start Block  Dbase      State Reloc Hot Spare
    clt0d0s5      0          No         Okay   Yes
d710: Submirror of d510
  State: Needs maintenance
  Invoke: metareplace d510 clt4d0s5 <new device>
  Size: 50340672 blocks (24 GB)
  Stripe 0:
    Device      Start Block  Dbase      State Reloc Hot Spare
    clt4d0s5      0          No         Maintenance Yes
d507: Mirror
  Submirror 0: d407
    State: Okay
  Submirror 1: d707
    State: Needs maintenance
  Pass: 1
  Read option: roundrobin (default)
  Write option: parallel (default)
  Size: 16780224 blocks (8.0 GB)
d407: Submirror of d507
  State: Okay
```

Chapter 2 Replace a Disk Drive on a Sun Fire V445 Server

Size: 16780224 blocks (8.0 GB)

Stripe 0:

Device	Start	Block	Dbase	State	Reloc	Hot	Spare
c1t0d0s7		0	No	Okay	Yes		

```

d707: Submirror of d507
  State: Needs maintenance
  Invoke: metareplace d507 clt4d0s7 <new device>
  Size: 16780224 blocks (8.0 GB)
  Stripe 0:
    Device      Start Block  Dbase      State Reloc Hot Spare
    clt4d0s7          0      No      Maintenance  Yes
d503: Mirror
  Submirror 0: d403
    State: Okay
  Submirror 1: d703
    State: Needs maintenance
  Pass: 1
  Read option: roundrobin (default)
  Write option: parallel (default)
  Size: 16780224 blocks (8.0 GB)
d403: Submirror of d503
  State: Okay
  Size: 16780224 blocks (8.0 GB)
  Stripe 0:
    Device      Start Block  Dbase      State Reloc Hot Spare
    clt0d0s3          0      No      Okay  Yes
d703: Submirror of d503
  State: Needs maintenance
  Invoke: metareplace d503 clt4d0s3 <new device>
  Size: 16780224 blocks (8.0 GB)
  Stripe 0:
    Device      Start Block  Dbase      State Reloc Hot Spare
    clt4d0s3          0      No      Maintenance  Yes
d501: Mirror
  Submirror 0: d401
    State: Okay
  Submirror 1: d701
    State: Needs maintenance
  Pass: 1
  Read option: roundrobin (default)
  Write option: parallel (default)
  Size: 16780224 blocks (8.0 GB)
d401: Submirror of d501
  State: Okay
  Size: 16780224 blocks (8.0 GB)
  Stripe 0:
    Device      Start Block  Dbase      State Reloc Hot Spare
    clt0d0s1          0      No      Okay  Yes

```

Chapter 2 Replace a Disk Drive on a Sun Fire V445 Server

```
d701: Submirror of d501
State: Needs maintenance
Invoke: metareplace d501 c1t4d0s1 <new device>
Size: 16780224 blocks (8.0 GB)
Stripe 0:
  Device      Start Block  Dbase      State Reloc Hot Spare
  c1t4d0s1    0           No         Maintenance Yes
```

```

d500: Mirror
  Submirror 0: d400
    State: Okay
  Submirror 1: d700
    State: Needs maintenance
  Pass: 1
  Read option: roundrobin (default)
  Write option: parallel (default)
  Size: 16780224 blocks (8.0 GB)
d400: Submirror of d500
  State: Okay
  Size: 16780224 blocks (8.0 GB)
  Stripe 0:
    Device      Start Block  Dbase      State Reloc Hot Spare
    clt0d0s0      0          No         Okay   Yes
d700: Submirror of d500
  State: Needs maintenance
  Invoke: metareplace d500 clt4d0s0 <new device>
  Size: 16780224 blocks (8.0 GB)
  Stripe 0:
    Device      Start Block  Dbase      State Reloc Hot Spare
    clt4d0s0      0          No         Maintenance Yes
d367: Soft Partition
  Device: d520
  State: Okay
  Size: 4194304 blocks (2.0 GB)
    Extent      Start Block      Block count
    0            71313920         4194304
d520: Mirror
  Submirror 0: d420
    State: Okay
  Submirror 1: d720
    State: Okay
  Pass: 1
  Read option: roundrobin (default)
  Write option: parallel (default)
  Size: 402603264 blocks (191 GB)
d420: Submirror of d520
  State: Okay
  Size: 402603264 blocks (191 GB)
  Stripe 0: (interlace: 32 blocks)
    Device      Start Block  Dbase      State Reloc Hot Spare
    clt1d0s0      0          No         Okay   Yes
    clt2d0s0     10176       No         Okay   Yes

```

Chapter 2 Replace a Disk Drive on a Sun Fire V445 Server

c1t3d0s0 10176 No Okay Yes

Verify Disk Drive Status

```

d720: Submirror of d520
  State: Okay
  Size: 402603264 blocks (191 GB)
  Stripe 0: (interlace: 32 blocks)
    Device      Start Block  Dbase      State Reloc Hot Spare
    clt5d0s0      0           No         Okay  Yes
    clt6d0s0    10176       No         Okay  Yes
    clt7d0s0    10176       No         Okay  Yes

d366: Soft Partition
  Device: d520
  State: Okay
  Size: 4194304 blocks (2.0 GB)
    Extent      Start Block      Block count
    0           67119584         4194304

d365: Soft Partition
  Device: d520
  State: Okay
  Size: 4194304 blocks (2.0 GB)
    Extent      Start Block      Block count
    0           62925248         4194304

d364: Soft Partition
  Device: d520
  State: Okay
  Size: 4194304 blocks (2.0 GB)
    Extent      Start Block      Block count
    0           58730912         4194304

d363: Soft Partition
  Device: d520
  State: Okay
  Size: 4194304 blocks (2.0 GB)
    Extent      Start Block      Block count
    0           54536576         4194304

d362: Soft Partition
  Device: d520
  State: Okay
  Size: 4194304 blocks (2.0 GB)
    Extent      Start Block      Block count
    0           50342240         4194304

d361: Soft Partition
  Device: d520
  State: Okay
  Size: 4194304 blocks (2.0 GB)
    Extent      Start Block      Block count
    0           46147904         4194304
  
```

Chapter 2 Replace a Disk Drive on a Sun Fire V445 Server

d360: Soft Partition

Device: d520

State: Okay

Size: 4194304 blocks (2.0 GB)

Extent	Start Block	Block count
0	41953568	4194304

Verify Disk Drive Status

```

d359: Soft Partition
  Device: d520
  State: Okay
  Size: 4194304 blocks (2.0 GB)
      Extent          Start Block      Block count
      0                37759232         4194304
d358: Soft Partition
  Device: d520
  State: Okay
  Size: 4194304 blocks (2.0 GB)
      Extent          Start Block      Block count
      0                33564896         4194304
d357: Soft Partition
  Device: d520
  State: Okay
  Size: 4194304 blocks (2.0 GB)
      Extent          Start Block      Block count
      0                29370560         4194304
d356: Soft Partition
  Device: d520
  State: Okay
  Size: 4194304 blocks (2.0 GB)
      Extent          Start Block      Block count
      0                25176224         4194304
d355: Soft Partition
  Device: d520
  State: Okay
  Size: 4194304 blocks (2.0 GB)
      Extent          Start Block      Block count
      0                20981888         4194304
d354: Soft Partition
  Device: d520
  State: Okay
  Size: 4194304 blocks (2.0 GB)
      Extent          Start Block      Block count
      0                16787552         4194304
d353: Soft Partition
  Device: d520
  State: Okay
  Size: 4194304 blocks (2.0 GB)
      Extent          Start Block      Block count
      0                12593216         4194304
d352: Soft Partition
  Device: d520

```

Chapter 2 Replace a Disk Drive on a Sun Fire V445 Server

State: Okay

Size: 4194304 blocks (2.0 GB)

Extent	Start Block	Block count
0	8398880	4194304

```

d351: Soft Partition
  Device: d520
  State: Okay
  Size: 4194304 blocks (2.0 GB)
      Extent          Start Block      Block count
      0                4204544         4194304
d350: Soft Partition
  Device: d520
  State: Okay
  Size: 4194304 blocks (2.0 GB)
      Extent          Start Block      Block count
      0                10208           4194304
Device Relocation Information:
Device  Reloc  Device ID
c1t5d0  Yes
id1,sd@SSEAGATE_ST973401LSUN72G_3510Y65J_____3LB0Y65J
c1t6d0  Yes    id1,sd@n500000e01822f150
c1t7d0  Yes    id1,sd@n5000c50006f495d3
c1t4d0  Yes    id1,sd@n5000c50006f49527
c1t1d0  Yes    id1,sd@n500000e012994210
c1t2d0  Yes
id1,sd@SSEAGATE_ST973401LSUN72G_3510Y69E_____3LB0Y69E
c1t3d0  Yes    id1,sd@n500000e017d4f470
c1t0d0  Yes    id1,sd@n500000e0129cf6e0
$

```

Note: Notice that device d510 (and others) has a submirror that is in a **Needs maintenance** state. For device d510, the c1t4d0 component has failed and needs to be replaced. Each device with a component in a **Needs maintenance** state needs to have the underlying target drive replaced. The metadevices for each mirror in a **Needs maintenance** state will also need to be resynchronized.

Important: For the rest of this chapter, we will focus on the *c1t4d0* defective component and will use that component in the examples that are developed in this chapter.

Prepare to Replace the Hard Drive

Note: All the commands in this section should be executed as **root** user.

Identify the Device That Needs to be Replaced

- 1 If necessary, open an xterm window on the DNCS.
- 2 Type the following command and press **Enter**.

```
metastat | grep Maintenance
```

Result: Using the example established in *Verify Disk Drive Status* (on page 51), output should look similar to the following:

```
c1t4d0s5          0      No      Maintenance  Yes
c1t4d0s7          0      No      Maintenance  Yes
c1t4d0s3          0      No      Maintenance  Yes
c1t4d0s1          0      No      Maintenance  Yes
c1t4d0s0          0      No      Maintenance  Yes
```

Observations: Though multiple component devices have failed, they are all on the same physical drive. The defective drive is **c1t4**.

Unconfigure the Defective Drive

- 1 From an xterm window on the DNCS, type the following command and press **Enter**.

```
metadb
```

Result: Output should look similar to the following:

```
  flags          first blk      block count
a m p lu0       16             8192        /dev/dsk/c1t0d0s4
a   p lu0       16             8192        /dev/dsk/c1t1d0s4
a   p lu0       16             8192        /dev/dsk/c1t2d0s4
a   p lu0       16             8192        /dev/dsk/c1t3d0s4
W   p l         16             8192        /dev/dsk/c1t4d0s4
a   p lu0       16             8192        /dev/dsk/c1t5d0s4
a   p lu0       16             8192        /dev/dsk/c1t6d0s4
a   p lu0       16             8192        /dev/dsk/c1t7d0s4
```

Notes:

- Any metadb that has a flag consisting of a capital letter (like "W" in this example) needs to be removed and then re-added when the condition has been corrected.

- Should there be multiple metadbs with flags consisting of a capital letter, never remove more than *one more than half* of the metadbs at any one time. Replace a single drive at a time and resynchronize the mirror to which it belongs before removing *more than one more than half* of the metadbs.
- 2 Type the following command and press **Enter** to remove the metadb from the defective drive (c1t4d0s4, in the example we have developed).

```
metadb -d /dev/dsk/c1t4d0s4
```

- 3 Type the following command and press **Enter** to unconfigure the drive, which prepares it for physical removal.

```
cfgadm -a c1
```

Result: Output should look similar to the following example:

Ap_Id	Type	Receptacle	Occupant	Condition
c1	scsi-bus	connected	configured	unknown
c1::dsk/c1t0d0	disk	connected	configured	unknown
c1::dsk/c1t1d0	disk	connected	configured	unknown
c1::dsk/c1t2d0	disk	connected	configured	unknown
c1::dsk/c1t3d0	disk	connected	configured	unknown
c1::dsk/c1t4d0	disk	connected	configured	unknown
c1::dsk/c1t5d0	disk	connected	configured	unknown
c1::dsk/c1t6d0	disk	connected	configured	unknown
c1::dsk/c1t7d0	disk	connected	configured	unknown

Note: This output shows that all drives are currently connected, even though we know that our target drive is defective.

- 4 Type the following command and press **Enter** to unconfigure the defective drive.
- ```
cfgadm -f -c unconfigure c1::dsk/c1t4d0
```
- 5 Type the following command and press **Enter** to confirm the success of the unconfigure command.

```
cfgadm -a c1
```

**Result:** Output should look similar to the following example:

| Ap_Id          | Type     | Receptacle | Occupant     | Condition |
|----------------|----------|------------|--------------|-----------|
| c1             | scsi-bus | connected  | configured   | unknown   |
| c1::dsk/c1t0d0 | disk     | connected  | configured   | unknown   |
| c1::dsk/c1t1d0 | disk     | connected  | configured   | unknown   |
| c1::dsk/c1t2d0 | disk     | connected  | configured   | unknown   |
| c1::dsk/c1t3d0 | disk     | connected  | configured   | unknown   |
| c1::dsk/c1t4d0 | disk     | connected  | unconfigured | unknown   |
| c1::dsk/c1t5d0 | disk     | connected  | configured   | unknown   |
| c1::dsk/c1t6d0 | disk     | connected  | configured   | unknown   |
| c1::dsk/c1t7d0 | disk     | connected  | configured   | unknown   |

**Note:** Notice that device c1t4d0 is now unconfigured.






## Avoid Electrostatic Discharge

**CAUTION:**

The electronic components of printed circuit boards and hard drives are sensitive to static electricity. Do not touch the components or any metal parts without taking proper antistatic precautions.


- 1 Leave the AC power cord connected to the server.  
**Note:** AC power cords provide a discharge path for static electricity. They should remain connected UNLESS you are performing one of the following procedures:
  - Removing and installing the motherboard
  - Removing and installing the power distribution board (PDB)
  - Removing and installing the Advanced Lights Out Manager (ALOM) card
  - Removing and installing the front input/output (I/O) card
  - Removing and installing the front indicator board
  - Removing and installing the hard drive backplane
- 2 Place the new hard drive on an antistatic mat.
- 3 Attach the appropriate end of an antistatic wrist strap to the sheet metal on the system chassis and attach the other end to your wrist.



## Remove the Hard Drive


You are now ready to physically remove the hard drive. If the drive has not failed to the point where there are no LED indicators on the front, then the failed drive will have a blue indicator light illuminated, indicating that this is the drive that we unconfigured. If no LED is lit, then remove the device according to the chart from *Hard Disk Layout* (on page 50). Reference the following illustration for an example of drive indicator lights.

Each hard drive has a series of status indicators.



- 1 Identify the hard drive to be removed and note the bay in which it is installed.
- 2 Press the button that releases the hard drive handle.

- 3 Hold the hard drive by the handle and slide it out of the hard drive bay.




**Note:** Note these important points about the removal of the hard drive:

- When you reinstall the hard drive (or a replacement drive for the hard drive you just removed), be certain to install it into the same bay from which the hard drive was just removed.
- Should you intend to operate the server without replacing the hard drive, be certain to install a hard drive blank so that the cooling system of the server can function as intended.

## Install the Hard Drive

- 1 Press the button to release the handle on the new hard drive.
- 2 Orient that hard drive so that the latch of the hard drive is on the bottom.
- 3 Insert the hard drive into the bay until it makes light contact with the backplane.



**Important:** If you are replacing the hard drive that you removed in *Remove the Hard Drive* (on page 66), be certain that you are installing the new hard drive into the bay from which the other hard drive was just removed.

- 4 Firmly press the center of the hard drive handle towards the hard drive until the latch closes. This secures the hard drive in place.

## Reconnect the New Hard Drive

- 1 As root user in an xterm window, type the following command and press **Enter**.

```
cfgadm -f -c configure c1::dsk/c1t4d0
```

- 2 To confirm that the new drive is connected, type the following command and press **Enter**.

```
cfgadm -a c1
```

**Result:** Output should look similar to the following example:

| Ap_Id          | Type     | Receptacle | Occupant   | Condition |
|----------------|----------|------------|------------|-----------|
| c1             | scsi-bus | connected  | configured | unknown   |
| c1::dsk/c1t0d0 | disk     | connected  | configured | unknown   |
| c1::dsk/c1t1d0 | disk     | connected  | configured | unknown   |
| c1::dsk/c1t2d0 | disk     | connected  | configured | unknown   |
| c1::dsk/c1t3d0 | disk     | connected  | configured | unknown   |
| c1::dsk/c1t4d0 | disk     | connected  | configured | unknown   |
| c1::dsk/c1t5d0 | disk     | connected  | configured | unknown   |
| c1::dsk/c1t6d0 | disk     | connected  | configured | unknown   |
| c1::dsk/c1t7d0 | disk     | connected  | configured | unknown   |

- 3 Type the following command and press **Enter**.

```
format
```

**Result:** Output should indicate a configuration similar to the configuration expressed in *Hard Disk Layout* (on page 50).

- 4 Did the system list all the disks?
  - If **yes**, the system recognizes the newly installed hard disk; go to step 5.
  - If **no**, contact Cisco Services for assistance.

- 5 Follow these instructions to determine which volume table of contents (VTOC) to apply to the new hard drive.

**Note:** A VTOC must be obtained by another drive in the enclosure that has a similar role as the hard drive you just replaced.

- a Refer to *Hard Disk Layout* (on page 50) and locate the line with the bad device.

**Note:** In the example we have been developing, the bad device is *c1t4* in slot 4.

- b Note the **Disk Role** of the bad device.

**Note:** The **Disk Role** of the *c1t4* device is *OS*.

- c Locate another device that is not currently in a bad state and that has the same **Disk Role** as the *c1t4* bad device.

**Note:** Device *c1t0* in slot 0 meets this requirement.

- d Type the following command and press **Enter**.

```
prtvtoc /dev/rdisk/[device with same role] > /tmp/[device
with same role].vtoc
```

**Example:** `prtvtoc /dev/rdisk/c1t0d0 > /tmp/c1t0d0.vtoc`

**Result:** The system creates a file with the partition layout needed by the new disk.

- 6 Type the following command and press **Enter** to apply the vtoc to the new drive.

**Note:** In the example we have developed, the device with the same **Disk Role** of the defective *c1t4* device is *c1t0*.

```
fmthard -s /tmp/[device with same role].vtoc
/dev/rdisk/[device name of replaced drive]
```

**Example:** `fmthard -s /tmp/c1t0d0.vtoc /dev/rdisk/c1t4d0`

**Result:** The new drive now has the same partitioning as the drive with which it is most similar in the disk mirror.

- 7 Was the drive that you removed a boot drive (slots 0 or 4)?

- If **yes**, you need to add boot blocks. This procedure is similar to *Install a Boot Block on a Mirrored Disk Drive (6-Disk Setup)* (on page 21).

**Note:** Then, continue with step 8.

- If **no**, go to step 8.

- 8 Type the following command and press **Enter** to restore the metadb.

```
metadb -a c1t4d0s4
```

- 9 Type the following command and press **Enter** to show each mirror and each subcomponent that needs to be resynchronized.

```
metastat | egrep "Mirror|Maintenance"
```

**Result:** Output should be similar to the following example:

```
d510: Mirror
 c1t4d0s5 0 No Maintenance Yes
d507: Mirror
 c1t4d0s7 0 No Maintenance Yes
d503: Mirror
 c1t4d0s3 0 No Maintenance Yes
d501: Mirror
 c1t4d0s1 0 No Maintenance Yes
d500: Mirror
 c1t4d0s0 0 No Maintenance Yes
d520: Mirror
```

**Note:** The line above each indented line represents the mirror to which the component should be reattached.

- 10 To reattach the mirrors, type the following commands and press **Enter**.

```
metareplace -e d510 c1t4d0s5
metareplace -e d507 c1t4d0s7
metareplace -e d503 c1t4d0s3
metareplace -e d501 c1t4d0s1
metareplace -e d500 c1t4d0s0
```

- 11 Type the following command and press **Enter** to verify that there are no additional disk failures.

```
syncwait.ksh
```

**Result:** Output should be similar to the following example:

```
Checking Metadevice status... Please wait.
Disks Syncing... Please wait.
d510: Mirror
 Submirror 0: d410
 Submirror 1: d710
 State: Resyncing
 Resync in progress: 5 % done
d507: Mirror
 Submirror 0: d407
 Submirror 1: d707
```

```
State: Resyncing
Resync in progress: 4 % done
```



```
d503: Mirror
 Submirror 0: d403
 Submirror 1: d703
 State: Resyncing
 Resync in progress: 12 % done
d501: Mirror
 Submirror 0: d401
 Submirror 1: d701
 State: Resyncing
 Resync in progress: 6 % done
d500: Mirror
 Submirror 0: d400
 Submirror 1: d700
 State: Resyncing
 Resync in progress: 7 % done
d520: Mirror
 Submirror 0: d420
 Submirror 1: d720
```

# 3

# Customer Information

## If You Have Questions

If you have technical questions, call Cisco Services for assistance. Follow the menu options to speak with a service engineer.

Access your company's extranet site to view or order additional technical publications. For accessing instructions, contact the representative who handles your account. Check your extranet site often as the information is updated frequently





Cisco Systems, Inc.  
5030 Sugarloaf Parkway, Box 465447  
Lawrenceville, GA 30042

678 277-1120  
800 722-2009  
[www.cisco.com](http://www.cisco.com)

This document includes various trademarks of Cisco Systems, Inc. Please see the Notices section of this document for a list of the Cisco Systems, Inc. trademarks used in this document.

Product and service availability are subject to change without notice.

© 2007, 2010, 2012 Cisco and/or its affiliates. All rights reserved.

March 2012 Printed in USA

Part Number 4005697 Rev C