

Lucent Technologies
Bell Labs Innovations



***CentreVu*[®] Call Management System**

Release 3 Version 8

Disk-Mirrored Systems

585-210-940
Comcode 108502352
Issue 1
December, 1999

Copyright © 1999 Lucent Technologies
All Rights Reserved
Printed in U.S.A.

Notice

Every effort was made to ensure that the information in this book was complete and accurate at the time of printing. However, information is subject to change.

Your Responsibility for Your System's Security

Toll fraud is the unauthorized use of your telecommunications system by an unauthorized party, for example, persons other than your company's employees, agents, subcontractors, or persons working on your company's behalf. Note that there may be a risk of toll fraud associated with your telecommunications system and, if toll fraud occurs, it can result in substantial additional charges for your telecommunications services.

You and your system manager are responsible for the security of your system, such as programming and configuring your equipment to prevent unauthorized use. The system manager is also responsible for reading all installation, instruction, and system administration documents provided with this product in order to fully understand the features that can introduce risk of toll fraud and the steps that can be taken to reduce that risk. Lucent Technologies does not warrant that this product is immune from or will prevent unauthorized use of common-carrier telecommunication services or facilities accessed through or connected to it. Lucent Technologies will not be responsible for any charges that result from such unauthorized use.

Lucent Technologies Fraud Intervention

If you *suspect that you are being victimized* by toll fraud and you need technical support or assistance, call Technical Service Center Toll Fraud Intervention Hotline at 1-800-643-2353.

Trademarks

CentreVu is a registered trademark of Lucent Technologies.
AnswerBook, *Enterprise*, *Solaris*, *SunLink*, and *Solstice DiskSuite* are trademarks of Sun Microsystems, Inc.
Sun and *Sun Microsystems* are registered trademarks of Sun Microsystems, Inc.
INFORMIX is a registered trademark of Informix Software, Inc.

Ordering Information

Call: Lucent Technologies Publications Center
Voice: 1-800-457-1235
International Voice: 317-322-6416
Fax: 1-800-457-1764
International Fax: 317-322-6699
Write: Lucent Technologies BCS Publications Center
2855 N. Franklin Road
Indianapolis, IN 46219
Order: Document No. 585-210-940
Comcode 108502352
Issue 1, December 1999

You can be placed on a Standing Order list for this and other documents you may need. Standing Order will enable you to automatically receive updated versions of individual documents or document sets, billed to account information that you provide. For more information on Standing Orders, or to be put on a list to receive future issues of this document, contact the Lucent Technologies Publications Center.

Disclaimer

Intellectual property related to this product (including trademarks) and registered to AT&T Corporation has been transferred to Lucent Technologies Incorporated.

Any references within this text to American Telephone and Telegraph Corporation or AT&T should be interpreted as references to Lucent Technologies Incorporated. The exception is cross references to books published prior to December 31, 1996, which retain their original AT&T titles.

Heritage Statement

Lucent Technologies—formed as a result of AT&T's planned restructuring—designs, builds, and delivers a wide range of public and private networks, communication systems and software, consumer and business telephone systems, and microelectronics components. The world-renowned Bell Laboratories is the research and development arm for the company.

Comments

To comment on this document, return the comment card at the front of the document.

Acknowledgment

This document was developed by the Lucent Technologies Product Publications group of the Global Learning Solutions organization.

Table of Contents

Chapter 1: Introduction to Disk Mirroring

Disk mirroring defined	1
Metadevices.	1
Concatenated metadevices	2
State databases	3
Mirrors.	3
How CMS implements disk mirroring	5
Mirror layouts	7
Mirror layout on a Sun Enterprise 3000 computer.	7
Mirror layout on a Sun Enterprise 3500 computer.	8
Required hardware	8
Required software	9
Additional references	9

Chapter 2: Installing Mirrored Systems

How to upgrade an Enterprise platform to disk mirroring.	11
Disk drive slots	11
Increasing system capacity	12
Step 1: Identify mirror devices.	12
Step 2: Install the mirror disks on an Enterprise 3500 system	14
Step 3: Install the mirror disks on an Enterprise 3000 system	16
Step 4: Partition the mirror disks	16
Step 5: Create the state database replicas	18
Step 6: Create the md.tab table entries	18
Step 7: Set up the mirrors for the root file system	20
Step 8: Set up the mirrors for the swap space	22
Step 9: Set up the /cms mirror.	24
Step 10: Verify disk space	25
Step 11: Create an alternate boot device	26
Factory installation procedures for disk mirroring	28
Step 1: Install the hardware	28
Step 2: Install the operating system.	28
Step 3: Install the application software	29
Step 4: Check the disk.type file	30
Step 5: Set up the disk subsystem for Solstice DiskSuite.	30
Step 6: Run the setup scripts	32
Step 7: Verify disk space.	33
Step 8: Install the CMS software.	35
Step 9: Create an alternate boot device.	35

Chapter 3: Maintaining Mirrored Systems

Using the alternate boot device	39
Maintaining the chkDisks crontab	40
Activating chkDisks	40
Verifying chkDisks.	40
Changing the scheduled run time	40
Canceling chkDisks	40
Doing a CMSADM backup on a mirrored system	41
Step 1: Print the vfstab file	41
Step 2: Run the backup	41
Step 3: Label and store the tapes	42
Restoring from a CMSADM Backup	43
Repairing or rebuilding the /cms file system	45
Repairing /cms.	45
Rebuilding /cms	46
Restoring a mirrored system from a CMSADM backup.	49
Step 1: Power off the system	49
Step 2: Install the new drives	49
Step 3: Boot the system from the compact disk	50
Step 4: Identify the system	51
Step 5: Set the Name Service options	52
Step 6: Set the date and time	53
Step 7: Select the operating system files to be installed	54
Step 8: Select the hard disks	55
Step 9: Set up the disk partitions.	55
Step 10: Install the operating system	57
Step 11: Complete the OS installation	58
Step 12: Alter the tape configuration file (Enterprise 3500 only).	58
Step 13: Restore the backup	59
Step 14: Reestablish the disk mirrors	59
Restoring specific files from a CMSADM backup	62
Adding new disks to a standard mirrored system	63
Adding new disks to a nonstandard mirrored system	66
Replacing a faulty disk	69
Step 1: Identify the faulty disk	69
Step 2: Identify the submirrors	70
Step 3: Detach the submirrors	71
Step 4: Remove state database replicas (boot disks only)	72
Step 5: Replace the faulty disk drive.	72
Step 6: Recreate the state database replicas	73
Step 7: Reattach the submirrors	74
Step 8: Reboot the system (boot disks only)	74
State Database Replicas	75
Setting up replicas.	75
Replacing replicas.	75

Partitioning disks with the format command	76
How to tell whether your mirrored system is standard	78
Examples of valid system files	79
Valid vfstab files	79
Valid md.tab files	80
 Chapter 4: Troubleshooting	
Solstice DiskSuite problems	81
The root filesystem runs out of space during a system install.	81
Excessively long resync	81
The system fails to recognize all disk drives	82
Fixing “needs maintenance” messages	84
Boot problems	85
Trying to boot from the wrong device	85
Primary boot device is not available	85
Excessive reboot time.	86
Log files	87
Error messages	88
Index	97

Introduction to Disk Mirroring

This chapter introduces you to “disk mirroring,” an optional feature of CentreVu® Call Management System (CMS) that provides you with a completely redundant set of data, helping to ensure data security. To use disk mirroring, you must have a Sun Enterprise* 3000 or Sun Enterprise 3500 platform running CMS r3v8.*

Disk mirroring defined

“Mirrors” are a feature of the *Solstice DiskSuite** software package. They allow you to build a hard disk system containing two—or even three—complete sets of data. Having such data redundancy greatly reduces the risk of data loss should a hard disk drive fail or your system crash. While mirrors greatly reduce the risk of losing data, however, they are not meant to be a substitute for regular backups. **Mirrored systems must be backed up just as often as unmirrored systems.**

Metadevices

The *Solstice DiskSuite* software package allows multiple disk partitions to be logically combined to create a single large partition. Using the *Solstice DiskSuite* package allows CMS databases to span multiple disks, and so grow quite large.


Solstice DiskSuite uses virtual disks to manage physical disks and their associated data. In *Solstice DiskSuite*, a virtual disk is called a **metadevice**. To a software application, a metadevice is identical to a physical disk drive. *Solstice DiskSuite* handles all I/O requests directed at a metadevice, converting them into I/O requests for the underlying disks.

Solstice DiskSuite metadevices are built from **slices** (disk partitions). A system controlled by *Solstice DiskSuite* may contain any number of metadevices, each of which may comprise any combination of disk partitions.

Once a metadevice has been set up, the underlying disk partitions can be accessed only through the metadevice.

*Sun, Enterprise, and *Solstice DiskSuite* are trademarks or registered trademarks of Sun Microsystems, Inc.


The illustration below shows a simple, two-slice metadvice.



Concatenated metadevices

A metadvice can be configured as any one of three basic types: *striped*, *concatenated*, or *concatenated stripes*. The type used by CMS systems is the concatenated metadvice. In a concatenated metadvice, data blocks, or *chunks*, are written sequentially across the slices, beginning with the first disk.

Consider, for example, a concatenated metadvice with three slices (see the illustration below.) In that scenario, disk A can be envisioned as containing logical chunks 1 through 4, disk B as containing logical chunks 5 through 8, and disk C as containing chunks 9 through 12.



The total capacity of this concatenated metadvice is the combined capacities of the three drives. If each drive is 4 gigabytes, for example, the metadvice has an overall capacity of 12 gigabytes.

State databases

The *Solstice DiskSuite* software tracks which disk partitions belong to which metadevices with a **state database**. A state database stores information on disk about the state of your *Solstice DiskSuite* configuration.


The state database consists of multiple copies of the basic database. The copies, referred to as **state database replicas**, ensure that the data in the database is always valid. Having multiple copies protects against data loss from single points-of-failure. The state database tracks the location and status of all state database replicas. *Solstice DiskSuite* cannot operate until you have created the state database and its replicas: the software must have an operating state database.

Mirrors

A **mirror** is a metadvice that can copy data from one metadvice to another. The metadevices containing the data are called **submirrors**. The process of copying the data between submirrors is called **mirroring**.


Mirroring provides redundant copies of data. To a software application, a mirror looks just like a physical disk. The mirror accepts I/O requests and converts them into I/O requests for the submirrors. The submirrors in turn—being metadevices themselves—convert I/O requests from the mirror into I/O requests for the underlying physical disks.

A mirror can be **two-way** or **three-way**. A two-way mirror is a mirror with two submirrors; a three-way mirror has three submirrors. CMS supports only two-way mirrors. The illustration below shows a typical two-way mirror, d21, consisting of the submirrors d19 and d20.




How CMS implements disk mirroring


A nonmirrored, CMS system with *Solstice DiskSuite* uses it to create a single metadvice, named d19, containing all the disk partitions used to store CMS data. For example:




To implement disk mirroring, a metadvice d20 is created as a duplicate of d19, and the two metadevices are configured as submirrors of d21, as shown in the following example:



In a CMS computer system with */cms* mirrored, *root* is also mirrored. The root mirror is d13; its submirrors are d11 and d12. For example:



If your computer system is running CMS r3v8, your swap partition is also mirrored. The swap mirror is d17; its submirrors are d15 and d16. For example:



Mirror layouts

Mirror layouts on a factory-installed mirrored system differ depending upon the platform. When a system is upgraded to mirroring in the field, mirror layouts are essentially unpredictable, and depend on how many disk drives are already installed, and what their target numbers are.

Mirror layout on a *Sun Enterprise 3000* computer

On a *Sun Enterprise 3000* platform with factory-installed mirroring, the mirror layout is the simple odd/even arrangement shown in the table below. A system upgraded to mirroring in the field probably has a different mirror layout.

	mirror d13 (/ (root) filesystem)		mirror d21 (/cms filesystem)		mirror d17 (swap partition) (CMS r3v8)	
No. Disks	submirror d11	submirror d12	submirror d19	submirror d20	submirror d15	submirror d16
2	c0t0d0s0	c0t1d0s0	c0t0d0s3	c0t1d0s3	c0t0d0s4	c0t1d0s4
4	c0t0d0s0	c0t1d0s0	c0t0d0s3 c0t2d0s1	c0t1d0s3 c0t3d0s1	c0t0d0s4	c0t1d0s4
6	c0t0d0s0	c0t1d0s0	c0t0d0s3 c0t2d0s1 c0t10d0s1	c0t1d0s3 c0t3d0s1 c0t11d0s1	c0t0d0s4	c0t1d0s4
8	c0t0d0s0	c0t1d0s0	c0t0d0s3 c0t2d0s1 c0t10d0s1 c0t12d0s1	c0t1d0s3 c0t3d0s1 c0t11d0s1 c0t13d0s1	c0t0d0s4	c0t1d0s4
10	c0t0d0s0	c0t1d0s0	c0t0d0s3 c0t2d0s1 c0t10d0s1 c0t12d0s1 c0t14d0s1	c0t1d0s3 c0t3d0s1 c0t11d0s1 c0t13d0s1 c0t15d0s1	c0t0d0s4	c0t1d0s4

Mirror layout on a Sun Enterprise 3500 computer

On a Sun Enterprise 3500 platform, the mirror layout must always be based upon the bays in which the disks are installed. Submirrors 11, 15, and 19 are the disks in the lower bay (targets 0 through 3 on controller 0); submirrors 12,16, and 20 are the disks in the upper bay (targets 4 through 7 on controller 1). The table below shows the mirror layout.

	mirror d13 / (root) filesystem		mirror d21 /cms filesystem		mirror d17 swap partition (CMS r3v8)	
No. Disks	submirror d11	submirror d12	submirror d19	submirror d20	submirror d15	submirror d16
2	c0t0d0s0	c1t4d0s0	c0t0d0s3	c1t4d0s3	c0t0d0s4	c1t4d0s4
4	c0t0d0s0	c1t4d0s0	c0t0d0s3 c0t1d0s1	c1t4d0s3 c1t5d0s1	c0t0d0s4	c1t4d0s4
6	c0t0d0s0	c1t4d0s0	c0t0d0s3 c0t1d0s1 c0t2d0s1	c1t4d0s3 c1t5d0s1 c1t6d0s1	c0t0d0s4	c1t4d0s4
8	c0t0d0s0	c1t4d0s0	c0t0d0s3 c0t1d0s1 c0t2d0s1 c0t3d0s1	c1t4d0s3 c1t5d0s1 c1t6d0s1 c1t7d0s1	c0t0d0s4	c1t4d0s4

- Required hardware
- In order for your CMS system to be mirrored, it must have the following hardware installed in addition to the hardware already installed:
- For a Sun Enterprise 3500 system, two GigaByte Interface Converter (GBIC) modules. One will be installed into the UA slot on the FC-AL Interface board, and the other will be installed into GBIC Port 1 on the Sbus I/O board (see the illustration on page 15).
 - For a Sun Enterprise 3500 system, a fiber cable to connect the UA port GBIC to GBIC Port 1 on the I/O board.
 - Twice the number of disk drives needed for an unmirrored system. All the disks must be the same size.

Required software

In order for your CMS system to be mirrored, it must be running the following software:

CMS R3V8
<i>Solaris 7</i> dated 3/99
<i>Solstice DiskSuite 4.2</i>
any CMS r3v8 load

For a list of other software required for your system, see the Lucent Technologies Software Installation manual for your version of CMS.

Additional references

Sun Microsystems^{*} documentation related to *Solstice DiskSuite* and disk maintenance and configuration includes:

Solstice DiskSuite 4.2 Administration Guide
Solstice DiskSuite Tool 4.2 User's Guide

^{*}*Sun Microsystems* is a registered trademark of Sun Microsystems, Inc.

Installing Mirrored Systems

This chapter contains procedures for upgrading existing computer systems to use disk mirroring and for installing a mirrored system under factory conditions.

How to upgrade an *Enterprise* platform to disk mirroring

This section tells how to upgrade *CentreVu*® Call Management System (CMS) on a *Sun Enterprise* 3000 or 3500 computer from a nonmirrored system to a mirrored system.

Disk drive slots

Before you begin, you must understand the drive slot arrangement in your system.

In an *Enterprise* 3000 computer, there are 10 slots, allowing up to five disks for each mirror. Each slot is labeled with a number 0 through 3 or 10 through 15; there are no slots numbered 4 through 9. All the drive slots are on controller 0.

In an *Enterprise* 3500 computer, there are eight disk drive slots, four in each of two bays. The slots in the lower bay are labeled 0 through 3 and are on controller 0; the slots in the upper bay are numbered 4 through 7 and are on controller 1. In a mirrored system, slots 0 through 3 are reserved for the original disks, and slots 4 through 7 are reserved for the mirror disks.

The slot number is the drive's target number, which becomes part of the device name. A drive in slot 1, for example, has the device name *c0t1d0*. The drive in slot 0 is always the primary boot disk.

In an *Enterprise* 3000 mirrored system, mirror disks are normally installed in target number order, beginning with the first empty slot. That makes it difficult to predict the mirror layout on an upgraded system. In an *Enterprise* 3500 mirrored system, however, mirror disks always go in the upper bay. If you are mirroring a two-disk system, for example, the two original disks would already be installed in slots 0 and 1 in the lower bay, and you would install the two new disks in slots 4 and 5 in the upper bay.

Increasing system capacity

If you intend to increase your system capacity as well as mirror your system, first install the new disks needed to increase capacity. For instructions, see the documentation that came with the new drives or the appropriate Lucent Technologies installation documents. After the system capacity has been increased, and it has been verified that the system works correctly, you may then install disk mirroring. **You must make certain the basic system works before you try to mirror the disks.**

Step 1: Identify mirror devices

1. Check the number of new disks to be installed as the mirror.

There should be as many new disks to be installed as there are old disks already on your system. That keeps the same system capacity and makes the two submirrors the same size.
2. Next you must identify the device names for your mirrored system and record them in the *Device/Metadevice Names* table on the next page.

If you have an *Enterprise* 3500 system, the device names have already been identified and recorded for you in the *Enterprise* 3500 table. Here's how you identify the device names for an *Enterprise* 3000 system:

- a. **Identify the secondary boot disk and record the device names.**

Find the first empty slot in your system. Whatever slot you find will hold the secondary boot disk. Then record the following device names for that disk:
 - Record the slice 0 device name as submirror 12.
 - Record the slice 3 device name as disk 1 of submirror 20.
 - If your system is running load r3v6aj.c or later of CMS—you can find out by executing a **pkginfo -x cms** command—record the slice 4 device as submirror 16.

Example: If the first empty drive slot in your system were slot 2, you would record the following values in the table:

Submirror d12 Device Name boot disk (2): <u>c0t2d0s0</u>	• • •	Submirror d20 Device Names CMS disk 1: <u>c0t2d0s3</u>
Submirror d16 Device (r3v6aj.c or later) swap <u>c0t2d0s4</u>		

- b. Record existing nonboot disks under submirror 19, and empty slots for new nonboot disks under submirror 20.

Example: Imagine a system that already has two disk drives in slots 0 and 1. In the previous step, you would have reserved slot 2 for the second boot disk. The remaining new disk is a nonboot disk and will be installed in slot 3. So you would record the following values in the table:

Submirror d19 Device Names		Submirror d20 Device Names	
CMS disk 1:	<u>c0t0d0s3</u>	CMS disk 1:	<u>c0t2d0s3</u>
CMS disk 2:	<u>c0t1d0s1</u>	CMS disk 2:	<u>c0t3d0s1</u>
CMS disk 3:	_____	CMS disk 3:	_____
CMS disk 4:	_____	CMS Disk 4:	_____
CMS disk 5:	_____	CMS disk 5:	_____

Device/Metadevice Names for an *Enterprise 3000* system

Submirror d11 Device Name	Submirror d12 Device Name	Submirror d19 Device Names	Submirror d20 Device Names
boot disk 1: <u>c0t0d0s0</u>	boot disk 2: _____s0	CMS disk 1: <u>c0t0d0s3</u>	CMS disk 1: _____s3
		CMS disk 2: _____s1	CMS disk 2: _____s1
Submirror d15 Device Name (r3v6aj.c and later)	Submirror d16 Device Name (r3v6aj.c and later)	CMS disk 3: _____s1	CMS disk 3: _____s1
		CMS disk 4: _____s1	CMS Disk 4: _____s1
swap, disk 1: <u>c0t0d0s4</u>	swap, disk 2: _____s4	CMS disk 5: _____s1	CMS disk 5: _____s1

Device/Metadevice Names for an *Enterprise 3500* system

Submirror d11 Device Name	Submirror d12 Device Name	Submirror d19 Device Names	Submirror d20 Device Names
boot disk 1: <u>c0t0d0s0</u>	boot disk 2: <u>c1t4d0s0</u>	CMS disk 1: <u>c0t0d0s3</u>	CMS disk 1: <u>c1t4d0s3</u>
Submirror d15 Device Name (r3v6aj.c and later)	Submirror d16 Device Name (r3v6aj.c and later)	CMS disk 2: <u>c0t1d0s1</u>	CMS disk 3: <u>c1t5d0s1</u>
		CMS disk 3: <u>c0t2d0s1</u>	CMS Disk 4: <u>c1t6d0s1</u>
swap, disk 1: <u>c0t0d0s4</u>	swap, disk 2: <u>c1t4d0s4</u>	CMS disk 4: <u>c0t3d0s1</u>	CMS disk 5: <u>c1t7d0s1</u>

Step 2: Install the mirror disks on an *Enterprise 3500* system

Do this step only if you have an *Enterprise 3500* system.

It is likely that your *Enterprise 3500* system is currently using controller 1 (c1) for the internal CD-ROM drive. This step associates the CD-ROM drive with controller 2 (c 2) so controller 1 will be free for the new disk drives. In this step you also install some additional hardware for controlling the new disks, and install the new disks themselves. You should be logged in as *root* before beginning this step.

1. Enter the following command:

```
# ls -l /dev/dsk/c1t6d0s0
```

The system should respond with a line similar to the following:

```
lrwxrwxrwx 1 root root 50 Nov 16 08:58
/dev/dsk/c1t6d0s0 -> ../../devices/sbus@3,0
/SUNW,fas@3,88000000/sd@6,0:a
```


Notice that c1t6d0s0 is pointing to sd@6,0:a. That indicates that controller 1 is being used by the CD-ROM drive. Remove that association by entering the following commands:

```
# rm /dev/dsk/c1*
# rm /dev/rdisk/c1*
# touch /reconfigure
# /usr/sbin/shutdown -y -g0 -i0
```


2. Install new hardware components as follows:

- Install a GigaByte Interface Converter (GBIC) into slot UA on the Interface board.
- Install another GBIC into port 1 on the SBus I/O board.
- Install a fiber cable between the two newly installed GBICs.

The illustration below shows where the GBICs are installed. For more detailed information, see *CentreVu CMS Sun Enterprise 3500 System (connectivity diagram)* (585-215-877), which was packed with your system.



3. Install the new disk drives in the upper drive bay. The illustration below shows where to install the new drives.



4. From the ok prompt, do a reconfigure boot by entering the following command:

```
ok boot -r
```

5. Log in as *root*.
6. Enter the following command:

```
# ls -l /dev/dsk/c2t6d0s0
```

The system should respond with a line similar to the following:

```
lrwxrwxrwx 1 root root 50 Nov 16 08:58  
/dev/dsk/c2t6d0s0 -> ../../devices/sbus@3,0  
/SUNW,fas@3,88000000/sd@6,0:a
```

Notice that c2t6d0s0 is pointing to sd@6,0:a, indicating that the CD-ROM drive is now on controller 2.

Step 3: Install the mirror disks on an *Enterprise 3000* system

Do this step only if you have an *Enterprise 3000* system.

Install the new disk drives, which will form the mirror, into the empty slots reserved in step 1. Follow the procedure in the “Add a Disk Drive” subsection in the Lucent Technologies hardware installation or hardware maintenance manual for your version of CMS. **DO NOT PARTITION THE DISKS JUST YET;** that happens next.

Step 4: Partition the mirror disks

Partition the newly installed disks using the format command. If you have an *Enterprise 3500*, make certain that the new drives are named using “c1,” as c1t4d0 or c1t5d0.

The procedure is as follows:

1. Enter: `# format`
A menu of disks displays. You are asked to select one.
2. Enter the menu number of the new disk to be partitioned.
The `format>` prompt displays.
3. Enter: `format> partition`
The `partition>` prompt displays.

4. Enter: `partition> print`

The partition table displays and the `partition>` prompt returns. Compare the display to the *Partitioning* table below to see whether you need to re-partition the disk. You probably will need to re-partition.

5. Enter the number of the partition to be reconfigured.

The system begins a series of prompts for partitioning information. Answer according to the CMS version and disk type you are partitioning (see the *Partitioning* table below).

Repeat this task for each partition on the disk.

Partitioning (Note: **Never** change the default size of the *overlap* slice!)

Disk Type	Part	Tag	Flag	CMS R3V8 4.2-GB Disk		CMS R3V8 9.1-GB Disk SCSI (E3000) and FC-AL (E 3500)	
				Size (cyl.)	Starting Cylinder	Size (cyl.)	Starting Cylinder
Boot	0	<i>boot disk 1: /</i> or root <i>boot disk 2:</i> unassigned	wm	1023	0	616	0
	1	unassigned	wm	7	1023	7	616
	2	overlap <i>or</i> backup	wm	3880	0	4924	0
	3	unassigned	wm	1879	1030	3716	623
	4	swap	wm	971	2109	585	4339
	5-7	unassigned	wm	0	0	0	0
Nonboot	0	unassigned	wm	2	0	2	0
	1	unassigned	wm	3878	2	4922	2
	2	overlap <i>or</i> backup	wm	3880	0	4924	0
	3-7	unassigned	wm	0	0	0	0

6. Enter: `partition> print`

The partition table displays and the `partition>` prompt returns. Compare the display to the *Partitioning* table to make certain you have repartitioned the disk properly. If not, repeat step 5.

7. Enter: `partition> label`

The system responds Ready to label disk, continue?

8. Enter: `y`
The `partition>` prompt returns.
9. Enter: `partition> q`
The `format>` prompt returns.
10. To select the next disk to be partitioned, enter: `format> disk` and return to item 2.

If you have finished partitioning **all** the new disks, enter: `format> q` and continue with the next step.

Step 5: Create the state database replicas

1. Alter your path by entering the following commands:

```
# export PATH=$PATH:/usr/opt/SUNWmd/sbin:/olds
```

2. Create two state database replicas on slice 1 of boot disk 2:

```
# metadb -a -c3 -f /dev/rdisk/<device>
```

(where *<device>* is the target number of boot disk 2)
3. Verify that the replicas were created by entering the following command:

```
# metadb -i
```

The system lists the state database replicas currently on the system. The following sample output shows two replicas on c0t2:

	flags	first blk	block count	
a	u	16	1034	/dev/dsk/c0t2d0s1
a	pc lu0	16	1034	/dev/dsk/c0t2d0s1

Make sure the list includes the two replicas you just created.

Step 6: Create the *md.tab* table entries

1. Make copies of essential system files by entering these commands:

```
# cp /etc/opt/SUNWmd/md.tab /etc/opt/SUNWmd/md.tab.orig
# cp /etc/vfstab /etc/vfstab.orig
# cp /etc/system /etc/system.orig
```

2. Edit */etc/opt/SUNWmd/md.tab* and insert entries for the *root* and */cms* metadevices. If your system is running CMS load *r3v6aj.c* or later, you must also create a section for the *swap* metadevice.

The root metadvice definitions go into the **#metaroot** section, which you must create for the occasion; d20 goes into the **#!/cms** section, which already exists in the file. The swap metadvice definitions go into the **#swap** section which, like **#metaroot**, must be created for the occasion.

The lines you add must be in the following general format:

```
d<x> <y> 1 <device>
```

d<x> is the metadvice name.

<y> denotes the total number of disks named in the line.

1 <device> indicates one disk with the device name **<device>**. You should have recorded the device names on [page 13](#).

The pattern **1 (space) <device>** repeats until all disks in the metadvice have been named. The d12 entry must specify one disk; the d20 entry must specify the same number of disks as the existing d19 entry. If your system has two disks in slots 0 and 1 and you are adding a two-disk mirror in slots 2 and 3, for example, you would add the *md.tab* lines shown in boldface in the following sample:

```
----- /etc/opt/SUNWmd/md.tab - EXAMPLE ONLY -----
.
.
.
#metaroot
d11 1 1 c0t0d0s0
d12 1 1 c0t2d0s0
d13 -m d11

#!/cms
d19 2 1 /dev/rdisk/c0t0d0s3 1 /dev/rdisk/c0t1d0s1
d20 2 1 /dev/rdisk/c0t2d0s3 1 /dev/rdisk/c0t3d0s1
d21 -m d19

#metaswap      (for load r3v6aj.c or later)
d15 1 1 c0t0d0s4
d16 1 1 c0t2d0s4
d17 -m d15
-----
```

3. Save the */etc/opt/SUNWmd/md.tab* file and quit the editor.

Step 7: Set up the mirrors for the *root* file system

Set up the mirrors for the *root* file system with the following steps:

1. Activate metadvice d11:

```
# metainit -f d11
d11: Concat/stripe is set up
#
```

2. Activate metadvice d13:

```
# metainit -f d13
d13: Concat/stripe is set up
#
```

3. Activate metadvice d12:

```
# metainit -f d12
d12: Concat/stripe is set up
#
```

4. Attach d12 to d13:

```
# metattach d13 d12
d13: submirror d12 is attached
#
```

5. Verify that root is set up and resynching:

```
# metastat d13
d13: Mirror
    Submirror 0: d11
        State: Okay
    Submirror 1: d12
        State: Resynching
    Pass: 1
    Read option: roundrobin (default)
    Write option: parallel (default)
    Size: 2209680 blocks
    .
    .
    .
```

Check the State: lines in the output list. They should read either Okay or Resynching. If they do, then everything is fine. If they indicate any other status, something has gone wrong with the setup: discontinue the upgrade and telephone the Lucent National Customer Care Center (1.800.242.2121) or contact your Lucent representative or distributor.

6. Set up d13 as the root mirror:

```
# metaroot d13
#
```

7. Reboot the system and log in as the root user:

```
# /usr/sbin/shutdown -i6 -g0 -y

<hostname> login: root
Password: <password>
#
```

Step 8: Set up the mirrors for the swap space

Set up the mirrors for the swap space with the following steps:

1. Activate metadvice d15:

```
# metainit -f d15
d15: Concat/stripe is set up
#
```

2. Activate metadvice d17:

```
# metainit d17
d17: Concat/stripe is set up
#
```

3. Activate metadvice d16:

```
# metainit d16
d16: Concat/stripe is set up
#
```

4. Attach d16 to d17:

```
# metattach d17 d16
d17: submirror d16 is attached
#
```

5. Verify that swap is set up and resynching:

```
# metastat d17
d17: Mirror
    Submirror 0: d15
        State: Okay
    Submirror 1: d16
        State: Resyncing
    Pass: 1
    Read option: roundrobin (default)
    Write option: parallel (default)
    Size: 2209680 blocks
    .
    .
    .
```

Check the State: lines in the output list. They should read either Okay or Resyncing. If they do, then everything is fine. If they indicate any other status, something has gone wrong with the setup: discontinue the upgrade and telephone the Lucent National Customer Care Center (1.800.242.2121) or contact your Lucent representative or distributor.

6. Edit */etc/vfstab*. The relevant portion of the file looks like this:

```
----- /etc/vfstab -----
. . .
/dev/dsk/c0t0d0s4    -    swap    -    no    -
/dev/dsk/c0t2d0s4    -    swap    -    no    -
-----
```

Do the following steps to modify the file:

- a. Copy the *c0t0d0s4* entry and insert the copy on the succeeding line.
- b. Comment out the original entry by inserting a pound sign at the beginning of the line.
- c. Alter the copy by:
 1. replacing *c0t0d0s4* with *d17*
 2. adding *md/* after the *dev/*
- d. Immediately after the line just added should be a line naming the secondary boot disk. Comment out that line by inserting a pound sign at the beginning of the line.

When you have finished *a* through *d*, the relevant lines of the file might look something like this:

```
----- /etc/vfstab -----
. . .
#/dev/dsk/c0t0d0s4    -    swap    -    no    -
/dev/md/dsk/d17       -    swap    -    no    -
#/dev/dsk/c0t2d0s4    -    swap    -    no    -
-----
```

- e. Save the file and quit the editor.

7. Set up the swap space by entering the following command:

```
# swap -a /dev/md/dsk/d17
#
```

Step 9: Set up the /cms mirror

1. Edit */etc/vfstab*. Copy the d19 entry to the succeeding line. Then comment out the original and change “d19” to “d21”. For example:

```
----- /etc/vfstab - EXAMPLE ONLY -----
# . . .
# /dev/md/dsk/d19 /dev/md/rdsk/d19 /cms ufs 2 yes -
# /dev/md/dsk/d21 /dev/md/rdsk/d21 /cms ufs 2 yes
-----
```

2. Save the */etc/vfstab* file and exit the editor.
3. Alter your path:

```
# export PATH=$PATH:/usr/opt/SUNWmd/sbin:/olds
```

4. Activate the new metadvice (d20):

```
# metainit d20
d20: Concat/stripe is set up
#
```

5. Verify that d20 is active:

```
# metastat d20
d20: Concat/Stripe
Size: 14398560 blocks
Stripe 0:
Device      Start Block  Dbase State Hot Spare
c0t1d0s3      0         No    Okay
Stripe 1:
Device      Start Block  Dbase State Hot Spare
c0t3d0s1      0         No    Okay
#
```

6. Initialize the newly created mirror (d21):

```
# metainit -f d21
d21: Concat/stripe is set up
#
```

7. Attach the new metadvice (d20) to the newly created mirror (d21):

```
# metattach d21 d20
d21: submirror d20 is attached
#
```

8. Reboot the system by entering the following command:

```
# /usr/sbin/shutdown -y -i6 -g0
```

Step 10: Verify disk space

1. Display available disk space by entering a `df -k` command. The system responds with a file system table. For example:

Filesystem	kbytes	used	avail	capacity	Mounted on
/dev/md/dsk/d13	1039555	355366	580239	38%	/
/proc	0	0	0	0%	/proc
fd	0	0	0	0%	/dev/fd
swap	1305272	344	1304928	1%	/tmp
/dev/md/dsk/d21	31956788	849300	31107488	3%	/cms

The figures on the `/dev/md/dsk/d21` line are critical. You will use them to determine whether *Solstice DiskSuite*^{*} is administering all your disks.

2. Complete the table below by filling in (1a) and (1b) as instructed.

Disk Space Verification

Total No. of Disks	/cms size in MB (4-GB disks)	/cms size in MB (9-GB disks)	Calculations:
2	3104	6208	(1a) Enter /cms line's 'kbytes' figure divided by 1000 (move decimal three places left): (1a)
4	7137	14274	
6	11170	22340	
8	15203	30406	
10	19236		(1b) Enter "/cms Size" (above left) corresponding to the number of hard disks and disk size: (1b)

3. Compare the figures in (1a) and (1b) of the table. They should be *approximately* equal. The two will not correspond exactly, but they should be within about 10% of the /cms size figure. A major discrepancy between the disk space you should have and the space recognized by your *Solstice DiskSuite* software usually indicates a connectivity problem. You can check connectivity by checking the output of a `metastat` command to see that all the disks are recognized. For example:

^{*}*Solstice DiskSuite* is a trademark of Sun Microsystems, Inc.

```
# metastat

d21: Mirror
  Submirror 0: d19
    State: Okay
  Submirror 1: d20
    State: Okay

d19: Submirror of d21
  State: Okay
  Size: 14398560 blocks
  Stripe 0:
    Device      Start Block  Dbase  State      Hot Spare
    c0t0d0s3      0         No     Okay
  Stripe 1:
    Device      Start Block  Dbase  State      Hot Spare
    c0t2d0s1      0         No     Okay
d20: Submirror of d21
  State: Okay
  Size: 14398560 blocks
  Stripe 0:
    Device      Start Block  Dbase  State      Hot Spare
    c0t1d0s3      0         No     Okay
  Stripe 1:
    Device      Start Block  Dbase  State      Hot Spare
    c0t3d0s1      0         No     Okay
```

If the stripes of d19 and d20 do not account for all the disk drives on your system, check your disk drive connections. If discrepancies persist, telephone the Lucent National Customer Care Center at 1.800.242.2121, or contact your Lucent representative or distributor.

Step 11: Create an alternate boot device

Set up an alternate boot device by doing the following procedure:

1. Enter the following command:

```
# ls -l /dev/rdisk/<newbootdev>
```

where *newbootdev* is the device name of the disk chosen as boot disk 2 (recorded as the d12 metadvice on [page 13](#)).

If your boot disk 2 were c0t2d0s0, for example, the system would respond with something like this:

```
lrwxrwxrwx 1 root root 83 Jun 18 15:23 /dev/rdisk/c0t2d0s0 ->
../../devices/sbus@3,0/SUNW,fas@3,8800000 /sd@2,0:a,raw
```

2. Record the part of the line immediately following */devices* and preceding *:a*. For the example above, you would record the part of the line beginning **/sbus@3,0** and ending **sd@2,0**.

3. Reboot the system with the following command:

```
# /usr/sbin/shutdown -y -i0 -g0
```

The system displays the `ok` prompt.

4. Create a device alias for the new boot device by entering the following command (for example):

```
ok nvalias bootdevice2 <device>
```

Replace *<device>* with what you recorded in item 2. For the example shown, you would replace *<device>* with the following:

```
/sbus@3,0/SUNW,fas@3,88000000/sd@2,0
```

5. Test the alternate boot device name by entering the following command:

```
ok boot bootdevice2
```

6. Reboot the system as usual and verify that everything is normal. Ignore any warnings you may get about a “forceload” failing.

7. Enter the following command:

```
# crontab -e
```

The system starts the editor and loads the `cron` file.

8. Find the line for */olds/chkDisks* and take one of the following actions:

If the line is commented out, uncomment it by deleting the pound sign at the beginning of the line.

If there is no line for */olds/chkDisks*, add the following line to the end of the file:

```
15 0 * * * /olds/chkDisks > /dev/null 2>&1
```

9. Save the file and exit the editor.

10. Enter the following command:

```
# chmod +x /olds/chkDisks
```

Factory installation procedures for disk mirroring

This section describes how a mirrored system is installed on a new *Sun Enterprise 3xxx* computer in the factory. You may use these instructions in the field to bring an existing system up to factory specifications.

Step 1: Install the hardware

Install hardware components as required. For details and instructions, see the appropriate Lucent hardware installation document and the *Sun Microsystems* documents shipped with your machine.

Step 2: Install the operating system

Install the *Solaris* operating system. Partitioning the disks will be part of that process.

To install the *Solaris* operating system, follow the instructions in the *CentreVu CMS R3V8 Software Installation and Setup* (585-210-941) document, **with the following exception:**

- Use the *Disks* and *Partitioning* tables below to determine disk partitioning details. *Disks* shows which disks are boot disks and which are not; *Partitioning* shows partitioning details.

Disks

Disk	<i>Enterprise 3000</i> Device Names	<i>Enterprise 3500</i> Device Names
boot disk 1	c0t0d0	c0t0d0
boot disk 2	c0t1d0 ^a	c1t4d0
other disks	c0t2d0 c0t3d0 c0t10d0 c0t11d0 c0t12d0 c0t13d0 c0t14d0 c0t15d0	c0t1d0 c0t2d0 c0t3d0 c1t5d0 c1t6d0 c1d7d0

^aTrue for factory-installed disk mirroring. Unlikely to be true for systems upgraded to disk mirroring in the field.

Partitioning (Note: **Never** change the default size of the *overlap* slice!)

Disk Type	Part	Tag	Flag	CMS R3V8 4.2-GB Disk		CMS R3V8 9.1-GB Disk SCSI (E3000) and FC-AL (E 3500)	
				Size (cyl.)	Starting Cylinder	Size (cyl.)	Starting Cylinder
Boot	0	<i>boot disk 1: / or root boot disk 2: unassigned</i>	wm	1023	0	616	0
	1	unassigned	wm	7	1023	7	616
	2	overlap <i>or</i> backup	wm	3880	0	4924	0
	3	unassigned	wm	1879	1030	3716	623
	4	swap	wm	971	2109	585	4339
	5-7	unassigned	wm	0	0	0	0
Nonboot	0	unassigned	wm	2	0	2	0
	1	unassigned	wm	3878	2	4922	2
	2	overlap <i>or</i> backup	wm	3880	0	4924	0
	3-7	unassigned	wm	0	0	0	0

Step 3: Install the application software

Install application software as instructed in the *CentreVu CMS R3V8 Software Installation and Setup* (585-210-941) document. That will mean installing the following packages in order:

- Common Desktop Environment (CDE)
- *SunLink** HSI/S drivers (if using HSI/S ports)
- NTS drivers (if using NTS ports)
- X.25 (if using X.25 switch links)
- *INFORMIX*† SQL (optional)
- *INFORMIX* SE
- *INFORMIX* ESQ
- *INFORMIX* ILS
- *Solstice DiskSuite*
- *Sun Solaris* patches

**SunLink* is a trademark of Sun Microsystems, Inc.

†*INFORMIX* is a registered trademark of Informix Software, Inc.

Step 4: Check the *disk.type* file

The *olds* commands, which you will be using to configure the *Solstice DiskSuite* software, obtain information about your disk configuration from the */olds/disk.type* file. If your boot disks are not as specified in the *Disks* table in Step 2, check the *disk.type* file to verify that it names the proper boot disks for your system.

The relevant portion of the *disk.type* file looks like this:

#arch	config	bootdisk1	disk2	scsi control
#Enterprise 3000				
sun4u	notmirrored	c0t0d0	c0t1d0	dontcare
sun4u	mirrored	c0t0d0	c0t1d0	controllers=0

The *bootdisk1* column designates the first boot disk; the *disk2* column designates the second boot disk in a mirrored configuration. In order for the *olds* commands to work, the *disks.type* file **must** match your physical configuration: you must have disks in the locations specified in the file.

Step 5: Set up the disk subsystem for *Solstice DiskSuite*

1. Set the environment to the Korn Shell by entering these commands:

```
# stty erase <Ctrl-H>
# ksh -o vi
```

(where *<ctrl-H>* indicates you should press and hold the Control key as you press the H key).

The *stty* command sets up your backspace key as an actual backspace. If you do not enter this command, you will have to use the Delete key as a backspace.

2. Make sure the CMS CD is loaded in the CD-ROM drive.
3. Copy the *Solstice DiskSuite* setup scripts from the CD to the boot disk by entering the following commands:

```
# mkdir /olds
# cp /cdrom/cdrom0/cms/reloc/rdonly/olds_install/* /olds
# cd /olds
# chmod +x /olds/olds
```

4. Alter your path:

```
# export PATH=$PATH:/usr/opt/SUNWmd/sbin:/olds
#
```

5. Check the disks:

```
# olds -mirrored -check_disks
number of external scsi controllers with disks
is = 0
number of disks is = 4

. . .
disk:c0t0d0 is partitioned ok
disk:c0t1d0 is partitioned ok
disk:c0t2d0 is partitioned ok
disk:c0t3d0 is partitioned ok
Success, checking disks for mirroring.
#
```

6. Create system files for the *Solstice DiskSuite* software:

```
# olds -mirrored -mk_files
#
```

7. Make sure all the disk drives on your system have been recognized.

To do that, read the file */olds/md.tab.new* into an editor and locate the *#/cms* section. It should reflect the precise number of disk drives on your system. The following example shows a total of four disk drives on the system, two in d19 and two in d20:

```
.
.
.
#metaroot
d11 1 1 c0t0d0s0
d12 1 1 c0t2d0s0
d13 -m d11

#/cms
d19 2 1 /dev/rdisk/c0t0d0s3 1 /dev/rdisk/c0t1d0s1
d20 2 1 /dev/rdisk/c0t2d0s3 1 /dev/rdisk/c0t3d0s1
d21 -m d19

#metaswap
d15 1 1 c0t0d0s4
d16 1 1 c0t2d0s4
d17 -m d15
```

If there is a discrepancy between what the system should have recognized and what it *did* recognize, complete the procedure in “The system fails to recognize all disk drives” on page 82, and then return to item 5.

Step 6: Run the setup scripts

1. Create the state database replicas by entering the following command:

```
# olds -mirrored -metadbs
.
```

Success, setting up metadb replicas

2. Mirror the root file system by entering the following command:

```
# olds -mirrored -setroot
.
```

Success, root mirrored successfully.

3. Check disks with “metastat” and wait until all disks are synched.

```
# /usr/sbin/shutdown -y -g0 -i0
.
```

ok



NOTE:

Ignore any FSCK error messages which may be generated.

4. Reboot and log in as the root user (boot -r):

```
ok boot -r
```

```
.
```

```
<hostname> console login: root
Password: <password>
#
```

5. Setup the */cms* metadevices by entering the following commands:

```
# ksh -o vi
# export PATH=$PATH:/olds
# olds -mirrored -setup
.
.
.
3962848, 13997552, 14032256, 14066960, 14101664,
14136368, 14171072,
  14205776, 14240480, 14275184, 14309888,
14344592, 14376992, 14411696,
  14446400, 14481104, 14515808, 14550512,
14585216, 14619920, 14654624,
  14689328, 14724032, 14758736, 14793440,
14828144, 14862848, 14897552,
  14929952,
ufs fsck: sanity check: /dev/md/rdsk/d21 okay
Success, /cms mirrored successfully
#
```

6. Create and mount the */cms* file system by entering the following commands:

```
# mkdir /cms
# mount /cms
```

Step 7: Verify disk space

1. Display available disk space by entering a `df -k` command. The system responds with a file system table. For example:

#	df -k					
	Filesystem	kbytes	used	avail	capacity	Mounted on
	/dev/md/dsk/d13	772327	236870	458227	35%	/
	/proc	0	0	0	0%	/proc
	fd	0	0	0	0%	/dev/fd
	/dev/md/dsk/d21	xxxxxxx	562168	14346985	0%	/cms

The figures on the */dev/md/dsk/d21* line are critical. You will use them to determine whether *DiskSuite* is administering all your disks.

2. Complete the table below by filling in (1a) and (1b) as instructed.

Disk Space Verification

Total No. of Disks	/cms size in MB (4-GB disks)	/cms size in MB (9-GB disks)	Calculations:	
2	3104	6208	(1a) Enter /cms line's 'kbytes' figure divided by 1000 (move decimal three places left):	
4	7137	14274		
6	11170	22340		
8	15203	30406		
10	19236		(1b) Enter "/cms Size" (above left) corresponding to the number of hard disks and disk size:	

3. Compare the figures in (1a) and (1b). They should be *approximately* equal. The two will not correspond exactly, but they should be within about 10% of the /cms size figure. A major discrepancy between the disk space you should have and the space recognized by the software usually indicates a connectivity problem. You can check connectivity by checking the output of a `metastat` command to see that all the disks are recognized. For example:

```
# metastat

d21: Mirror
  Submirror 0: d19
    State: Okay
  Submirror 1: d20
    State: Okay

d19: Submirror of d21
  State: Okay
  Size: 14398560 blocks
  Stripe 0:
    Device      Start Block  Dbase  State      Hot Spare
    c0t0d0s3      0         No      Okay
  Stripe 1:
    Device      Start Block  Dbase  State      Hot Spare
    c0t2d0s1      0         No      Okay

d20: Submirror of d21
  State: Okay
  Size: 14398560 blocks
  Stripe 0:
    Device      Start Block  Dbase  State      Hot Spare
    c0t1d0s3      0         No      Okay
  Stripe 1:
    Device      Start Block  Dbase  State      Hot Spare
    c0t3d0s1      0         No      Okay
```

If the stripes of d19 and d20 do not account for all the disk drives on your system, check your disk drive connections. If discrepancies persist, telephone the Lucent National Customer Care Center at 1.800.242.2121, or contact your Lucent representative or distributor.

Step 8: Install the CMS software

Now you will complete the installation of the usual software by installing CMS. In general, the steps are as listed below. Detailed installation instructions can be found in *CentreVu CMS R3V8 Software Installation and Setup* (585-210-941).

The steps to be completed are these:

- Install CMS software
- Install CMS patches (if needed)
- Install CMS Supplemental Services software
- Set up CMS
- Install Open Database Connectivity (ODBC) software (optional)
- Set up the remote console
- Administer the NTS (if using NTS ports)

Step 9: Create an alternate boot device

1. Enter the following command:

```
# ls -l /dev/rdisk/newbootdev
```

where *newbootdev* is the device name of the disk chosen as boot disk 2 (as **c0t1d0s0**, for example). You partitioned that disk earlier in the installation procedure.

The system responds (for example):

```
lrwxrwxrwx  1 root    root          83 Jun 18 15:23
/dev/rdisk/c0t1d0s0 -> ../../devices/sbus@3,0/SUNW,
fas@3,88000000/sd@1,0:a,raw
```

2. Record the device definition (the part after `/devices` and before `:a,raw`) . For the example above, you would record the part beginning **/sbus@3,0** and ending **sd@1,0**.
3. Reboot the system by entering the following command:

```
# /usr/sbin/shutdown -y -i0 -g0
```

The system displays the `ok` prompt.

4. Create a device alias for the new boot device by entering the following command (for example):

```
ok nvalias bootdevice2 <device>
```

Replace *<device>* with what you recorded in item 2. For the example shown, you would replace *<device>* with the following:

```
/sbus@3,0/SUNW, fas@3,88000000/sd@2,0
```

5. Enter the following command:

```
ok devalias
```

6. Check the output for a line like this:

```
bootdevice2 <device definition>
```

where, again, <device definition> is the string recorded in item 2.

7. Verify that the system boots normally off the alternate boot device, and log in as the root user:

```
ok boot bootdevice2
.
.
.
<hostname> console login: root
Password: <password>
.
.
.
#
```

8. Reboot with a shutdown command and again log in as *root*:

```
# /usr/sbin/shutdown -y -g0 -i6
.
.
.
<hostname> console login: root
Password: <password>
.
.
.
#
```

9. Verify that everything is normal.

10. Enter the following command:

```
# crontab -e
```

The system starts the editor and loads the cron file.

11. Add the following line to the end of the file:

```
15 0 * * * /olds/chkDisks > /dev/null 2>&1
```

12. Save the file and exit the editor.

13. Enter the following command:

```
# chmod +x /olds/chkDisks
```

14. Do a CMSADM backup. For instructions, see the *CentreVu CentreVu CMS R3V8 Administration* (585-210-910) document or the maintenance and troubleshooting document for your hardware platform.

When the CMSADM backup is finished, the factory installation procedure for disk mirroring is complete.

Maintaining Mirrored Systems

This chapter contains procedures you are likely to need to maintain a disk-mirrored system, including replacing disk drives, recovering from disk failures, and using the olds scripts.

Using the alternate boot device

When you set up your mirrored system, you created an alternate boot device named *bootdevice2*.

The alternate boot device is a fail-safe device. If for some reason you are unable to boot from the primary boot disk—a corrupted boot disk, for example—you can tell the system to boot from the alternate device until further notice.

You do that by entering one of the following commands, depending upon whether you are at the boot prompt or in root:

```
ok setenv boot-device bootdevice2
# eeprom boot-device=bootdevice2
```

To reset the boot device to the primary disk, enter one of the following commands, again depending upon whether you are at the boot prompt or in root:

```
ok setenv boot-device disk
# eeprom boot-device=disk
```

Maintaining the chkDisks crontab

The chkDisks crontab runs each night and checks to see whether any potential or actual drive problems have been logged. The results of the search are logged to */olds/elog* and mailed to the *root* user. This section shows how to start and stop chkDisks, and how to verify that it will run.

Activating chkDisks

Activate chkDisks by doing the following steps:

1. Enter the command: `# crontab -e`

The system starts the editor and loads the cron file.

2. Find the line for */olds/chkDisks*. If it is commented out, uncomment it by deleting the pound sign at the beginning of the line. If there is no line for */olds/chkDisks*, add the following line to the end of the file:

```
15 0 * * * /olds/chkDisks > /dev/null 2>&1
```

Verifying chkDisks

To verify that cron is running, enter the following command:

```
# crontab -l
```

Check the listing to see that there is an entry for chkDisks.

Changing the scheduled run time

The */olds/chkDisks* line in the cron file is generally in the following format:

```
15 0 * * * /olds/chkDisks > /dev/null 2>&1
```

This line tells the system to run chkDisks every day at 15 minutes past hour zero, or 12:15 AM. You can change that schedule by changing the first five fields as necessary. The fields, in order of appearance, are: minute, hour, day of the month, month of the year, and day of the week. An asterisk means “all legal values.” For more information, see the manual (man) page for the crontab command.

Canceling chkDisks

To stop cron from running, enter the following command:

```
# crontab -e
```

With the file loaded in the editor, comment out the entry for chkDisks and write and quit the file.

Doing a CMSADM backup on a mirrored system

This section describes how to perform a traditional CMSADM backup on a mirrored system.

NOTE:

Mirroring is intended as a way to minimize data loss should a system crash occur, not as a way to avoid backing up data. Mirrored systems must be backed up just as often as unmirrored systems.

Step 1: Print the *vfstab* file

Print out the contents of the *vfstab* file before beginning the backup. You will need the information in that file to restore files from the backup tape.

To print the *vfstab* file, do the following:

1. Enter the following command:

```
# lp /etc/vfstab
```

2. Retrieve the printout and save it. You will need it again when the backup finishes.

Step 2: Run the backup

1. At the system console, log in as *root*.
2. Verify that the computer is in a *Solaris* multi-user state (a run level of 2 or 3) by entering a `who -r` command. The system responds by listing the run level and date. For example:

```
# who -r
. run-level 3  Feb 2 16:52 3 0 S
#
```

If the system is in some other run level, reboot and log in again as the root user:

```
# /usr/sbin/shutdown -y -g0 -i6
. . .
<hostname> login: root
password: <password>
. . .
#
```

3. Display the *CentreVu* CMS Administration menu by entering a `cmsadm` command. For example:

```
# cmsadm
Lucent Technologies CentreVu(R) Call Management
System Administration Menu
Select a command from the list below.
  1) acd_create  Define a new ACD
      . . . . .
```

4. Enter the number of the backup option. The system prompts for the tape drive.
5. Enter the appropriate option number. The system begins calculating the approximate number of tapes required and informs you of its calculation. For example:

```
The backup will need approximately 2 tapes.
You will be prompted for additional tapes.
Be sure to number the cartridge tapes
consecutively in the order they will be
inserted.
Please insert the first cartridge tape into
/dev/rmt/<xxx>
Press ENTER when ready:
```

6. Insert the tape, wait for it to rewind and reposition, and press Enter.

The backup begins.

During the backup, you may receive a prompt about CMS being on or prompts to insert more tapes. Respond to each prompt as appropriate, and then press Enter.

When the backup completes, it will be verified. If you used more than one tape, you will have to reinsert each tape.

Any time you need to insert a tape, allow it to rewind and reposition before you press Enter.

Step 3: Label and store the tapes

After the tape verification, the system prompts you to label the tapes and then returns you to the system prompt:

```
Please label the backup tape(s) with the date and
the current CMS version (<version>).
#
```

Label the tapes as instructed. Bundle them with the *vfstab* printout and put them away in a safe place.

Restoring from a CMSADM Backup

This section contains the procedure for restoring /cms data files from a traditional CMSADM backup tape.

You may want to perform a CMSADM restore when, for example, your file systems are still intact, but some data has been corrupted.

If your CMS is running on a *Sun Enterprise* 3500, and you have had a crashed or corrupted disk, check the `/kernel/drv/st.conf` file before you begin the restore. Make certain the file contains the following entry:

```
# Begin CMS tape configuration list.
tape-config-list=
"EXABYTE  EXB-8900", "Mammoth  EXB-8900  8mm  Helical
Scan", "EXB-8900",
"TANDBERG TDC 4200", "Tandberg 2.5 Gig QIC", "TAND-25G-
FIXED",
"TANDBERG SLR5", "Tandberg 8 Gig QIC", "TAND-8G-FIXED";
EXB-8900  = 1,0x29,0,0xce39,4,0x7f,0x7f,0x7f,0x7f,0;
TAND-25G-FIXED  = 1,0x37,512,0x867a,1,0x00,0;
TAND-8G-FIXED                                     =
1,0x37,512,0x963a,4,0xA0,0xD0,0xD0,0xD0,3;
# End CMS Tape configuration list.
```

If the file does not contain the entry, add it. The 3500's Mammoth 8mm tape drive will not operate if the entry is absent.

To restore a CMSADM backup, do these steps:

1. Obtain the latest CMSADM backup tape.
2. Load the backup tape into the tape drive.
3. Enter the following command:

```
# cpio  -icmudv  -C 10240 -I <device>
```

where `<device>` is one of the following:

<code>/dev/rmt/0</code>	The internal, noncompressing tape drive (14-GB, 8-mm drive) with the lowest target address.
<code>/dev/rmt/1</code>	The external, noncompressing tape drive (QIC-150 or 5-GB, 8-mm drive) with the second lowest target address.
<code>/dev/rmt/0c</code>	The internal, compression-capable tape drive (usually a 14-GB or 40-GB) with the lowest target address.
<code>/dev/rmt/1c</code>	The external, compression-capable tape drive (either a QIC 2.5-GB or a 14-GB tape drive) with the second lowest target address.

 NOTE:

You may get error messages concerning the */home* directory. These errors display when the directory is already present. Ignore them.

4. Restore any *CentreVu* CMS maintenance backups dated after the latest CDMADM backup. See the *CentreVu CMS R3V8 Administration* (585-210-941) document for details on restoring a maintenance backup.

Repairing or rebuilding the /cms file system

This section contains procedures for reinitializing the /cms file system. You may need to perform this procedure in case of disk corruption or some other catastrophic system problem. The point of the procedure is to repair /cms or, failing that, to rebuild it and restore the latest available CMS data. Try to repair the file system first: if you can repair it, you will save considerable time and trouble.

Repairing /cms

This procedure attempts to repair /cms.

1. Turn off CMS.
 - a. Enter a cmsadm or cmssvc command.
 - b. Select the run_cms option.
 - c. Select the Turn off CMS option.

2. Unmount /cms:

```
# umount /cms
#
```

The system prompt should return without messages, as shown.

3. **Do this step only if the umount command returned the response /cms: device busy.**

- a. Load the /etc/vfstab file into your editor.
- b. Find the /cms and /cms/swap lines. Comment out both lines. (To comment out a line, insert a pound sign at the beginning.)
- c. Save the file and exit the editor.
- d. Reboot by entering the following command:

```
# /usr/sbin/shutdown -i6 -g0 -y
```

4. Attempt to repair /cms:

```
# fsck -y /dev/md/rdisk/d21
** /dev/md/dsk/d21
** Last mounted on /cms
** Phase 1 - Check Blocks and Sizes
** Phase 2 - Check Pathnames
** Phase 3 - Check Connectivity
** Phase 4 - Check Reference Counts
** Phase 5 Check Cyl groups
2060 files, 564453 used, 6468748 free (1900
frags, 808356 blocks, 0.0%
#
```

The system prompt should return without error, as shown above. (Of course, file and block counts in the last line will vary.) In that case, continue with the next step.

If you get an error message, however, repeat the `fsck` command. If the command still returns errors after 10 repetitions, stop this procedure and skip ahead to “Rebuilding /cms,” below.

Continue with step 6 only if you entered the `fsck` command and it returned the system prompt without error messages.

5. This step remounts `/cms` and reallocates the swap file. How that is done depends upon whether you altered the `vfstab` file earlier in this procedure.

If you altered the `vfstab` file earlier (step 4), do the following:

- a. Load the `/etc/vfstab` file into your editor.
- b. Find the `#/cms` line. Uncomment the line. (To uncomment a line, delete the initial pound sign.)
- c. Save the file and exit the editor.
- d. Reboot by entering the following command:

```
# /usr/sbin/shutdown -i6 -g0 -y
```

If you did NOT alter the `vfstab` file, do the following:

- a. Remount `/cms`:

```
# mount /cms
```

6. Turn on CMS:

- a. Enter a `cmsadm` or `cmssvc` command.
- b. Select the `run_cms` option.
- c. Select the Turn on CMS option.

The file system has now been repaired and should be operating normally.

If this procedure completed successfully, do NOT do the “Rebuild /cms” procedure.

Rebuilding /cms

Do this procedure only if you have tried the `fsck` command at least ten times and it is still returning errors.

If you are running a *Sun Enterprise 3500* platform, do not begin this procedure until you have checked the `/kernel/drv/st.conf` file to make sure the 3500's 8mm tape drive is properly defined. See [“Restoring from a CMSADM Backup” on page 43](#) for details.

To complete this procedure, you need the latest CMSADM backup. Also, `/cms` should still be unmounted at this point. If anything has been done to alter that, repeat steps 2 through 4 of the “Repairing /cms” procedure.

1. Add the *Solstice DiskSuite* directories to your path:

```
# export PATH=$PATH:/usr/opt/SUNWmd/sbin:/olds
#
```

2. Detach the d20 submirror:

```
# metadetach d21 d20
d21: submirror d20 is detached
#
```

3. Clear the d20 submirror:

```
# metaclear d20
d20: Concat/Stripe is cleared
#
```

4. Clear the cms mirror:

```
# metaclear d21
mirror is cleared
#
```

5. Clear the d19 submirror:

```
# metaclear d19
d19: Concat/Stripe is cleared
#
```

6. Reinitialize the d19 submirror:

```
# metainit d19
d19: Concat/Stripe is set up
#
```

7. Reinitialize the d20 submirror:

```
# metainit d20
d20: Concat/Stripe is set up
#
```

8. Reinitialize the cms mirror:

```
# metainit d21
mirror is set up
#
```

9. Prepare the cms mirror for a new file system (this will take 15 to 30 minutes to complete):

```
# newfs -m 0 /dev/md/rdisk/d21
.
.
.
14550512, 14585216, 14619920, 14654624,
14689328, 14724032, 14758736, 14793440,
14828144, 14862848, 14897552, 14929952,
ufs fsck: sanity check: /dev/md/rdisk/d21 okay
Success, /cms mirrored successfully
#
```

10. This step remounts /cms. How that is done depends upon whether you had to alter the *vfstab* file when you attempted to repair the file system.

If you altered the *vfstab* file during the repair procedure, do the following:

- Load the */etc/vfstab* file into your editor.
- Find the *#/cms* line. Uncomment the line. (To uncomment a line, delete the initial pound sign.)
- Save the file and exit the editor.
- Reboot by entering the following command:

```
# /usr/sbin/shutdown -i6 -g0 -y
```

If you did NOT alter the *vfstab* file, do the following:

- Remount /cms:

```
# mount /cms
```

11. Verify that you are in the root directory.
12. Load the latest CMSADM backup tape into the tape drive.
13. Restore CMS data from the tape:

```
# cpio -icmudv -C 10240 -I /dev/rmt/0c -M
"Insert Tape %d" "cms/*"
```

14. Restore the latest full maintenance backup.

Restoring a mirrored system from a CMSADM backup

This section shows you how to restore an entire mirrored system from a CMSADM backup. You might have to do this procedure, for example, if you have a disk crash that disables the operating system disk drive.

To do this procedure, you will need the “*Solaris 7 3/99*” compact disc and the most recent CMSADM backup tape. You will also need to know your system’s host name, host ID, and IP address.

Step 1: Power off the system

Power off the system by doing the following tasks:

1. Turn off the system unit.
2. Turn off the system monitor.
3. Turn off all external devices starting with the device closest to the system unit and working toward the farthest device.

Step 2: Install the new drives

CAUTION:

You must wear an ESD wrist strap when installing or removing hard disk drives to prevent electrical discharge that may harm system components.

In Sun Enterprise 3xxx systems, the internal hard disk drives install in the front of the system. To expose the disk drive bays, open the front access door.

Each disk drive has a retainer latch/drive handle on the front of the drive. To install a drive, unlatch and extend the retainer and slide the drive carefully into the appropriate drive slot. When the drive stops, gently apply pressure to the face of the drive until you hear the connectors engage. Then lower the retainer latch and snap it in place, making certain the drive is secure and does not extend beyond the front of the chassis.

Do not use excessive force to seat the connector or to close the retainer latch.

For additional information about hard disks and disk drives, see the *Sun System Reference Manual* appropriate to your hardware platform.

Step 3: Boot the system from the compact disk

The following tasks boot the system from the compact disc.

Task	Action	Result
1	Turn on the power to the system units in the opposite order in which you powered them off.	The system begins to boot.
2	After the display console banner appears but before the operating system begins to boot, interrupt the process by pressing Stop A .	The ok prompt is displayed
3	Load the “ <i>Solaris 7 3/99</i> ” compact disc into the CD-ROM drive.	
4	Enter the following command: ok boot cdrom	Booting takes from 2 to 10 minutes. When it completes, Select Language and Locale screen is displayed:
5	Click the Continue button.	The Solaris Installation Program screen is displayed
6	Click the Continue button.	The Identify this System screen is displayed
7	Click the Continue button.	The Host Name screen is displayed

Step 4: Identify the system

The following tasks identify your system for the installation.

Task	Action	Result
1	<i>No action required.</i>	The Host Name screen is displayed.
2	a. Click the <i>Host name</i> box b. Type the host name for the workstation. c. Click the Continue button.	The Network Connectivity screen is displayed.
3	a. Click the Yes option. b. Click the Continue button. c. If your system has more than one network board, a <i>Primary Network Interface</i> screen displays. In that case, select hme0 and click Continue.	The IP Address screen is displayed.
4	a. Click the <i>IP address</i> box b. Type your system's IP address. If you are unsure, use the one that is NOT commented out. c. Click the Continue button.	The Confirm Information screen is displayed.
5	a. Check the information displayed on the screen. If the information is not correct, click Change and return to task 1 of this procedure. b. If the information is correct, click Continue.	The Name Service screen is displayed.

Step 5: Set the Name Service options

The following tasks set the name service options.

Task	Action	Result
1	<i>No Action Required.</i>	The Name Service screen is displayed.
2	a. Select the None option. b. Click the Continue button.	The Confirm Information screen is displayed.
3	a. Check the information displayed on the screen. If it is incorrect, click Change and return to task 1 of this table. b. If the information is correct, click Continue.	The Subnets screen is displayed.
4	Do this task only if the switch is connected to this computer via a serial port or HSI card. a. Select No. b. Click Continue. c. Skip the remainder of the tasks in this table.	The Time Zone screen is displayed.
5	Do this task and the next only if the switch is connected to this computer via a network. a. Select Yes. b. Click Continue.	The Netmask screen is displayed.
6	a. Enter the appropriate net mask. The default of 255.255.255.0 is recommended. b. Click Continue.	The Time Zone screen is displayed.

Step 6: Set the date and time

The following tasks set the date and time.

Task	Action	Result
1	No action required.	The Time Zone screen is displayed.
2	a. Select Geographic region. b. Click Set.	The Geographic Region screen is displayed.
3	a. Select the region where this system is located. b. Select the time zone where this system is located. c. Click the Continue button.	The Date and Time screen is displayed.
4	a. Change the date and time as needed to reflect the current local date and time. b. Click Continue.	The Confirm Information screen is displayed.
5	a. Check the information displayed on the screen. If it is incorrect, click Change and return to task 1. b. If the information is correct, click Continue.	The Solaris Interactive Installation - Initial screen is displayed.

Step 7: Select the operating system files to be installed

This procedure selects *Solaris* system files for the installation.

Task	Action	Result
1	No Action Required.	The Solaris Interactive Installation - Initial screen is displayed.
2	Click the Initial button.	The Solaris Interactive Installation screen is displayed.
3	Click the Continue button.	The Allocate Client Services? screen is displayed.
4	Click the Continue button.	The Select Languages screen is displayed.
5	a. Select the languages you want to see displayed in the user interface. b. Click the Continue button.	The Select Software screen is displayed.
6	a. Select End User System Support. b. Click the Customize button.	The Customize Software screen is displayed.
7	Select the packages listed below: — Basic Networking — On-Line Manual Pages — X Windows system online user man pages (under Open Windows Version 3) — Point-to-Point Protocol — CCS tools bundled with SunOS and Solaris bundled tools (under Programming tools and libraries) — System Accounting — Terminal Information Click OK when finished.	The Select Software screen is displayed.
8	Click the Continue button.	The Select Disks screen is displayed.

Step 8: Select the hard disks


This procedure selects hard disks for partitioning.

Task	Action	Result
1	No Action Required.	The Select Disks screen is displayed.
2	Check the list of disk drives in the <i>Available</i> column: it should list every disk in your system. If it doesn't, you probably have a connectivity problem. In that case stop the installation, power down, secure all drive connections, and restart the installation.	
3	a. Click the disk name to highlight an available disk. b. Click Add to select the disk. c. Repeat <i>a</i> and <i>b</i> until all available disks are in the Selected column. d. Click Continue.	The Preserve Data? screen is displayed.
4	Click Continue.	The Automatically Layout File Systems? screen is displayed.
5	Click Manual Layout.	The File System and Disk Layout screen is displayed.
6	Click Customize.	The Customize Disks screen is displayed.

Step 9: Set up the disk partitions

This procedure sets up the disk partitions.

Task	Action	Result
1	No action required.	The Customize Disks screen is displaying. This screen lists partitioning for two disks, one in the left panel and another (if another exists) in the right panel.
2	Click the cylinders icon for the disk you want to partition.	The Customize Disks by Cylinder screen is displayed.

3	a. Click the space at the right of the slice number. b. Type the slice name (see the table below). c. Click the next space to the right and type the size (see the table below). d. Repeat a through c until the disk is properly partitioned.	
4	Click the OK button.	The Customize Disks screen is displayed.
5	Repeat the tasks in this table for every disk drive in the system. Do not go on until all drives in the system are partitioned.	

Standard Mirrored-disk Partitioning

(Note: **Never** change the default size of the overlap slice!)

			Size (Cylinders)	
Device	Slice	Name	CMS R3V8 4.2-GB Disk	CMS R3V8 9.1-GB Disk SCSI (E 3000) and FC-AL (E 3500)
c0t0d0 c0t1d0 (3000 only) c1t4d0 (3500 only)	0	c0t0d0: / c0t1d0: (leave blank) c1t4d0: (leave blank)	1023	616
	1	(leave blank)	7	7
	2	overlap	3880	4924
	3	(leave blank)	1879	3716
	4	swap	971	585
	5-7	(leave blank)	0	0
All others	0	(leave blank)	2	2
	1	(leave blank)	3878	4922
	2	overlap	3880	4924
	3-7	(leave blank)	0	0

Step 10: Install the operating system

This procedure installs the *Solaris 7 3/99* operating system using the system software and disk partitioning already specified.

Task	Action	Result
1	No action required.	The Customize Disks screen is displayed.
2	Click OK.	The File System and Disk Layout screen is displayed.
3	Click Continue.	The Mount Remote File System? screen is displayed.
4	Click Continue.	The Profile screen is displayed.
5	Click Begin Installation.	The Reboot/No Reboot screen is displayed.
6	Click Reboot.	
7	<i>No action required.</i> As the disks are partitioned and the system files are copied to the disk, the progress gauge is updated to indicate the progress of the installation, and takes about 40 minutes to reach 100%. At that point the progress window disappears. A console window appears in the upper left corner of the screen and the installation continues. The remainder of the installation can take several hours, depending upon the number of disks installed.	

Step 11: Complete the OS installation

This procedure completes the operating system installation.

Task	Action	Result
1	<i>No action required.</i> When the installation completes, the machine reboots and prompts for a root password.	On this screen you can create a root password. . . . Root password: Press Return to continue.
2	Enter the same password you had before you began the installation.	Please reenter your root password? Press Return to continue.
3	Enter the password again.	 syslog services starting. Print services started. volume management starting. The system is ready. <hostname> console login:
4	Log in as the root user.	#

Step 12: Alter the tape configuration file (*Enterprise 3500 only*)

If you are running on a *Sun Enterprise 3500* platform, you must alter the */kernel/drv/st.conf* file so the operating system recognizes the 3500's 8mm tape drive. To accomplish that, do the following steps:

1. Start an editor such as vi or the CDE Text Editor, and load the */kernel/drv/st.conf* file.
2. Find the `#tape-config-list=` line. It should be in the vicinity of line number 40.
3. Insert, immediately before that line, the following:

```
# Begin CMS tape configuration list.  
tape-config-list=  
"EXABYTE EXB-8900","Mammoth EXB-8900 8mm Helical  
Scan","EXB-8900",  
"TANDBERG TDC 4200", "Tandberg 2.5 Gig QIC", "TAND-25G-  
FIXED",  
"TANDBERG SLR5", "Tandberg 8 Gig QIC", "TAND-8G-FIXED";  
EXB-8900 = 1,0x29,0,0xce39,4,0x7f,0x7f,0x7f,0x7f,0;  
TAND-25G-FIXED = 1,0x37,512,0x867a,1,0x00,0;  
TAND-8G-FIXED = 1,0x37,512,0x963a,4,0xA0,0xD0,0xD0,0xD0,3;  
# End CMS Tape configuration list.
```

4. Save the file and exit the editor.

Step 13: Restore the backup

With the system prompt displaying, insert the CMSADM backup tape into the tape drive and enter the following command:

```
# cpio -icmudfv -C 10240 -I /dev/rmt/0c  
-M "Insert Tape %d" "/etc/vfstab" "/cms/*"  
"/etc/path_to_inst" "/etc/mnttab" "/swap*"
```

This command restores essential system files. As the restoration proceeds, you may ignore all of the following messages:

- cpio cannot create temporary files
- errno 18
- errno 89
- Cross-device link

Step 14: Reestablish the disk mirrors

1. Clean up the system files by entering the following command:

```
# /olds/olds -cleanup
```

The system responds (for example):

```
number of external scsi controllers with disks  
is = 0  
number of disks is = 4
```

All data on CMS filesystems will be removed and filesystems will be unmirrored, are you sure?

2. Enter: `y`
The CMS file systems are cleaned up and the system prompt returns.
3. Reboot the system and log in as the root user:

```
# /usr/sbin/shutdown -y -i6 -g0  
  
.  
.  
.  
<hostname> login: root  
password: <password>  
  
.  
.  
.  
#
```

4. Add the *Solstice DiskSuite* directories to your path:

```
# PATH=$PATH:/usr/opt/SUNWmd/sbin:/olds
# export PATH
```

5. Create md.tab.new:

```
# olds - mirrored -mk_files
.
.
.
Success, creating md.tab.new
#
```

6. Check the disk:

```
# olds -mirrored -check_disk
.
.
.
Success, checking md.tab.new
#
```

7. Activate the state database:

```
# olds -mirrored -meatdbs
.
.
.
Success, setting up metadb replicas
#
```

8. Set up the root mirror:

```
# olds -mirrored -setroot
.
.
.
Success, root mirrored successfully.
#
```

9. Get to the boot prompt:

```
# /usr/sbin/shutdown -y -i0 -g0
.
.
.
ok
```

10. Reboot using a boot -r command and log in as the root user:

```
ok boot -r
.
.
.
<hostname> login: root
Password: <root password>
.
.
.
#
```

11. Add the *Solstice DiskSuite* directories to your path:

```
# PATH=$PATH:/usr/opt/SUNWmd/sbin:/olds
# export PATH
```

12. Activate the */cms* metadvice:

```
# olds -mirrored -setup
.
.
.
14550512, 14585216, 14619920, 14654624,
14689328, 14724032, 14758736, 14793440,
14828144, 14862848, 14897552, 14929952,
ufs fsck: sanity check: /dev/md/rdisk/d21 okay
Success, /cms mirrored successfully
#
```

13. Mount */cms*:

```
# mount /cms
```

14. Load the latest CMSADM backup tape into the tape drive.

15. Restore CMS data from the tape:

```
# cpio -icmudv -C 10240 -I /dev/rmt/0c -M
"Insert Tape %d" "/cms/*"
```

16. Turn on CMS:

- a. Enter a cmsadm or cmssvc command.
- b. Select the run_cms option.
- c. Select the Turn on CMS option.

17. Do a full maintenance backup.

Restoring specific files from a CMSADM backup

To restore specific files from CMSADM backup tapes, enter the following commands at the system prompt (for example):

```
# cd /
# cpio -icmudv -C 10240 -I /dev/rmt/0c -M "Please
remove the current tape, insert tape number %d, and
press ENTER" <full_path_name>
```

(where *<full_path_name>* is the full path name of the file to be restored)

Note the device name in the example (*/dev/rmt/0c*). That name is usually correct. Some other device name may be required, however, depending upon the drive's SCSI ID and compression capability. Tape device names are as follows:

- | | |
|--------------------|---|
| <i>/dev/rmt/0</i> | The internal, noncompressing tape drive (14-GB, 8-mm drive) with the lowest target address. |
| <i>/dev/rmt/1</i> | The external, noncompressing tape drive (QIC-150 or 5-GB, 8-mm drive) with the second lowest target address. |
| <i>/dev/rmt/0c</i> | The internal, compression-capable tape drive (usually a 14-GB or 40-GB) with the lowest target address. |
| <i>/dev/rmt/1c</i> | The external, compression-capable tape drive (either a QIC 2.5-GB or a 14-GB tape drive) with the second lowest target address. |

Adding new disks to a standard mirrored system

The procedure in this section adds new disks to a system that is already mirrored. In order to use this procedure, your assigned boot disks must be the same as those assigned during a factory installation of disk mirroring. If your boot disk assignments differ from the standard factory arrangement (see the table immediately below), you must add the disks using the procedure in [“Adding new disks to a nonstandard mirrored system” on page 66](#).

Boot-Disk Assignments for Factory-Installed Mirroring

Boot Disk 1	Boot Disk 2 (E3000)	Boot Disk 2 (E3500)
c0t0d0	c0t1d0	c1t4d0

New *Enterprise* 3000 internal disks are “hot swappable,” meaning they can be inserted into a running system. The operating system, however, will not recognize the new disks until the system has been rebooted.

1. Enter the command:

```
# /usr/opt/SUNWmd/sbin/metastat
```

Examine the output for the status of each metadevice. (It’s indicated in the `Status: . . . .` line.) **All metadevices must have a status of okay before you may proceed. Do not attempt this procedure while any metadevice is resynching or needs maintenance.**

2. Do a CMSADM backup. For detailed instructions, see the Lucent Technologies Administration or Upgrades document for your version of CMS.
3. Turn off CMS:
 - a. Enter a `cmsadm` or `cmssvc` command.
 - b. Select the `cms_run` option.
 - c. Select the `Turn off cms` option.

4. Halt the system:

```
# /usr/sbin/shutdown -y -i0 -g0
```

5. Install the new disk drives by plugging them into the appropriate empty drive slots. Follow the instructions you received with the disk drives or with your basic system, or see the Lucent Technologies installation or maintenance document for your version of CMS.

6. Reboot the system and log in as the root user:

```
ok boot -r

. . .
<hostname> login: root
password: <password>
. . .
#
```

7. Partition the new disks by following the instructions in [“Partitioning disks with the format command” on page 76](#).
8. Check disk space by entering a `df -k` command. For example:

```
# df -k
Filesystem      kbytes  used   avail  capacity Mounted on
/dev/md/dsk/d13  772038  538683 156155    78%      /
/proc            0         0         0     0%      /proc
fd               0         0         0     0%      /dev/fd
/dev/md/dsk/d21 3059955 192497 2867458    7%      /cms
#
```

Record the *avail* figure for */cms* here: _____

9. Create new versions of the setup files by entering the following commands:

```
# cd /olds
# olds -mirrored -mk_files <newdisk1> <newdisk2>
```

<newdisk1> and *<newdisk2>* represent the device names of the disk drives added (remember that in a mirrored system, disk drives are always added in pairs). For example:

```
# olds -mirrored -mk_files c0t10d0 c0t11d0
```

10. Load the file */etc/opt/SUNWmd/md.tab.new* into an editor and make certain it names all of your disks, including the ones you just added. (For an example of the *md.tab* file format, see [“Examples of valid system files” on page 79](#).)
11. Save a copy of the existing *md.tab* file by entering the following command:

```
# cp /etc/opt/SUNWmd/md.tab /etc/opt/SUNWmd/md.tab.mirror
```

12. Replace the existing *md.tab* file with the one you created in [Step 10](#):

```
# cp /etc/opt/SUNWmd/md.tab.new /etc/opt/SUNWmd/md.tab
```

13. Add the new disks to the existing metadevices by entering the following command:

```
# olds -mirrored -setup <newdisk1> <newdisk2>
```

(where, again, <newdisk1> and <newdisk2> are the device names of the disks you have added).

The system responds with the following:

```
Warning: Current Disk has mounted partitions.
device: c0t6d0 will not be used

. . .
d19: component is attached
d20: component is attached

. . .
d21: submirror d20 is attached
    Success, growing d21
    Success, the capacity of the /cms mirror has increased.
#
```

14. Restart CMS by entering the `cmsadm` or `cmssvc` command and selecting the `cms_run` option. When the run option menu displays, select `start cms`.
15. Check the available disk space by entering a `df -k` command. For example:

```
# df -k
Filesystem      kbytes    used   avail capacity  Mounted on
/dev/md/dsk/d13  772038  538683  156155    78%      /
/proc              0         0         0     0%    /proc
fd                 0         0         0     0%    /dev/fd
/dev/md/dsk/d21 3059955  192497 4967458     7%    /cms
```

Compare the *avail* figure for */cms* to the figure you recorded in [Step 8](#); it should show an increase roughly equal to one-half of the total capacity of the drives added.

Adding new disks to a nonstandard mirrored system

A “nonstandard” disk drive addition is one that cannot take advantage of advanced features of the *olds* setup scripts because the disk configuration is not standard. The procedure is as follows:

1. Turn off CMS. (From the command line, enter the `cmsadm` or `cmssvc` command, and select the `cms_run` option. When the run option menu displays, select `Turn off cms.`)
2. Halt the system by entering the following command:

```
# /usr/sbin/shutdown -y -g0 -i0
```

3. Install the new disk drives by plugging them into the appropriate empty drive slots.

Follow the instructions you received with the disk drives or with your basic system, or see the hardware installation or maintenance and troubleshooting document for your platform and version of CMS.
4. Reboot the system by entering the following command:

```
ok boot -r
```

5. Determine the total amount of disk space available to `/cms` by entering the following command:

```
# df -k /cms
```

The system responds by displaying file system information for `/cms`. For example:

Filesystem	kbytes	used	avail	capacity	Mounted on
/dev/md/dsk/d21	xxxxxxx	9	yyyyyy	0%	/cms

Check the response. Look for the “avail” figure for `/cms` (shown as “yyyyyy” in the example above). Record that figure here:

/cms “avail” figure: _____

6. Partition the new disks. Follow the instructions in [“Partitioning disks with the format command” on page 76](#). Double check the partitioning before continuing.

7. Add the *olds* directories to your path:

```
# PATH=$PATH:/usr/opt/SUNWmd/sbin:/olds
# export PATH
```

8. Create new versions of the setup files by entering the following command:

```
# olds -mirrored -mk_files <newdisk1> <newdisk2>
```

<newdisk1> and <newdisk2> represent the device names of the disk drives added (remember that in a mirrored system, disk drives are always added in pairs). For example:

```
# olds -mirrored -mk_files c0t10d0 c0t11d0
```

9. Load the file */olds/md.tab.new* into an editor and check to make certain it names all of your disks, including the ones you just added. (For an example of the format of an *md.tab* file, see [“Examples of valid system files” on page 79.](#))

10. Save a copy of the existing *md.tab* file by entering the following command:

```
# cp /etc/opt/SUNWmd/md.tab /etc/opt/SUNWmd/md.tab.mirror
```

11. Replace the existing *md.tab* file with the one you created in Step 8:

```
# cp /olds/md.tab.new /etc/opt/SUNWmd/md.tab
```

12. Attach one of the newly installed disk drives to each of the d21 submirrors (d19 and d20):

```
# metattach d19 <newdisk1>
# metattach d20 <newdisk2>
```

where <newdisk1> and <newdisk2> represent the device names of the disk drives you have just added. For example:

```
# metattach d19 c0t10d0s1
# metattach d20 c0t11d0s1
```

13. Grow the */cms* file system by entering the following command:

```
# growfs -M /cms /dev/md/rdisk/d21
```

If the operation should fail, do the following:

- a. Restore the original *md.tab* file by entering the following command:

```
# cp /etc/opt/SUNWmd/md.tab.mirror /etc/opt/SUNWmd/md.tab
```

- b. Look up the error message in Chapter 4, “Troubleshooting.”

14. Enter the following command:

```
# metastat
```

Examine the output to make certain the disk drives you just added are part of the configuration.

15. Enter the following command:

```
# df -k /cms
```

The system responds by displaying file system information for */cms*. For example:

Filesystem	kbytes	used	avail	capacity	Mounted on
/dev/md/dsk/d21	xxxxxxx	9	yyyyyy	0%	/cms

Check the response. Look for the “avail” figure for */cms* (shown as “yyyyyy” in the example above). That figure should be roughly twice the figure recorded in Step 5.

Replacing a faulty disk

This section explains how to replace a faulty disk. The system does not need to be powered off or rebooted to perform this procedure.

Step 1: Identify the faulty disk

A disk problem is usually indicated in the `/var/adm/messages` file. The following lines, for example, indicate a disk problem:

```
Jun 12 16:27:08 leopard unix: WARNING:
Jun 12 16:27:08 leopard unix: Error for command
'read(10)' Error Level: R
Jun 12 16:27:09 leopard unix: retryable
Jun 12 16:27:09 leopard unix: Requested Block 0,
Error Block: 0
Jun 12 16:27:09 leopard unix: Sense Key: Media Error
Jun 12 16:27:09 leopard unix: Vendor 'SEAGATE':
Jun 12 16:27:09 leopard unix: ASC = 0x31 (medium
format corrupted), ASCQ
= 0x0, FRU = 0x9
Jun 12 16:27:09 leopard unix: WARNING:
/sbus@3,0/SUNW,fas@3,88000000/sd@2,0
```

Those lines, however, do not tell us which disk has the problem. One way of finding out is to enter an `ls -l /dev/dsk/c*` command and search the output for a device description matching that in the warning message. For example:

```
# ls -l /dev/dsk/c*
.
.
.
lrwxrwxrwx  1 root    root          50 Apr 24 15:21
/dev/dsk/c0t2d0s0 -> ../../devices/sbus@3,0/SUNW,fas
@3,88000000/sd@2,0:a
lrwxrwxrwx  1 root    root          50 Apr 24 15:21
/dev/dsk/c0t2d0s1 -> ../../devices/sbus@3,0/SUNW,fas
@3,88000000/sd@2,0:b
.
.
.
#
```

That the `"/iommu@0..."` information matches the same information in the warning message indicates that disk `c0t2d0` is the faulty disk.

Step 2: Identify the submirrors

1. Add the *Solstice DiskSuite* directories to your path:
- # export PATH=\$PATH:/usr/opt/SUNWmd/sbin:/olds
2. Enter a metastat command:
- # metastat
3. Scan the output of the metastat command to find the name of the faulty disk. The following lines, for example, indicate that the faulty c0t2d0 disk is in d20:
- # metastat
- .
.
d20: Submirror of d21
Size: 14960160 blocks
Stripe 0:
Device Start Block Dbase State Hot Spare
c0t2d0s1 0 No Okay
Stripe 1:
Device Start Block Dbase State Hot Spare
c0t4d0s3 0 No Okay
.
.

Be sure to check all the submirrors. If the disk is also listed under d11 or d12, it is a boot disk belonging to two different submirrors. Replacing a boot disk involves some special considerations; those considerations are noted in this procedure where they apply.

4. Record below the metadevices to which the disk belongs (for example, **d11** and **d19**):
- d_____ d_____

Step 3: Detach the submirrors

To detach a submirror means to put it off line so you can access its constituent disk drives. You detach a submirror with a `metadetach` command. The command is in the format:

```
metadetach <metamirror> <submirror>
```

where *<metamirror>* is the mirror and *<submirror>* is the submirror to be detached.

To continue our example, since `c0t2d0` is part of `d20`, which is a submirror of `d21`, the detach command would be as follows:

```
# metadetach d21 d20
d21: submirror d20 is detached
```

If the faulty disk is a boot disk, a `metastat` command also shows the disk as part of the `d11` or `d12` submirror. Be sure to detach both the submirrors of which it is a member. For example:

```
# metadetach d13 d12
d13: submirror d12 is detached
# metadetach d21 d20
d21: submirror d20 is detached
```

Step 4: Remove state database replicas (boot disks only)

When you replace a boot disk, you must remove the state database replicas from the old disk and recreate them on the new one. A boot disk is one that is a member of metadvice d11 or d12 (see [“Step 2: Identify the submirrors” on page 70](#)).

If you fail to remove old replicas, the software will have an inaccurate picture of the replicas existing on your system and will attempt to write to nonexistent database replica files. Match up the device description lines to determine which disk drive has problems.

In general, the procedure is a three-step process:

1. Issue a `metadb -i` command to find out which replicas to remove.
2. Issue a `metadb -d` command to do the actual deletion.
3. Issue another `metadb -i` command to verify the deletion.

If the faulty disk were `c0t1d0`, for example, the replica deletion might resemble the following series of commands and responses:

```
# metadb -i
  flags      first blk    block count
a m p   lu0    16         1034      /dev/dsk/c0t0d0s1
a      p   lu0   1050         1034      /dev/dsk/c0t0d0s1
a      p   lu0   2084         1034      /dev/dsk/c0t0d0s1
a      p   lu0   3118         1034      /dev/dsk/c0t0d0s1
a      p   lu0    16         1034      /dev/dsk/c1t4d0s1
a      p   lu0   1050         1034      /dev/dsk/c1t4d0s1
a      p   lu0   2084         1034      /dev/dsk/c1t4d0s1
# metadb -d c1t4d0s1
# metadb -i
  flags      first blk    block count
a m p   lu0    16         1034      /dev/dsk/c0t0d0s1
a      p   lu0   1050         1034      /dev/dsk/c0t0d0s1
a      p   lu0   2084         1034      /dev/dsk/c0t0d0s1
a      p   lu0   3118         1034      /dev/dsk/c0t0d0s1
```

Step 5: Replace the faulty disk drive

1. Remove the faulty disk drive.
2. Install the new disk drive into the same slot the faulty drive came out of. For installation instructions, see the documentation that came with the disk drive or the Lucent hardware installation manual.
3. Partition the new disk using the `format` command. See [“Partitioning disks with the format command” on page 76](#). **Make sure you partition the new disk drive!**

Step 6: Recreate the state database replicas

1. Recreate the primary boot disk replicas you removed earlier by entering the following command:

metadb -a -c4 -f /dev/rdisk/devname
2. Recreate the secondary boot disk replicas you removed earlier by entering the following command:

metadb -a -c3 -f /dev/rdisk/devname
3. Verify that the replicas were created:

metadb -i

The system lists the state database replicas currently on the system. The following, for example, shows three replicas on c0t4:

flags				first blk	block count	
a	m	p	luo	16	1034	/dev/dsk/c0t0d0s1
a		p	luo	1050	1034	/dev/dsk/c0t0d0s1
a		p	luo	2084	1034	/dev/dsk/c0t0d0s1
a		p	luo	3118	1034	/dev/dsk/c0t0d0s1
a		p	luo	16	1034	/dev/dsk/c1t4d0s1
a		p	luo	1050	1034	/dev/dsk/c1t4d0s1
a		p	luo	2084	1034	/dev/dsk/c1t4d0s1

In the above example, there are 4 replicas listed for the primary boot disk (c0t0d0s1), and 3 replicas listed on the secondary disk (c1t4d0s1)

 **WARNING:**

Do not try to reboot a system when there are fewer than two state database replicas.

Step 7: Reattach the submirrors

1. Enter a metainit command to reinitialize the submirror. For example:

```
# metainit d20
#
```

If you receive an “already set up” message, ignore it.

2. Reattach the submirror by issuing a metattach command. For example:

```
# metattach d21 d20
#
```

If the disk you are replacing is a boot disk, be sure to reattach both the submirrors of which it is a member. For example:

```
# metattach d13 d12
# metattach d21 d20
```

Step 8: Reboot the system (boot disks only)

You may delay this step until a more convenient time.

If the replaced disk is a boot disk, you will have to reboot the system at your convenience, typically during low busy hours. The reboot is necessary because the */etc/system* file was modified when the state database replicas were recreated.

State Database Replicas

Setting up replicas

1. Verify that the /etc/opt/SUNWmd/md.tab file is correct (see [“Examples of valid system files” on page 79](#)).
2. Enter the metadb command:

```
# metadb -a -c4 -f /dev/dsk/c0t0d0s1
# metadb -a -c3 -f /dev/dsk/c0t1d0s1
```

Replacing replicas

Sometimes replicas may have to be replaced. A replica can become corrupted, for example, and in that case needs to be removed and then recreated. The procedure to do that is as follows:

1. Remove all existing replicas with the following command:

```
# metadb -d -f /dev/dsk/c0t0d0s1
```



WARNING:

Do not reboot the system after removing replicas! With replicas gone, the system may not boot!

2. Recreate two replicas on the same partition:

```
# metadb -d -c4 -f /dev/dsk/c0t0d0s1
```

Partitioning disks with the *format* command

To use the format command to partition a disk, do the following steps:

1. Enter: `# format`
- A menu of disks displays. You are asked to select one.
2. Enter the menu number of the disk to be partitioned. A list of commands displays, followed by the `format>` prompt.
3. Enter: `format> partition`
- The `partition>` prompt displays.
4. Enter: `partition> print`
- The partition table displays and `partition>` returns.
5. Enter the number of the partition to be configured. The system begins a series of prompts for partitioning information. Answer according to the CMS version, disk type (boot or nonboot), and disk size you are partitioning.

Partitioning (Note: **Never** change the default size of the *overlap* slice!)

				CMS R3V8 4.2-GB Disk		CMS R3V8 9.1-GB Disk SCSI (E3000) and FC-AL (E 3500)	
Disk Type	Part	Tag	Flag	Size (cyl.)	Starting Cylinder	Size (cyl.)	Starting Cylinder
Boot	0	<i>boot disk 1: /</i> or root <i>boot disk 2: unassigned</i>	wm	1023	0	616	0
	1	unassigned	wm	7	1023	7	616
	2	overlap <i>or</i> backup	wm	3880	0	4924	0
	3	unassigned	wm	1879	1030	3716	623
	4	swap	wm	971	2909	585	4339
	5-7	unassigned	wm	0	0	0	0
Nonboot	0	unassigned	wm	2	0	2	0
	1	unassigned	wm	3878	2	4922	2
	2	overlap <i>or</i> backup	wm	3880	0	4924	0
	3-7	unassigned	wm	0	0	0	0

6. Enter: `partition> print`

The partition table displays and the `partition>` prompt returns. Compare the display to the *Partitioning* table to make certain you have repartitioned the disk properly. If not, repeat step 5.

7. Enter: `partition> label`

The system responds `Ready to label disk, continue?`.

8. Enter: `y`

9. The `partition>` prompt returns.

10. Enter: `partition> q`

The `format>` prompt displays.

11. To select another disk to partition, enter: `format> disk` and return to item 2.

To quit, enter: `format> q` and continue with the next step.

12. When the system prompt returns, verify the partitioning.

You can verify the exact partitioning of the new disk with an `olds -mirrored -check_disks` command. For example:

```
# olds -mirrored -check_disks c0t2d0
scsi=c1
number of external scsi controllers with disks
is = 1
number of disks is = 6
valid disks are c0t2d0
disk:c0t2d0 is partitioned ok
#
```

How to tell whether your mirrored system is standard

To tell whether mirroring was installed as a “standard” or “nonstandard” configuration, you can do one of two things. Either:

- a. Enter the command: `# metastat d11 d12`
- b. Edit the `/etc/opt/SUNWmd/md.tab` file

Examine the command output or the contents of the file. You are looking for the d11 and d12 metadevices and their constituent disk drives.

Standard mirroring—mirroring that was installed at the factory or upgraded to a standard factory configuration—will have the structure indicated in the following table:.

Root Metadevice Device Names

Metadevice	Device Name
metadevice d11	c0t0d0
metadevice d12	c0t1d0 (E3000) c1t4d0 (E3500)

Examples of valid system files

Valid *vfstab* files

Valid <i>vfstab</i> file for a non mirrored R3V8 CMS system:						
#device #to mount #	device to fsck	mount point	FS type	fsck pass	mount at boot	mount options
/dev/dsk/c1d0s2	/dev/rdisk/c1d0s2	/usr	ufs	1	yes	-
fd	-	/dev/fd	fd	-	no	-
/proc	-	/proc	proc	-	no	-
/dev/dsk/c0t0d0s4	-	-	swap	-	no	-
/dev/dsk/c0t0d0s0	/dev/rdisk/c0t0d0s0	/	ufs	1	no	-
swap	-	/tmp	tmpfs	-	yes	-
/dev/md/dsk/d19	/dev/md/rdisk/d19	/cms	ufs	2	yes	-

Valid <i>vstab</i> file for an R3V8 CMS system with mirrored disks:						
#device #to mount #	device to fsck	mount point	FS type	fsck pass	mount at boot	mount options
fd	-	/dev/fd	fd	-	no	-
/proc	-	/proc	proc	-	no	-
/dev/md/dsk/d17	-	-	swap	-	no	-
/dev/md/dsk/d13	/dev/md/rdisk/d13	/	ufs	1	no	-
/dev/md/dsk/d21	/dev/md/rdisk/d21	/cms	ufs	2	yes	-

Valid *md.tab* files

Valid *md.tab* file for a 3 disk non mirrored R3V8 CMS system:

```
#state database replicas
mddb00 c0t0d0s1
mddb01 c0t1d0s0

#/cms
d19 3 1 /dev/rdisk/c0t0d0s3 1 /dev/rdisk/c0t1d0s1 1 /dev/rdisk/c1t0d0s1
d21 -m d19
```

Valid *md.tab* file for a 4 disk mirrored R3V8 CMS system:

```
#state database replicas
mddb00 c0t0d0s1
mddb01 c1t4d0s1
mddb02 c0t1d0s0
mddb03 c1t5d0s0

#metaroot
d11 1 1 c0t0d0s0
d12 1 1 c1t4d0s0
d13 -m d11

#/cms
d19 2 1 /dev/rdisk/c0t0d0s3 1 /dev/rdisk/c0t1d0s1
d20 2 1 /dev/rdisk/c1t4d0s3 1 /dev/rdisk/c1t5d0s1
d21 -m d19
```

Troubleshooting

This chapter contains procedures for troubleshooting problems you may experience while installing or maintaining systems that use disk mirroring.

***Solstice DiskSuite* problems**

The *root* filesystem runs out of space during a system install

If you run out of space on the */*(*root*) file system—while trying to install *Solaris* patches, for example—you probably installed your */cms* swap file on *root*. Do the following to recover:

1. Unmount */cms*: `umount /cms`
2. Enter the command: `ls -l /cms`

If the output of that command lists a swap file, do the following:

- a. Enter: `mount /cms`
Remounts /cms.
- b. Enter: `df -k`
Verifies that space is available.

Excessively long resync

When using `metadetach` and `metattach` to detach and then reattach a submirror, it may take a long time—sometimes hours—for resync to complete. You cannot do any maintenance (adding or replacing disks and so on) while the resync is in progress. That behavior is normal. Try again later.

The system fails to recognize all disk drives

If the system should fail to recognize all disk drives, do the following:

- 1. Reboot the system by entering the following command:

```
# /usr/sbin/shutdown -y -g0 -i0
```

The system reboots and displays the ok prompt:

- 2. Turn off the system unit.
- 3. Turn off all external devices starting with the system monitor and working toward the device farthest from the system unit.
- 4. Check all disk drive connections to make sure they are secure.
- 5. Turn on the power to the system components in the opposite order in which you powered them off. Power on the SCSI devices first, beginning with the device at the end of the chain and working your way toward the system unit. Then power on other devices, again working your way toward the system unit.

Finally, power on the system unit itself and the system monitor. When you power on the system unit, the system begins to boot. Interrupt the boot by pressing Stop-A. The system displays the ok prompt.

- 6. To verify that the system sees all the disk devices, including any new drives, enter one of the following commands:

```
ok probe-scsi-all           ← Enterprise 3000
ok probe-fcal-all           ← Enterprise 3500
```

If you receive a “This command may hang the system....” message, respond with the following steps:

- a. Enter: n (to prevent the probe from continuing).
- b. Enter: setenv auto-boot? false (to prevent a reboot when you do the next command)
- c. Enter: reset -all (it may take a minute to complete).

Now repeat the probe-scsi-all or probe-fcal-all.

The program responds to the probe as shown below. The specific response, of course, differs depending on the platform.

```
/sbus@3,0/SUNW,fas@3,8800000
Target 0
  Unit 0   Disk   SEAGATE ST34371W
SUN4.2G74629737K6748
                                Copyright(c) 1997 Seagate
                                All rights reserved
.
.
.
Target 6
  Unit 0   Disk   Removable Read Only device
TOSHIBA XMS5701TASUN12XCD099704/09/97
ok
```

7. **IMPORTANT! If you had to do a reset -all command**, reset *autoboot?* to *true* by entering the following command:

```
ok setenv auto-boot? true
```



WARNING:

Failure to reset *autoboot?* before the next reboot will cause the reboot to stop at the boot prompt instead of proceeding through the normal boot-up.

8. When you have verified that the system is recognizing all its disk drives, reboot the system with a `boot -r` command and log in as the root user. For example:

```
ok> boot -r
.
.
.
console login: root
Password:
```

Fixing “needs maintenance” messages

On rare occasions mirrors may get severely out of sync. An extreme example is when a disk is removed from a powered system without first being detached with a `metadetach` command. The situation is indicated by the appearance of a *Needs maintenance* message in a *State* line of a `metastat` response. For example:

```
# metastat
d13: Mirror
    Submirror 0: d11
        State: Needs maintenance
    Submirror 1: d12
        State:
    Pass: 1
    Read option: roundrobin (default)
    Write option: parallel (default)
    Size: 2209680 blocks

d11: Submirror of d13
    State: Needs maintenance
    Invoke: metareplace d13 c0t0d0s0
    Size: 2209680 blocks
    Stripe 0:
        Device      Start Block Dbase State Hot Spare
        c0t0d0s0      0      No    Okay
        .      .      .
```

```
#
```

In such situations the `metareplace` command can force the recovery of the failed slice. The *Invoke* line of the `metastat` response leads the way. Given the response above, for example, you would enter the following:

```
# metareplace -e d13 c0t0d0s0
# metastat
d13: Mirror
    Submirror 0: d11
        State: Resyncing
    Submirror 1: d12
        State: Needs maintenance
    .      .      .
```

```
#
```

Boot problems

Generally speaking, most boot problems are the result of one of two problems: either you are trying to boot from the wrong device or the primary boot device is not available. This section suggests remedies for those situations.

Trying to boot from the wrong device

If you are trying to boot from the wrong device, you must change the boot device. Here is a quick way to do that:

1. Reboot: `# /usr/sbin/shutdown -y -g0 -i0`
2. From the ok prompt enter: `ok> printenv boot-device`

The system responds: `boot-device <disk>`

where *<device>* is either the word *disk*, indicating the primary boot device, or the word *bootdevice2*, indicating the alternate boot device. (This may be followed by additional device names.)

To change the boot device, enter the following command:

```
ok setenv boot-device bootdevice2
```

- or -

```
ok setenv boot-device disk
```

You should have set up *bootdevice2* as the alternate boot device when you installed disk mirroring.

Primary boot device is not available

If the primary boot device is not available, you must change to the secondary boot device.

Changing to the secondary boot device is a matter of entering, from the ok prompt, the following command:

```
ok> setenv boot-device bootdevice2
```

You should have set up *bootdevice2* as the alternate boot device when you installed mirroring.

Excessive reboot time

Problem: The system takes a long time to come back up after a reboot and seems to stop at the *configuring multicast* point.

Possible workaround: The *Enterprise 3000*, by default, tries to configure itself as a router/gateway. To avoid that problem, execute the following command before rebooting:

```
touch /etc/notrouter
```

That command tells the startup scripts that this system should not be configured as a router.

Log files

A disk-mirrored system contains the following system log files:

File Name	Description
<code>/var/adm/messages</code>	Contains all console error messages, including hardware messages. (For an example of how to use this log to detect a bad disk drive, see <i>Replacing Faulty Disks</i> in Chapter 3.)
<code>/olds/elog</code>	Contains <i>olds</i> error messages. Most messages are logged when setting up the system or when making changes to the disk setup.

If a log seems to indicate a problem, it may take an experienced troubleshooter to determine the nature of the problem.

The logs are checked each night by the `chkDisks` crontab, which is able to determine whether any potential or actual drive problems have been logged. The results of the determination are logged to `/olds/elog` and mailed to the *root* user.

If `chkDisks` is not running each night, verify that it is available in the `/olds` directory by entering the following command:

```
# ls /olds/chkDisks
```

If the file is not found, load the CMS CD into the CD-ROM drive and enter the following commands:

```
# cp /cms/ronly/olds_install/chkDisks /olds
# chmod +x /olds/chkDisks
```

To test `chkDisks`, simply execute it by entering the following command:

```
# /olds/chkDisks
```

Error messages

The table that follows lists and explains error messages you can receive related to *Solstice DiskSuite* and disk mirroring. The messages are presented in alphabetical order, ignoring special characters and replacement strings. (For example, you would find the message "*filename* restored" in alphabetical order by the word *restored*.)

Error Messages

Message:	All data on cms filesystems will be removed and filesystems will be unmirrored, are you sure?
Explanation:	Standard warning message for olds -mirrored -cleanup command.
Response:	Always use an olds -mirrored -cleanup with caution, keeping in mind that a backup will be needed to recover any CMS data in the /cms filesystem. Use the command only to set up your system from scratch.

Message:	Attempt an operation on a submirror that has erred components.
Explanation:	Component needs maintenance before the operation can be carried out.
Response:	1. Execute a metastat command. 2. Examine the output for the <i>Invoke</i> line and execute the command named in that line. For example: <pre># metareplace -e d13 c0t0d0s0</pre> 3. Reboot system.

Message:	CDROM must not be present, remove CD
Explanation:	In order for the command to execute, the CDROM drive must be empty.
Response:	1. Enter the command: <code>eject cdrom</code> 2. Remove the CD from the drive tray. 3. Close the drive tray.

Error Messages (Continued)

Message:	component is attached metattach: <i>hostname</i> : d21: growing of metadevice delayed
Explanation:	Submirrors are out of sync.
Response:	Enter the command: <code>metasync d21</code> This command effectively resyncs submirrors.

Message:	device: <i>devicename</i> cannot be setup or does not exist
Explanation:	The device named is a CD-ROM drive or is not available and will not be used in the <i>Solstice DiskSuite</i> setup.
Response:	Verify that the device should not be used.

Message:	device: <i>devicename</i> will not be used
Explanation:	The device named is a CD-ROM drive or is not available and will not be used in the <i>Solstice DiskSuite</i> setup.
Response:	If the device is c0t6d0, it is a CD-ROM drive and no further action is required. If the device is a disk, verify the sanity of the message: check device availability, for example, or use a <code>format</code> command to view partition sizes, etc.

Message:	disk <i>diskname</i> already setup, exiting...
Explanation:	You are trying to attach a disk that is already attached or is part of the base configuration. No response is necessary.

Message:	disk: <i>diskname</i> slice 1 is not partitioned correctly
Explanation:	Disk is partitioned incorrectly, or the <i>olds</i> script was run without the <i>-mirrored</i> option.
Response:	Try rerunning the command with the <i>-mirrored</i> option. If that doesn't work, repartition the disk. If <i>diskname</i> is the primary boot disk, you must reload <i>Solaris</i> , repartitioning the disk correctly in the process. If <i>diskname</i> is not the primary boot disk, you can repartition it with the <i>Solaris</i> <code>format</code> command.

Error Messages (Continued)

Message:	/etc/system has been updated since the last reboot. CMS cannot run without an up-to-date /etc/system file.
Explanation:	/etc/system can change when a particular Solaris patch is applied to the system or when state database replicas are removed and re-added during a boot disk replacement.
Response:	Reboot the system.

Message:	Failed activating new replicas, replicas already active, or a second disk does not exist
Explanation:	The system does not have all the replicas it needs to be functional (it needs at least three), or the replicas you attempted to create already exist.
Response:	Use <code>metadb -i</code> to verify that at least three replicas are active. If not, add replicas (see Chapter 3).

Message:	Failed: disk <code>cxtyd0</code> or <code>cx'ty'd0</code> already setup, exiting ...
Explanation:	You are trying to set up a disk that is already a part of the configuration. You probably mistyped the name.
Response:	Set up the disk with the appropriate device name.

Message:	Illegal Request, cannot setup
Response:	Check your hardware configuration and your usage of the <code>olds</code> command.

Message:	in order to attach disk, /cms must already be mounted, exiting...
Response:	Mount cms with this command: <code>mount /cms</code>

Message:	<code>ksh: command: not found</code>
Explanation:	Command does not exist or is not in the execution path.
Response:	Use the entire path—as <i>pathname/command</i> —or add it to the execution path with the following command: <pre># export PATH=\$PATH:pathname</pre> The path of the <code>metastat</code> , <code>metadb</code> , <code>metattach</code> , and <code>metadetach</code> commands is <code>/usr/opt/SUNWmd/sbin</code> ; the path of the <code>olds</code> command is <code>/olds</code>

Error Messages (Continued)

Message:	metadb: <i>hostname</i> : cxytd0sz has a metadatabase replica
Explanation:	The slice named in the message already has a replica.
Response:	If you really need to add another replica, choose a different slice.

Message:	metadetach: <i>systemname</i> : d21: resync in progress
Explanation:	You cannot detach a metadatabase while a resync is in progress.
Response:	Try again later. To find out how far along the resync has progressed, enter a metastat command.

Message:	must have at least 2 databases
Response:	Add a state database replica and retry the command.

Message:	No disks found! Internal error.
Response:	Add the state database replicas and retry the command. If the message persists, telephone the National Customer Care Center (1-800-242-2121) or contact your Lucent distributor or customer representative.

Message:	No more space on device
Explanation:	Occurs during a restore cpio command: cpio -icmudf -C10240 -I input-device -M "Enter tape.. " ../cms*.. The "/cms" entry has probably been mistyped.
Response:	Remove the /cms directory and all files and try again.

Message:	<i>command</i> : not found
Explanation:	Command does not exist or is not in the execution path.
Response:	Use the entire path—as <i>pathname/command</i> —or add it to the execution path with the following command: # export PATH=\$PATH: <i>pathname</i> The path of the metastat, metadb, metattach, and metadetach commands is <i>/usr/opt/SUNWmd/sbin</i> ; the path of the olds command is <i>/olds</i>

Error Messages (Continued)

Message:	panic: vfs_mountroot: cannot mount root rebooting... resetting....
Explanation:	The system is misconfigured.
Response:	There are two possible causes: 1. <i>You attempted to boot from a nonbootable device.</i> Check the device name and try again. The alternate boot device may be set up incorrectly. 2. <i>You failed to reboot after mirroring root</i> (olds -mirrored -setroot), leaving the system unstable. In that case, do the following: a. Boot from the CD. b. Mount /dev/dsk/c0t0d0s0 /a (or whatever partition is your root partition). c. Remove <i>Solstice DiskSuite</i> entries from /a/etc/system (i.e. everything between tags "mddb ...") d. Remove /a/etc/opt/SUNWmd/md* files, i.e.: rm /a/etc/opt/SUNWmd/md* e. Replace /dev/md/rdisk/d21 entries with /dev/rdisk/c0t3d0s0 (or whatever your boot device name is). f. Boot: /usr/sbin/shutdown -y -g0 -i6 g. If the system still refuses to reboot, reload <i>Solaris</i>.
Message:	<i>filename</i> restored from <i>filebackup</i>
Explanation:	The action failed, and the md.tab file was restored from the previous version. Consequently, the configuration files reflect the previous system setup.
Response:	Determine the cause of the problem and try again.
Message:	resync in progress
Response:	The command entered cannot be executed because the mirrors are currently being resynched. Try again later. You can tell whether a resync is in progress by entering a metastat command.

Error Messages (Continued)

Message:	stale databases
Response:	The state database contains old information. Recreate (see “Recreating the Replicas and Reattaching the Disk” in Chapter 3).

Message:	state database information not found, DiskSuite must not be installed
Response:	Install and set up <i>Solstice DiskSuite</i> before doing this task.

Message:	Submirror too small to attach
Explanation:	The submirrors d11 and d12 (or d19 and d20) do not match. You have forgotten to label one or more disks, or you have partitioned one or more disks incorrectly.
Response:	1. Determine the specific cause of the message by examining, via the <i>format</i> and <i>partition</i> commands, the disks involved. 2. If the message was a response to a metattach d13 d12 command, enter the following commands: <pre># metaclear d12 # metadb -d -f /dev/dsk/c0t<x>d0s1</pre> (where <x> is the second boot disk) 3. If the message was a response to a metattach d21 d20 command, enter the following command: <pre># metaclear d20</pre> 4. Label or repartition the disks in error, as needed. 5. Restart the mirror upgrade process. Note that you need not recreate the md.tab entries the second time through the procedure.

Message:	syntax error
Response:	Check the syntax and usage of the command. Reenter the command, correcting syntax errors you have made.

Error Messages (Continued)

Message:	The disk is not recognized by cms.
Explanation:	The disk is either not a standard Sun disk drive, or is larger than 4.2 gigabytes.
Response:	Replace the disk with a standard Sun disk drive of not more than 4.2 gigabytes capacity.

Message:	The DiskSuite Software must be installed to run this script.
Response:	Install <i>Solstice DiskSuite</i> software before attempting to run this script.

Message:	The file <i>filename</i> could not be restored.
Explanation:	The previous action failed, and the md.tab or vfstab file could not be copied back. The existing files may not accurately reflect the system environment.
Response:	Check the file and repair it if necessary. See “How to Set Up State Database Replicas” in Chapter 3.

Message:	The /cms filesystem needs to be mounted
Explanation:	/cms must be mounted for the command to work.
Response:	Mount /cms with the command: mount /cms

Error Messages (Continued)

Message:	This command may hang the system if a Stop-A or halt command has been executed. Please type reset-all to reset the system before executing this command. Do you wish to continue?
Response:	1. Enter: N (to prevent the probe from continuing). 2. Enter: setenv auto-boot? false (to keep the system from rebooting) 3. Enter: reset-all The reset may take a minute to complete. Once it does, you may do the <i>probe-scsi</i> or <i>probe-scsi-all</i> and perform any other boot prom level diagnostics. 4. IMPORTANT: Before you reboot again, enter: <pre>setenv auto-boot? true</pre> Failure to do so will cause the reboot to stop at the boot prompt.

Message:	touch: /cms/db/unix_start cannot create
Explanation:	A CMSADM backup was done when CMS was still running. An attempt is made to restart cms, but cms files are not yet available.
Response:	No response required. The message will disappear after you have restored and migrated CMS.

Message:	trying to add two new disks, but a mirrored setup needs to be in place.
Explanation:	You have tried to do an olds -mirrored -setup on an unmirrored system.
Response:	If you want a mirrored system, set it up as one. See the appropriate section in Chapter 2 for instructions on how to set up a mirrored system.

Message:	Unbalanced configurations not supported.
Explanation:	You tried to set up a system with an odd number of disk drives, or you tried to add disks without having rebooted via a "boot -r" command.
Response:	If necessary, add another disk drive to make the drive count even. Then reboot the system with "boot -r".

Error Messages (Continued)

Message:	Usage: olds -mirrored -check_disks -mirrored -mk_files -metadbs -cleanup -mirrored -setup [disk name]
Explanation:	You have mistyped a command.
Response:	Retry the command. See Chapter 3 for information about <i>olds</i> command syntax.

Message:	Warning: inode blocks/cyl group (230 >= data blocks (135) in lost cylinder group. This implies 2160 sector(s) cannot be allocated.
Explanation:	Some sectors will not be used by the filesystem. This is just a warning; the filesystem should be fine.

Message:	Warning: script does not support this configuration, configure manually
Explanation:	The olds script is supported only with the hardware described in Chapter 1.
Response:	Verify that you have appropriate hardware. If you need extra disks for purposes other than mirroring CMS, you may attach them after disk mirroring has been installed and has been verified to be running correctly.

Message:	x25netd: failed to open driver /dev/hihx (Bad file number [n])
Response:	Enter the following command: /cms/toolsbin/lnSBusdev

Message:	You must be root in order to run this command.
Explanation:	You are not logged in as the <i>root</i> user, but you must be in order to run the command.
Response:	Log in as <i>root</i> and retry the command.

Message:	You need to have at least one disk setup, before attaching one.
Response:	Install and configure <i>Solstice DiskSuite</i> before attempting this task.

Index

Symbols

/cms mirror 5

A

alternate boot device
 creating 26, 35
 testing 27
 using 39

B

backing up
 CMSADM backup 41–42
 CMSADM restore 43–44
 boot device
 creating alternate 26, 35
 testing alternate 27
 using the alternate 39
 boot devices 63
 boot disk
 device designations 28
 boot disk designations 78
 boot problems 85

C

capacity, system 12
 chkDisks 36, 40, 87
 chunks 2
 CMSADM backup 41–42
 CMSADM restore 43–44
 concatenated metadvice 2
 capacity 3
 cron file 27, 36

D

d11 metadvice 6
 d12 metadvice 6
 d17 metadvice 6
 d19 metadvice 5
 d20 metadvice 5
 d21 metadvice 5
 disk devices 28
 disk drives
 adding 63–65
 recognition problem 82
 replacing 69–74
 restoring 49
 disk mirroring
 definition 1
 factory installation 28–37
 upgrade procedure 11–27

disk partitioning 56, 76
 disk partitions 1
disks.type file 30
DiskSuite
 see *Solstice DiskSuite*
 drive slots 11

E

error messages 88–96

F

factory installation procedures 28–37
 faulty disk
 identifying 69
 replacing 69–74

H

hardware, required 8

I

increasing system capacity 12
 installation procedures, factory 28–37
 installing mirrored systems 11

K

Korn shell 30

L

layout, mirror
 Enterprise 3000 7
 Enterprise 3500 8
 log files 87

M

Mammoth 8mm tape drive 43
md.tab file 18
 metadvice
 concatenated 2
 definition 1
 mirror disks
 partitioning 16
 mirrors
 and metadvice 3
 definition 1
 identifying 12
 installing 11
 layout 7
 maintaining 39
 required hardware 8
 required software 9

N

needs maintenance messages 84

P

partitioning
 procedure 16
partitioning disks 76
partitioning, disk 56
printing *vfstab*. 41

R

replacing a faulty disk. 69–74
replicas
 replacing 75
 setup 75–76
restoring a CMSADM backup. 43–44
resync problem 81
root mirror. 6

S

slices 1
software, required. 9
Solstice DiskSuite. 1
 installation problems 81
 setup scripts 30
st.conf file 43
standard configuration 78
state database
 definition 3
 state database replicas 3
state database replicas
 replacing 75
 setup 75–76
submirrors 3

T

tape drives
 st.conf file. 43
three-way mirror 3
two-way mirror. 3

U

upgrading
 and increasing system capacity 12
 to disk mirroring 11–27

V

vfstab file 23, 24
 printing 41

How Are We Doing?

Document Title: *CentreVu*[®] Call Management System Release 3 Version 8
Disk-Mirrored Systems

Document No.: 585-210-940 Issue 1 Date: December 1999

Lucent Technologies welcomes your feedback on this document. Your comments are of great value in helping us to improve our documentation.

1. Please rate the effectiveness of this document in the following areas:

	Excellent	Good	Fair	Poor	Not Applicable
Ease of Use					////////////////////
Clarity					////////////////////
Completeness					////////////////////
Accuracy					////////////////////
Organization					////////////////////
Appearance					////////////////////
Examples					////////////////////
Illustration					
Overall Satisfaction					////////////////////

2. Please check the ways you feel we could improve this document:

- ☐ Improve the overview/introduction
- ☐ Improve the table of contents
- ☐ Improve the organization
- ☐ Include more figures
- ☐ Add more examples
- ☐ Add more detail
- ☐ Make it more concise/brief
- ☐ Add more step-by-step procedures/tutorials
- ☐ Add more troubleshooting information
- ☐ Make it less technical
- ☐ Add more/better quick reference aids
- ☐ Improve the index

Please provide details for the suggested improvement. _____

3. What did you like most about this document?

4. Feel free to write any comments below or on an attached sheet.

If we may contact you concerning your comments, please complete the following:

Name: _____ Telephone Number: (____) _____
Company/Organization: _____ Date: _____

When you have completed this form, please fax to +1-303-538-1741.

