

Lucent Technologies
Bell Labs Innovations



***CentreVu*® Call Management System**
Disk-Mirrored Systems

585-215-841
Comcode 108-056-755
Issue 1
May, 1998

Copyright © 1997, 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 Innovative Software, Inc.

Trademarks

Sun, *Solstice*, *SPARCserver*, *Enterprise*, *SPARCstorage*, *MultiPack*, and *DiskSuite* are trademarks of Sun Microsystems, Inc.

Federal Communications Commission Statement

Part 15: Class A Statement. This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his own expense.

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. 585-215-841
Comcode 108-056-755
Issue 1, November 1997

For additional documents, refer to the section entitled "Related Documents" in "About This Book."

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, please 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 Lucent Technologies Network Systems Customer Training and Information Products 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	4
Mirror Layout	5
Required Software	6
Additional References	6

Chapter 2: Installing Mirrored Systems

Upgrading an <i>Enterprise 3000</i> System to Disk Mirroring	7
Disk Drive Slots	7
Increasing System Capacity	7
Step 1: Identify Mirror Devices	8
Step 2: Install the Mirror Disks	9
Step 3: Partition the Mirror Disks	9
Step 4: Create the State Database Replicas	11
Step 5: Create md.tab Table Entries	11
Step 6: Set up the Mirror for root	12
Step 7: Create vfstab Entries	12
Step 8: Set up the Mirror for /cms	13
Step 9: Create an Alternate Boot Device	13
Factory Installation Procedures for Disk Mirroring	15
Step 1: Install the Hardware	15
Step 2: Install the Operating System	15
Step 3: Install Application Software	16
Step 4: Check The disk.type File	17
Step 5: Set Up the Disk Subsystem for DiskSuite	17
Step 6: Run the Setup Scripts	18
Step 7: Verify Disk Space	19
Step 8: Install the CMS Software	20
Step 9: Create an Alternate Boot Device	21

Chapter 3: Maintaining Mirrored Systems

Using the /olds Script	23
Example olds Commands.	24
Using the Alternate Boot Device	25
Maintaining the chkDisks crontab	26
Activating chkDisks	26
Verifying chkDisks	26
Changing the Scheduled Run Time	26
Canceling chkDisks	26
Doing a CMSADM Backup on a Mirrored System	27
Step 1: Print the vfstab File.	27
Step 2: Run the Backup.	27
Step 3: Label and Store the Tapes.	28
Restoring from a CMSADM Backup	29
Recreating the /cms File System	30
Restoring an Entire System from a CMSADM Backup	34
Restoring Specific Files from a CMSADM Backup	35
Adding New Disks to an Already-mirrored System	36
Adding New Disk Drives Manually	39
Replacing a Faulty Disk	42
Step 1: Identify the Faulty Disk.	42
Step 2: Identify the Submirrors.	43
Step 3: Detach the Submirrors.	43
Step 4: Remove State Database Replicas (Boot Disks Only).	44
Step 5: Install the New Disk Drive	45
Step 6: Recreate the State Database Replicas (Boot Disks Only)	45
Step 7: Reattach the Submirrors.	46
Step 8: Reboot the System (Boot Disks Only)	46
State Database Replicas	47
Setting Up Replicas	47
Replacing Replicas	47
Adding Replicas	47
Partitioning Disks Via the format Command	49
How to Tell Whether Mirroring was a Field Upgrade	51

Chapter 4: Troubleshooting

<i>Solstice DiskSuite</i> Problems	53
The root Filesystem Runs Out of Space During a System Install	53
Excessively Long Resync.	53
The System Fails to Recognize All Disk Drives	54
Fixing “Needs Maintenance” Messages	56
Boot Problems	57
Trying to Boot From the Wrong Device	57
Primary Boot Device is Not Available	57
Excessive Reboot Time.	57
Log Files	58
Problem Notification.	58
Error Messages	60

Index

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. In order to use disk mirroring, you must have a Sun Enterprise 3000 platform running CMS load r3v5uc.d or later.

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.

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

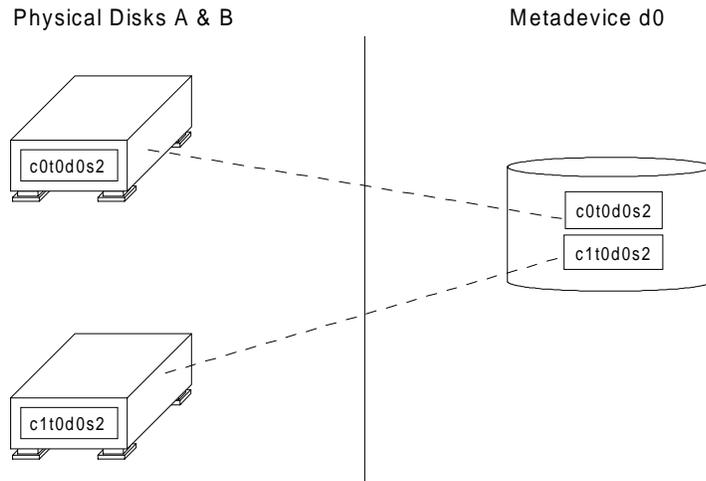
DiskSuite metadevices are built from **slices** (disk partitions). A system controlled by *Solstice DiskSuite* may contain any number of metadevices, each of which may comprise any combination of disk partitions.

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

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

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

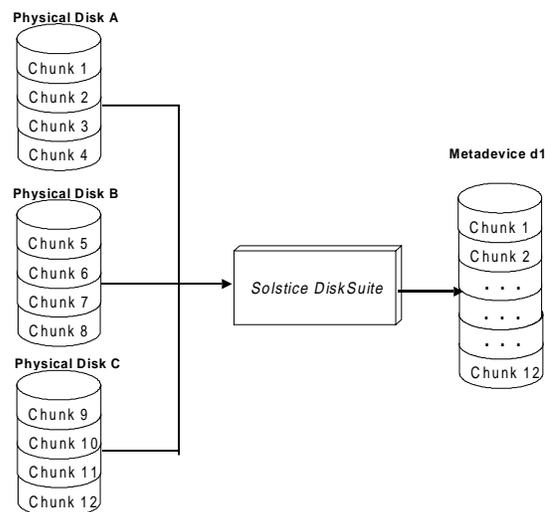
Disk Mirroring Defined



Concatenated Metadevices

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

Consider, for example, a concatenated metadevice 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 metadevice is the combined capacities of the three drives. If each drive is 4 gigabytes, for example, the metadevice 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 *DiskSuite* configuration.

The state database is actually a collection 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.

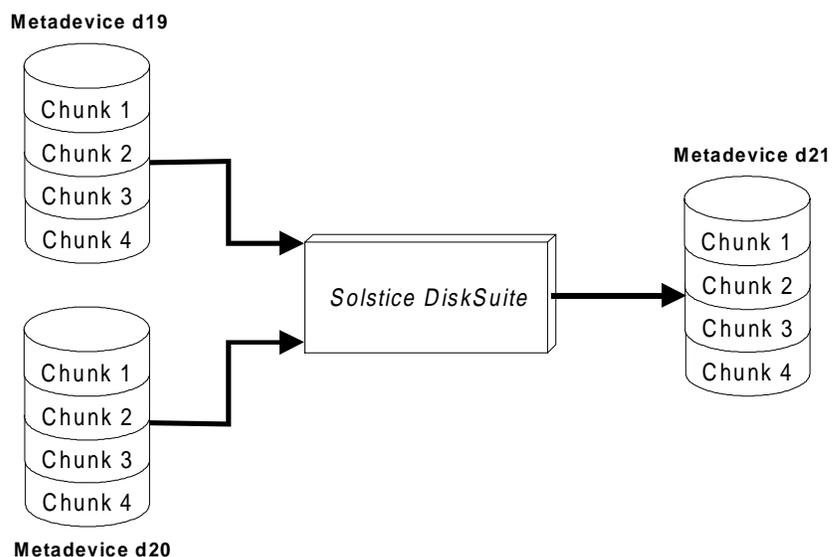
DiskSuite cannot operate until you have created the state database and its replicas: the software must have an operating state database.

Mirrors

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

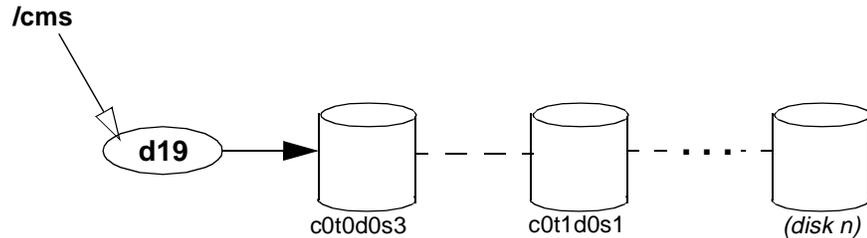
Mirroring provides redundant copies of data. To a software application, a mirror looks just like a physical disk. The mirror accepts I/O requests directed at it, and converts them into I/O requests for the submirrors. The submirrors in turn—being metadevices themselves—convert I/O requests from the mirror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

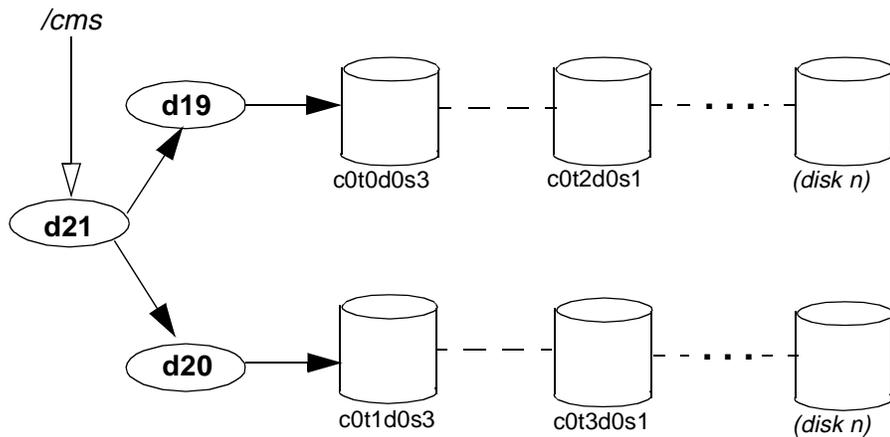


How CMS Implements Disk Mirroring

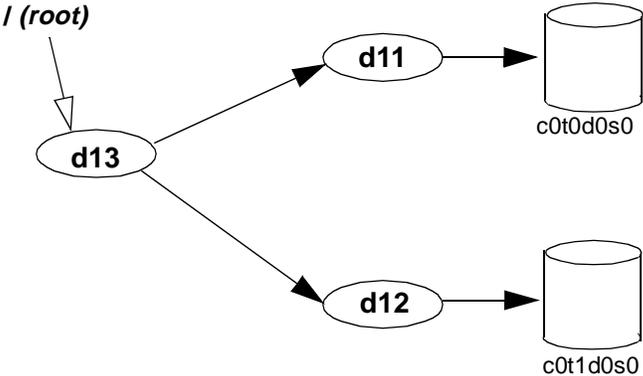
A nonmirrored, CMS system with *Solstice DiskSuite* uses *DiskSuite* to create a single metadvice, named d19, containing all the disk partitions used to store CMS data. For example:



To implement disk mirroring, a metadvice d20 is created as a duplicate of d19, and the two metadivices 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:



Mirror Layout

On a system 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.

No. Disks	/ (root) filesystem		/cms filesystem	
	submirror1	submirror2	submirror1	submirror2
2	c0t0d0s0	c0t1d0s0	c0t0d0s3	c0t1d0s3
4	c0t0d0s0	c0t1d0s0	c0t0d0s3 c0t2d0s1	c0t1d0s3 c0t3d0s1
6	c0t0d0s0	c0t1d0s0	c0t0d0s3 c0t2d0s1 c0t10d0s1	c0t1d0s3 c0t3d0s1 c0t11d0s1
8	c0t0d0s0	c0t1d0s0	c0t0d0s3 c0t2d0s1 c0t10d0s1 c0t12d0s1	c0t1d0s3 c0t3d0s1 c0t11d0s1 c0t13d0s1
10	c0t0d0s0	c0t1d0s0	c0t0d0s3 c0t2d0s1 c0t10d0s1 c0t12d0s1 c0t14d0s1	c0t1d0s3 c0t3d0s1 c0t11d0s1 c0t13d0s1 c0t15d0s1

How CMS Implements Disk Mirroring

Required Software

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

- *Solaris* 2.5.1 or later
- *Solstice DiskSuite* 4.0 or later
- *CentreVu*[®] CMS load r3v5ud.f or later

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 *DiskSuite* and disk maintenance and configuration includes:

Solstice DiskSuite 4.0 Administration Guide

Solstice DiskSuite Tool 4.0 User's Guide

^{*}*Sun* and *Sun Microsystems* are registered trademarks 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.

Upgrading an *Enterprise 3000* System to Disk Mirroring

This section tells how to upgrade *CentreVu*[®] Call Management System (CMS) on a *Sun*^{*} *Enterprise*[†] 3000 computer from a nonmirrored system to a mirrored system.

Disk Drive Slots

Before you begin, make certain you understand the drive slots in an *Enterprise 3000* system.

All disk drives in an *Enterprise 3000* computer are on controller 0. There are 10 slots, allowing up to five disks for each */cms* submirror. Each slot is labeled with a number 0 through 3 or 10 through 15; there are no slots numbered 4 through 9. 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; remaining slots are normally filled in target number order.

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 attempt to mirror the disks.

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

†*Enterprise* is a trademark of Sun Microsystems, Inc.

Step 1: Identify Mirror Devices

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

There should be as many new disks to be installed as there are old disks already on your system, in order to keep the same system capacity and make the two submirrors the same size.

2. Identify the device names for your mirrored system and record them in the *Device/Metadevice Names* table on the next page.

Here's how you identify the device names:

- a. Find the first empty slot in your system; it will hold the secondary boot disk. Record the device name of slice 0 as submirror 12, and the device name of slice 3 as disk 1 of submirror 20.

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	...	Submirror d20 Device Names
boot disk (2): <u>c0t2d0s0</u>		CMS disk 1: <u>c0t2d0s3</u>

- b. Record existing nonboot disks under submirror 19. Record 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

Submirror d11 Device Name	Submirror d12 Device Name	Submirror d19 Device Names	Submiror 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
		CMS disk 3: _____s1	CMS disk 3: _____s1
		CMS disk 4: _____s1	CMS Disk 4: _____s1
		CMS disk 5: _____s1	CMS disk 5: _____s1

Step 2: Install the Mirror Disks

Install the new disk drives that will be the mirror. Install them into the empty slots reserved in step 1. For installation instructions, see the documentation that came with the drives or the Lucent Technologies hardware installation manual for your version of CMS.

Step 3: Partition the Mirror Disks

Partition the newly installed disks using the `format` command. 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 repartition the disk. You probably will need to repartition.
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 for each partition on the disk.

Upgrading an Enterprise 3000 System to Disk Mirroring

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

Disk Type	Part	Tag	Flag	CMS R3V5 4.2-GB Disks		CMS R3V6 4.2-GB Disks	
				Starting Cylinder	Size	Starting Cylinder	Size
Boot	0	<i>boot disk 1: /</i> or root <i>boot disk 2: unassigned</i>	wm	0	761c	0	1023c
	1	unassigned	wm	761	71c	1023	7c
	2	overlap <i>or</i> backup	wm	0	3880c	0	3880c
	3	unassigned	wm	832	3048c	1030	2788c
	4	<i>R3V5: unassigned</i> <i>R3V6: swap</i>	wm	0	0	3818	62c
	5-7	unassigned	wm	0	0	0	0
Nonboot	0	unassigned	wm	0	2c	0	2c
	1	unassigned	wm	2	3878c	2	3878c
	2	overlap <i>or</i> backup	wm	0	3880c	0	3880c
	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 4: 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 by entering the following command:

```
# metadb -a -c2 -f /dev/rdisk/c0t<x>d0s1
```

(where **<x>** is the target number of boot disk 2)

3. Verify 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 m pc luo	16	1034	/dev/dsk/c0t2d0s1
a pc luo	16	1034	/dev/dsk/c0t2d0s1

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

Step 5: Create *md.tab* Table Entries

1. Make a copy of the *md.tab* file by entering the following command:

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

2. Edit */etc/opt/SUNWmd/md.tab* and insert entries for the root and d20 metadevices. The root metadevice 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 lines you add must be in this general format:

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

d<x> is the metadevice name (d12 or d20).

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

1 <device> indicates one disk with the device name *<device>*. You should have recorded the device names on page 9.

The pattern **1 (space) <device>** repeats until all disks in the metadevice have been named. The d12 entry must specify one disk drive; the d20 entry must specify the same number of drives 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:

Upgrading an Enterprise 3000 System to Disk Mirroring

```
----- /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  
-----
```

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

Step 6: Set up the Mirror for *root*

Set up the mirror for the root file system with the following steps:

1. Activate the first submirror (metadevice d11) and the mirror (d13) by entering the following commands:

```
# metainit -f d11  
# metainit -f d13
```

2. Activate the second submirror (d12) and attach it to d13::

```
# metainit -f d12  
# metattach d13 d12
```

3. Set up d13 as the root mirror:

```
# metaroot d13
```

4. Reboot the system:

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

Step 7: Create *vfstab* Entries

1. Make a copy of file `/etc/vfstab` to by entering the following command:

```
# cp /etc/vfstab /etc/vfstab.old
```

2. Edit `/etc/vfstab`. Comment out the d19 entry and add an entry for d21. For example:

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

3. Save the `/etc/vfstab` file and exit the editor.

Step 8: Set up the Mirror for /cms

1. Alter your path by entering the following commands:

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

2. Activate the new metadevice (d20) with a metainit command:

```
# metainit d20
```

3. Verify that the new submirror is active:

```
# metastat d20
```

4. Initialize the newly created mirror (d21):

```
# metainit -f d21
```

5. Attach the new metadevice (d20) to the newly created mirror (d21):

```
# metattach d21 d20
```

Step 9: Create an Alternate Boot Device

To set up the alternate boot device, do 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 metadevice on page 9).

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
```

2. Record the part of the line immediately following `/devices`. For the example above, you would record the part of the line beginning `/sbus@3,0` and ending `sd@2,0:a`.
3. Reboot the system with the command: `init 0`
The system displays the `ok` prompt.
4. Create a device alias for the new boot device by entering the following command (for example):

```
# nvalias bootdevice2 <device>
```

Upgrading an Enterprise 3000 System to Disk Mirroring

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@3,8800000/sd@2,0:a
```

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

```
# boot bootdevice2
```

6. Reboot the system as usual and verify that everything is normal.

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

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

8. 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
```

Factory Installation Procedures for Disk Mirroring

This section describes how to install an R3V5 CMS mirrored system on a new *Enterprise 3000* computer. These instructions may also be followed 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.

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 R3V5 Enterprise 3000 Software Installation (585-215-836)* or *CentreVu CMS R3V6 Software Installation and Maintenance document (585-215-866)* , **with the following exceptions:**

- In the *Customize Software* screen, select *SUNWmfrun* (the Motif Runtime Kit) ***in addition to*** any other software packages you are instructed to select.
- 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	Device Names
boot disk 1	c0t0d0
boot disk 2	c0t1d0 ^a
other disks	c0t2d0 c0t3d0 c0t10d0 c0t11d0 c0t12d0 c0t13d0 c0t14d0 c0t15d0

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

Factory Installation Procedures for Disk Mirroring

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

Disk Type	Part	Tag	Flag	CMS R3V5 4.2-GB Disks		CMS R3V6 4.2-GB Disks	
				Starting Cylinder	Size	Starting Cylinder	Size
Boot	0	<i>boot disk 1: /</i> <i>boot disk 2: unassigned</i>	wm	0	761c	0	1023c
	1	unassigned	wm	761	71c	1023	7c
	2	overlap <i>or</i> backup	wm	0	3880c	0	3880c
	3	unassigned	wm	832	3048c	1030	2788c
	4	<i>R3V5: unassigned</i> <i>R3V6: swap</i>	wm	0	0	3818	62c
	5-7	unassigned	wm	0	0	0	0
Nonboot	0	unassigned	wm	0	2c	0	2c
	1	unassigned	wm	2	3878c	2	3878c
	2	overlap <i>or</i> backup	wm	0	3880c	0	3880c
	3-7	unassigned	wm	0	0	0	0

Step 3: Install Application Software

Install application software as instructed in *CentreVu CMS R3V5 Enterprise 3000 Software Installation* (585-215-836) or *CentreVu CMS R3V6 Software Installation and Maintenance* (585-215-866). That means installing the following packages in order:

- Common Desktop Environment (CDE)
- *Sun Online Validation Test Suite (VTS) (R3V6 only)*
- *Solaris and hardware AnswerBooks (R3V5 only)*
- *SunLink** HSI/S drivers (if using HSI/S ports)
- NTS drivers (if using NTS ports)
- X.25
- *INFORMIX†* SQL (optional)
- *INFORMIX SE*
- *INFORMIX ILS (R3V6 only)*
- *Solstice DiskSuite*
- *Sun Solaris patches*

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

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

Step 4: Check The *disk.type* File

The `olds` commands, which you will be using to configure the *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 *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/ronly/olds_install/* /olds
# cd /olds
# chmod +x /olds/olds
```

Factory Installation Procedures for Disk Mirroring

4. Create system files for the *Solstice DiskSuite* software by entering the following commands::

```
# export PATH=$PATH:/usr/opt/SUNWmd/sbin:/olds
# olds -mirrored -check_disks
# olds -mirrored -mk_files
#
```

5. Check to 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. This example shows a total of four disk drives on the system, two in d19 and two in d20:

```
# vi /olds/md.tab.new

.
  (file contents display)
.

#/cms
d19 3 1 /dev/rdisk/c0t0d0s3 1 /dev/rdisk/c0t2d0s1
d20 3 1 /dev/rdisk/c0t1d0s1 1 /dev/rdisk/c0t3d0s1
d21 -m d19
```

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 54, and then return to item 4.

Step 6: Run the Setup Scripts

1. Create the state database replicas and mirror the *root* file system by entering the following commands::

```
# /olds/olds -mirrored -metadbs
# /olds/olds -mirrored -setroot
# init 6 ← IMPORTANT! Do Not Skip This Step!
```

2. Mirror */cms* by entering the following commands::

```
# /olds/olds -mirrored -setup
# mkdir /cms
# mount /cms
# /olds/olds -addswapfile /* R3V5 only */
# /olds/olds -addswapfile /cms/* R3V5 & R3V6 */
```

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/dsk/c0t0d0s0 772327 236870 458227   35%      /
/proc           0         0         0     0%      /proc
fd              0         0         0     0%      /dev/fd
/dev/md/dsk/d21 xxxxxxxx 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. To verify the disk space, go through the calculation in the table below.

Disk Space Verification

No. of 4-GB Disks	/cms Size (MB)
2	3104
4	7137
6	11170
8	15203
10	19236

Calculations:

(1a) Enter /cms line's 'kbytes' figure divided by 1000 (move decimal three places left):

(1a)

(1b) Enter "/cms Size" (above left) corresponding to the number of hard disks in your system:

(1b)

(1a) and (1b) should be *approximately* equal. The two will not correspond exactly, but they should be reasonably close—within about 10% of the /cms size figure.

3. Check the results of your verification.

A major discrepancy between the disk space you should have and the space recognized by your *DiskSuite* software usually indicates a connectivity problem.

You can check connectivity by entering a `metastat` command, and then checking the output of that command to see that all the system's disks have been recognized by the *DiskSuite* software.

For example:

Factory Installation Procedures for Disk Mirroring

```
# metastat
d13: Mirror
. . .
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 R3V5 Enterprise 3000 Software Installation* (585-215-836) or *CentreVu CMS R3V6 Software Installation and Maintenance* (585-215-866).

The steps to be completed are these:

- Install CMS R3V6 software
- Install CMS patches (if needed)
- Install CMS Supplemental Services software (*R3V6 only*)
- Set up CMS
- Install Open Database Connectivity (ODBC) software (optional)
- Set up the remote console
- Administer the NTS (if using NTS ports)
- Create an alternate boot device (see Step 9 below)
- Do a CMSADM backup.

Step 9: Create an Alternate Boot Device

The procedure for creating an alternate boot device is as follows:

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 with something like this:

```
lrwxrwxrwx  1 root    root          83 Jun 18 15:23
/dev/rdisk/c0t1d0s0 -> ../../devices/iommu@0, 10000000
/sbus@0,10001000/dma@1,81000/esp@1,80000/sd@1,0:b,raw
```

2. Record the device definition (the part of the line following */devices*). For the example above, you would record the part beginning */iommu@0,10000000* and ending *sd@1,0:b,raw*.
3. Reboot the system with the command: `# init 0`
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 definition>
```

Replace *<device definition>* with the string you recorded in Step 2.

5. Enter the command: `# boot bootdevice2`
(to verify that the system boots normally).
6. Reboot the system with the command: `# init 6`
Verify that everything is normal.
7. Enter the command: `# crontab -e`
The system starts the editor and loads the `cron` file.
8. 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
```

The factory installation procedure for disk mirroring is now complete.

Factory Installation Procedures for Disk Mirroring

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 `/olds` Script

The `/olds` script performs common functions associated with managing disk drives administered by the *DiskSuite* software. You can run an `olds` script by executing a command in the following format:

```
# olds [-mirrored] -param ...
```

where:

`olds`

is the script name. If the `/olds` directory is not in your path, precede the name with `/olds` (as: `/olds/olds`).

`-mirrored`

indicates the disks in the system are, or will be, mirrored. (This parameter alters the effect of other parameters to make them specifically applicable to mirroring.)

`-param ...`

represents a list of parameters separated by blanks. Each parameter must be preceded by a hyphen. Some parameters may be followed by the device names of the disks to which they are to be applied.

Legal parameters are as follows:

Using the /olds Script

<code>-check_disks [disk]</code>	Check the disk partitioning of the named disk. If no disk is named, all disks are checked.
<code>-mk_files [disk disk']</code>	Generate <i>md.tab.new</i> and <i>vfstab.new</i> files to include the named disks. If no disks are named, all disks are included.
<code>-metadbs</code>	Activate state database replicas.
<code>-setup [disk disk']</code>	Activate the <i>/cms</i> metadvice. If the two disk device names are included, the first is added to d19 and the second is added to d20
<code>-addswapfile [/cms]</code>	Add a swap file. If <i>/cms</i> is included, adds a swap file to <i>/cms</i> . If <i>/cms</i> is omitted, adds a swap file to <i>root</i> .
<code>-mirrored -cleanup</code>	CAUTION. USE WITH CARE. Removes all <i>DiskSuite</i> information from the system. On a mirrored system, <i>-cleanup</i> should only be used with <i>-mirrored</i> .

Example *olds* Commands

The table that follows lists and explains example *olds* commands

Example *olds* Mirroring Commands

Example Command	Explanation
<code>olds -mirrored -check_disks c0t3d0</code>	Checks the partitioning of disk c0t3d0.
<code>olds -mirrored -mk_files c0t1d0 c0t2d0</code>	Generates <i>md.tab.new</i> and <i>vfstab.new</i> files to include only the disks c0t1d0 and c0t2d0.
<code>olds -mirrored -mk_files</code>	Generates <i>md.tab.new</i> and <i>vfstab.new</i> files to include all the disks on the system.
<code>olds -mirrored -metadbs</code>	Activates metadatabase replicas (if not already active).
<code>olds -mirrored -setup</code>	Activates or grows the <i>/cms</i> mirror.
<code>olds -mirrored -cleanup</code>	Removes all <i>DiskSuite</i> information from the system. Use this option with <u>extreme</u> caution. Before you do a <i>cleanup</i>, always make sure you have a backup.
<code>olds -addswapfile /cms</code>	Adds a swap file to the <i>/cms</i> filesystem.
<code>olds -addswapfile</code>	Adds a swap file to the <i>/ (root)</i> filesystem.

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 failsafe 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 ← Following an init 0  
# eeprom boot-device=bootdevice2 ← From the command line
```

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 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 save and write 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
#
```

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
    .
    .
Enter choice (1-9) or q to quit:
```

Doing a CMSADM Backup on a Mirrored System

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.

NOTE:

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.

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 -icmud -C 10240 -I <device>
-M "Please remove the current tape, insert tape
number %d and then press ENTER" "/cms" "/cms/*"
```

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

- | | |
|-------------|---|
| /dev/rmt/0 | The internal, noncompressing tape drive (14-GB, 8-mm drive) with the lowest target address. |
| /dev/rmt/1 | The external, noncompressing tape drive (QIC-150 or 5-GB, 8-mm drive) with the second lowest target address. |
| /dev/rmt/0c | The internal, compression-capable tape drive (usually a 14-GB) with the lowest target address. |
| /dev/rmt/1c | 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 *CenterVu CMS R3V5 Sun Enterprise Software Installation (585-215-836)* document for details on restoring a maintenance backup.

Recreating the /cms File System

This section contains procedures for reinitializing the /cms file system and restoring CMSADM backup data. You may need to perform this procedure in case of disk corruption or some other catastrophic system problem.

WARNING:

The first command in this procedure is a cleanup command. This command effectively removes everything from /cms. Be certain you need to do this procedure before you go about destroying the file system.

1. Remove stale information from system files, and clear and reinitiate the metadevice, by entering the following command:

```
# olds -cleanup -mirrored
```

2. Reboot the system with the following command:

```
# init 6
```

3. When the system prompt reappears, log on as the *root* user.

4. Check disk partitioning by entering the following command:

```
# olds -check_disks -mirrored
```

The system responds (for example):

```
.  
. .  
disk:cot0d0 is partitioned ok  
disk:cot1d0 is partitioned ok  
disk:cot2d0 is partitioned ok  
Warning: Current Disk has mounted partitions  
disk:cot0d0 is partitioned ok  
Success, checking disks  
#
```

5. Enter the following command::

```
# /olds/olds -mk_files -mirrored
.
.
.
Success, creating md.tab.new and/or vfstab.new
#
```

6. Set up the /cms metadvice with an `olds -setup` command. For example:

```
# nohup ./olds/olds -setup -mirrored | tee
.
.
.
super-block backups (for fsck -F ufs -o b=#) at:
32, 16240, 32448, 48656, 64864, 81072, 97280,
113488,
.
.
.
10532656, 10548864, 10565072, 10580000, 10596208
ufs fsck: sanity check: /dev/md/rdisk/d19 okay
Success, activating or growing /cms metadvice
#
```

7. Mount the file system with the command: `mount /cms`

8. Verify disk space by entering the command: `df -k`

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

Filesystem	kbytes	used	avail	capacity	Mounted on
/dev/dsk/c0t3d0s0	10772308	310956	383882	45%	/
/proc	0	0	0	0%	/proc
fd	0	0	0	0%	/dev/fd
/dev/md/dsk/d19	xxxxxxxx	9	yyyyyy	0%	/cms

Using the information in the response, perform the calculations shown in Table to verify that *Solstice DiskSuite* is administering all the disks in your system.

If there seems to be a problem, check to see that the file `/etc/opt/SUNWmd/md.tab` accounts for all your disk drives. The `d19` line of the `#/cms` section should reflect the precise number of disk drives on your system.

Disk Space Verification (4-GB Disks)

Calculation	Result
(1a) Divide the /cms line's 'kbytes' figure by 1000 (move the decimal point three places to the left):	(1a)
(1b) Enter the appropriate figure from the "/cms Size (MB)" column below:	(1b)
Note: (1a) and (1b) should be <i>approximately</i> equal (the correspondence will not be exact, but must be within 4000MB)	

No. of 4-GB Disks	/cms Size (MB)
1	3104
2	7137
3	11170
4	15203
5	19236
6	23269
7	27302
8	31335
9	35368
10	39401

9. Install the swap file on the CMS file system with the following command:

```
# /olds/olds -addswapfile /cms
```

10. Reboot the system with the command: `init 6`

11. To restore the latest available CMSADM backup, load the tape into the tape drive and enter a command similar to this:

```
# cpio -icmud -C 10240 -I <device>
-M "Please remove the current tape, insert tape
number %d and then press ENTER" "/cms" "/cms/*"
```

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

`/dev/rmt/0` Indicates the internal, noncompressing tape drive (14-GB, 8-mm drive) with the lowest target address.

- `/dev/rmt/1` Indicates the external, noncompressing tape drive (QIC-150 or 5-GB, 8-mm drive) with the second lowest target address.
- `/dev/rmt/0c` Indicates the internal, compression-capable tape drive (usually a 14-GB tape drive) with the lowest target address
- `/dev/rmt/1c` Indicates 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.

12. Restore any *CentreVu* CMS maintenance backups dated after the latest CDMADM backup.

For details on restoring a maintenance backup, see the software installation document appropriate to your version of CMS.

Restoring an Entire System from a CMSADM Backup

If you have a disk crash that disables the operating system disk drive, you must perform the procedure outlined in this section. The procedure is essentially a new install, with some variation to accommodate the probable availability of a CMSADM backup tape.

1. Replace the crashed disk drive. For instructions see the Lucent Technologies hardware installation document appropriate to your hardware and the version of CMS you are running.
2. Perform Steps 2 through 7 of the factory installation procedure in Chapter 2 (page 15 through page 19).

3. Do this step only if you have a CMSADM backup available.

If you have a CMSADM backup available, do the following:

- a. Restore the backup. For detailed instructions see the Lucent Technologies installation or maintenance document appropriate to your CMS version.
- b. Run CMS setup. For detailed instructions see the Lucent Technologies software installation document appropriate to your CMS version.

4. Do this step only if you have no CMSADM backup available.

If you have no CMSADM backup available, perform steps 8 and 9 of the factory installation procedure in Chapter 2 (page 20).

5. Restore CMS data from the latest available CMS maintenance backup. For detailed instructions see the Lucent Technologies maintenance document appropriate to your version of CMS.
6. If you were unable to restore the system from a CMSADM backup as instructed in step 3, readminister your terminals, printers, modems, and other peripherals as needed.

Restoring Specific Files from a CMSADM Backup

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

```
# ulimit unlimited
# 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 this example is */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. Device names are as follows:

- **/dev/rmt/0** indicates the noncompressing tape drive with the lowest target number in the SCSI chain.
- **/dev/rmt/1** indicates the noncompressing tape drive with the second lowest target number in the SCSI chain.
- **/dev/rmt/0c** indicates the compression-capable tape drive with the lowest target number in the SCSI chain.
- **/dev/rmt/1c** indicates the compression-capable tape drive with the second lowest target number in the SCSI chain (the QIC 2.5-GB tape drives support compressed mode).

Adding New Disks to an Already-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, you must add the disks manually. For factory boot disk assignments, see the table immediately below; for instructions on adding disks manually, see “Adding New Disk Drives Manually” on page 39.

Boot-disk Assignments for Factory-installed Mirroring

Boot Disk 1	Boot Disk 2
c0t0d0	c0t1d0

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.

Before you begin, install the new disk drives into your system. For instructions, see the documentation that accompanied your unit or system. Remember that, in a mirrored system, disk drives must be added in pairs in order to increase system capacity and use all the disk space available.

Once the drives are in place, configure your system for the new drives by doing the following steps:

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. You may not do this procedure while any metadevice is resynching.**

2. Reboot the system by entering the following commands:

```
# init 0
# boot -r
```

3. Partition the new disks by following the instructions in “Partitioning Disks Via the `format` Command” on page 49.
4. Stop CMS (via the `cms_run` option of `cmsadm` or `cmssvc`).

5. 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: _____

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

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

7. Load the file `/etc/opt/SUNWmd/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 `md.tab` file format, see the table “Valid Example Files” on page 48.)

8. 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.orig
```

9. Replace the existing `md.tab` file with the one you created in Step 7 by entering the following command:

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

10. 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:

Adding New Disks to an Already-mirrored System

```
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.
#
```

11. Restart CMS by entering the `cmsadm` or `cmssvc` command and selecting the `cms_run` option. When the run option menu displays, select `start cms`.
12. 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 4; it should show an increase roughly equal to one-half of the total capacity of the drives added.

Adding New Disk Drives Manually

A “manual” 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. Stop CMS. (From the command line, enter the `cmsadm` or `cmssvc` command, and select the `cms_run` option. When the `cms_run` option menu displays, select `stop cms`.)

2. Halt the system by entering the following command:

```
# init 0
```

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 *CentreVu CMS R3V5 Sun Enterprise 3000 Installation and Maintenance (585-215-837)* document.

4. Reboot the system by entering the following command:

```
# 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
/proc	0	0	0	0%	/proc
fd	0	0	0	0%	/dev/fd
/dev/md/dsk/d21	xxxxxxxx	9	yyyyyy	0%	/cms

Check the `/cms` line of the response. Look for the “kbytes” figure in that line (shown as “xxxxxxxx” in the example above). Record that figure here:

/cms “kbytes” figure: _____

6. Partition the new disks. Follow the instructions in “Partitioning Disks Via the format Command” on page 49. Make sure to double check the partitioning before continuing.

Adding New Disk Drives Manually

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

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

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 */etc/opt/SUNWmd/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 the table “Valid Example Files” on page 48.)

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.orig
```

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

```
# cp /etc/opt/SUNWmd/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.

13. Deactivate the */cms* swap file by entering the following command:

```
# swap -d /cms/swap
```

14. 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 command:

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

- b. Reactivate the */cms* swap file by entering the command:

```
# swap -a /cms/swap
```

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

15. Enter the following command:

```
# metastat
```

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

16. 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
/proc	0	0	0	0%	/proc
fd	0	0	0	0%	/dev/fd
/dev/md/dsk/d21	xxxxxxx	9	yyyyyy	0%	/cms

Check the */cms* line of the response. Look for the “kbytes” figure in that line (shown as “xxxxxxx” in the example above). Compare that figure to the figure recorded in Step 5; they should be the same.

17. Reactivate the */cms* swap file by entering the following command:

```
# swap -a /cms/swap
```

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:
/iommu@0,10000000/sbus@0,10001000/dma@1,81000/esp@1,80000/sd@2,0
(sd17):
```

Match up the device description lines to determine which disk drive has problems

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          79 Jun 10 11:58
/dev/dsk/c0t2d0s7 ->
../../../../devices/iommu@0,10000000/sbus@0,10001000/dma@1,81000/esp@1,
80000/sd@2,0:h
lrwxrwxrwx  1 root    root          79 Jun 10 11:58
/dev/dsk/c0t2d0s6 ->
.
.
.
#
```

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 *olds* directories to your path:

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

2. Use the `metastat` command to list the metadevices.
3. Scan the output 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:

Replacing a Faulty Disk

```
# metadetach d21 d20
```

If the faulty disk is a boot disk, a `metastat` command would have also shown 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
# metadetach d21 d20
```

Step 4: Remove State Database Replicas (Boot Disks Only)

When you are replacing 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 43).

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.

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 luo   16      1034      /dev/dsk/c0t0d0s1
a   p luo   1050   1034      /dev/dsk/c0t0d0s1
a       u    16      1034      /dev/dsk/c0t1d0s1
a       u   1050   1034      /dev/dsk/c0t1d0s1
# metadb -d c0t1d0s1
# metadb -i
  flags      first blk  block count
a m p luo   16      1034      /dev/dsk/c0t0d0s1
a   p luo   1050   1034      /dev/dsk/c0t0d0s1
```

Step 5: Install the New Disk Drive

1. 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.
2. Partition the new disks using the `format` command. See “Partitioning Disks Via the `format` Command” on page 49. **Make sure you partition the correct drive!**

Step 6: Recreate the State Database Replicas (Boot Disks Only)

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

```
metadb -a -c2 -f /dev/rdisk/devname
```

Replace *devname* with the device name of the new, second boot disk. For example:

```
metadb -a -c2 -f /dev/rdisk/c0t4d0s1
```

2. 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, for example, shows three replicas on `c0t4`:

flags	first blk	block count	
a m pc luo	16	1034	/dev/dsk/c0t4d0s1
a pc luo	1050	1034	/dev/dsk/c0t4d0s1
a pc luo	16	1034	/dev/dsk/c0t4d0s1

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

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 the table “Valid Example Files” on page 48).
2. Enter the `metadb` command:

```
metadb -a -c2 -f /dev/dsk/c0t0d0s1
metadb -a -c2 -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 -c2 -f /dev/dsk/c0t0d0s1
```

Adding Replicas

Every nonboot disk has space for adding a replica on partition `s0`. If you must add more replicas do the following steps:

1. Identify a nonboot disk.

The disk `c0t14d0`, for example, is a nonboot disk on an *Enterprise 3000* system.

2. Enter the following command:

```
# metadb -d -c1 -f /dev/dsk/c0t14d0s0
```

- Verify that a new replica was created by entering the following command:

```
# metadb -i
```

Valid Example Files

Valid *vfstab* file for a nonmirrored R3V5 CMS system with *Solstice DiskSuite*:

#device	device	mount	FS	fsck	mount	mount
#to mount	to fsck	point	type	pass	at boot	options
#						
#/dev/dsk/c0d0s2	/dev/rdisk/c0d0s2	/usr	ufs	1	yes	-
/proc	-	/proc	proc	-	no	-
fd	-	/dev/fd	fd	-	no	-
#swap	-	/tmp	tmpfs	-	yes	-
<hr/>						
/dev/dsk/c0t3d0s0	/dev/rdisk/c0t0d0s0	/	ufs	1	no	-
/dev/md/dsk/d19	/dev/md/rdisk/d19	/cms	ufs	2	yes	-

Valid *vstab* file for an R3V5 CMS system with mirrored disks:

#device	device	mount	FS	fsck	mount	mount
#to mount	to fsck	point	type	pass	at boot	options
#						
#/dev/dsk/c0d0s2	/dev/rdisk/c0d0s2	/usr	ufs	1	yes	-
/proc	-	/proc	proc	-	no	-
fd	-	/dev/fd	fd	-	no	-
#swap	-	/tmp	tmpfs	-	yes	-
<hr/>						
/dev/md/dsk/d13	/dev/md/rdisk/d13	/	ufs	1	no	-
/dev/md/dsk/d21	/dev/md/rdisk/d21	/cms	ufs	2	yes	-

Standard Mirrored DiskSuite System

Valid *md.tab* file for a system with factory-installed mirroring and four internal disks:

```
state database replicas
mddb00 c0t0d0s1
mddb01 c0t1d0s1
#metaroot
d11 1 1 c0t0d0s0
d12 1 1 c0t1d0s0
d13 -m d11
#/cms
d19 2 1 /dev/rdisk/c0t0d0s3 1 /dev/rdisk/c0t2d0s1
d20 2 1 /dev/rdisk/c0t1d0s3 1 /dev/rdisk/c0t3d0s1
d21 -m d19
```

Partitioning Disks Via 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!)

Disk Type	Part	Tag	Flag	CMS R3V5 4.2-GB Disks		CMS R3V6 4.2-GB Disks	
				Starting Cylinder	Size	Starting Cylinder	Size
Boot	0	<i>boot disk 1: root</i> <i>boot disk 2: unassigned</i>	wm	0	761c	0	1023c
	1	unassigned	wm	761	71c	1023	7c
	2	overlap <i>or</i> backup	wm	0	3880c	0	3880c
	3	unassigned	wm	832	3048c	1030	2788c
	4	<i>R3V5: unassigned</i> <i>R3V6: swap</i>	wm	0	0	3818	62c
	5-7	unassigned	wm	0	0	0	0
Nonboot	0	unassigned	wm	0	2c	0	2c
	1	unassigned	wm	2	3878c	2	3878c
	2	overlap <i>or</i> backup	wm	0	3880c	0	3880c
	3-7	unassigned	wm	0	0	0	0

Partitioning Disks Via the format Command

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 -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 Mirroring was a Field Upgrade

To tell whether mirroring was installed at the factory or as a field upgrade, you can do one of two things. Either:

- a. Enter the command: `# metastat`
- 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.

Mirroring that was installed at the factory will have the structure indicated in the following table:

Root Metadevice Device Names

Metadevice	Device Name
metadevice d11	c0t0d0
metadevice d12	c0t1d0

How to Tell Whether Mirroring was a Field Upgrade

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: `swap -d /cms/swap`
Removes the logical file.
 - b. Enter: `rm /cms/swap`
Removes the physical file.
 - c. Enter: `mount /cms`
Remounts /cms.
 - d. Enter: `olds -addswapfile /cms`
Recreates the swap file.
 - e. 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 (add or replace disks, etc.) 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 procedure:

1. Reboot the system with an `init 0` command. The system reboots and displays the `ok` prompt:

```
# init 0
. . .
ok
```

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 SCSI devices, including the new disk drives, enter the following command:

```
ok probe-scsi-all
```

The program should respond as follows:

```
/iommu@f,e0000000/sbus@f.e0001000/esp@3,200000
Target 1
Unit 0 Disk SEAGATE ST14801 SUN04246266 Copyright (C) 1991
Target 3
Unit 0 Disk SEAGATE ST14801 SUN04246266 Copyright (C) 1991
.
.
.
Target 6
Unit 0 Disk Removable Read Only Device SONY CD-ROM CDU-8012
ok
```

NOTE:

If you receive the following warning 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?
```

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).

You may now repeat the `probe-scsi-all` and perform any other boot prom level diagnostics.

- d. **IMPORTANT!** Enter: `setenv auto-boot? true`

WARNING:

Failure to do a `setenv` command before the next reboot will cause the reboot to stop at the boot prompt instead of proceeding through the normal boot-up.

7. 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
.
.
.
hostname 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 the *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

d12: Submirror of d13
  State:
  Invoke:
  Size: 2209680 blocks
  Stripe 0:
    Device    Start Block Dbase State Hot Spare
    c0t1d0s0    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. For example:

```
# 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 with the command: `# init 0`
2. From the `ok` prompt enter: `ok> printenv boot-device`
The system responds: `boot-device <disk>`
3. If `disk` is not the appropriate boot device, change it with the following command:

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

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:

Log Files

File Name	Description
<code>/var/adm/messages</code>	Contains all console error messages, including hardware messages.
<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.

The messages sometimes indicate problems. For example:

```

Jun 12 16:27:08 leopard unix: WARNING:
/iommu@0,10000000/sbus@0,10001000/dma@1,81000/esp@1,80000/sd@2,0
(sd17):
Jun 12 16:27:08 leopard unix:  Error for command 'read(10)' Error
Level: R
. . .
Jun 13 10:02:27 leopard unix: WARNING: md: d20: read error on
/dev/dsk/c0t2d0s1
Jun 13 10:16:03 leopard unix:  Polled cmd failed
Jun 13 11:00:06 leopard unix: WARNING:
/iommu@0,10000000/sbus@0,10001000/dma@1,81000/esp@1,80000/sd@2,0
(sd17):
Jun 13 11:00:06 leopard unix:  SCSI transport failed: reason
'reset': givi
. . .

```

The messages illustrated above (for the host named *leopard*) indicate a potential disk failure. In particular, disk `c0t2d0` seems to have a problem.

At this point some troubleshooting will be necessary in order to determine the nature of the problem. If a disk needs to be replaced see “Replacing Faulty Disks” in Chapter 3.

Problem Notification

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, check to see 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 maintenace before the operation can be carried out.
Response:	<ol style="list-style-type: none"> 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 in order for the command to execute, the CDROM drive must be empty.
Response:	<ol style="list-style-type: none"> 1. Enter the command: <code>eject cdrom</code> 2. Remove the CD from the drive tray. 3. Close the drive tray.

Message:	cms does not support DiskSuite with disks > 4.2Gb
Response:	Don't use disks bigger than 4.2GB.

Error Messages (Contd)

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>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>DiskSuite</i> setup.
Response:	If the device is <code>c0t6d0</code> , 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 (Contd)

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 readded 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:	System needs at least three replicas to be functional.
Response:	Use <code>metadbs -i</code> to verify that at least three replicas are active. If not, add replicas (see Chapter 3).

Message:	Failed: disk cxytd0 or cx'ty'd0 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:	ksh: <i>command</i> : not found
Explanation:	Command does not exist. Probably the location of the command file has not been added to your path.
Response:	Use the entire path to execute the command—as <code>pathname/command</code> —or add the pathname to your path with the following command: <pre># export PATH=\$PATH:pathname</pre> <p>The path of the <code>metastat</code>, <code>metadb</code>, <code>metattach</code>, and <code>metadetch</code> commands is <code>/usr/opt/SUNWmd/sbin</code>; the path of the <code>olds</code> command is <code>/olds</code></p>

Error Messages (Contd)

Message:	metadb: <i>hostname</i> : cxyd0sz has a metadvice database 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 metadvice while a resync is in progress.
Response:	Try again later. To find out how far along the resync has progressed, enter a <code>metastat</code> 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 <code>cpio</code> command: <pre>cpio -icmudf -C10240 -I input-device -M "Enter tape.. " .."/cms*..</pre> The "/cms" entry has probably been mistyped.
Response:	Remove the /cms directory and all files and try again.

Message:	Overlapping swap files are not allowed
Explanation:	You are trying to add a swap file when one already exists.
Response:	No response is necessary. The system ignores your request and keeps the old swap file. To recreate a swap file, first remove the old file with the following command: <pre>swap -d /cms/swap</pre>

Error Messages (Contd)

Message:	panic: vfs_mountroot: cannot mount root rebooting... resetting....
Explanation:	The system is misconfigured.
Response:	There are two possible causes: <ol style="list-style-type: none"> 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: <ol style="list-style-type: none"> a. Boot from the CD. b. Mount <code>/dev/dsk/c0t0d0s0 /a</code> (or whatever partition is your root partition). c. Remove <i>DiskSuite</i> entries from <code>/a/etc/system</code> (i.e. everything between tags "mddb ...") d. Remove <code>/a/etc/opt/SUNWmd/md*</code> files, i.e.: <code>rm /a/etc/opt/SUNWmd/md*</code> e. Replace <code>/dev/md/rdisk/d21</code> entries with <code>/dev/rdisk/c0t3d0s0</code> (or whatever your boot device name is). f. Boot with the command: <code>init 6</code> g. If the system still refuses to reboot, reload <i>Solaris</i>.

Message:	<i>filename restored from 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 resynced. Try again later. You can tell whether a resync is in progress by entering a <code>metastat</code> command.

Error Messages (Contd)

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>DiskSuite</i> before doing this task.

Message:	Swap file cannot be placed on xx filesystem
Explanation:	Tried to add swap file to a system other than root or /cms.
Response:	Reenter command with either /cms (to add the /cms swap file) or no name at all (to add the <i>root</i> swap file).

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

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 <i>md.tab</i> or <i>vfstab</i> 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:	The /cms filesystem must be mounted for the command to work.
Response:	Mount /cms with the command: <code>mount /cms</code>

Error Messages (Contd)

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:	<ol style="list-style-type: none"> 1. Enter: N (to prevent the probe from continuing). 2. Enter: setenv auto-boot? false (to keep the system from rebooting when you do the reset-all) 3. Enter: reset-all <p>It may take a minute to complete. Once it does, you may do the probe-scsi or probe-scsi-all and perform any other boot prom level diagnostics.</p> <ol style="list-style-type: none"> 4. IMPORTANT: Before you reboot again, enter the command: <pre style="margin-left: 40px;">setenv auto-boot? true</pre> <p>Failure to do so will cause the reboot to stop at the boot prompt instead of proceeding through the normal boot-up.</p>

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.

Error Messages (Contd)

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”.

Message:	Usage: olds -mirrored -check_disks -mirrored -mk_files -metadbs -cleanup -mirrored -setup [disk name]
Explanation:	You have mistyped a command. Accompanied by a number of “example” messages.
Response:	Find out how you have misused the command and retry it. See “Using the olds Scripts” in Chapter 3.

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

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 the appropriate hardware for your system (internal disks only for the <i>Enterprise 3000</i>). If you need other types of disks for other purposes, you may attach them after the disk mirroring scripts have run successfully with the supported hardware.

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

Error Messages

Error Messages (Contd)

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>DiskSuite</i> before attempting this task.

Index

A	
adding disk drives manually	39
adding replicas	47
alternate boot device	
creating	13, 21
testing	14
using	25

B	
backing up	
CMSADM backup	27–28
CMSADM restore	29–35
boot devices	36
boot disk designations	15, 51
boot problems	57

C	
chkDisks	21, 26, 58, 59
chunks	2
CMSADM backup	27–28
CMSADM restore	29–35
concatenated metadevice	2
capacity	2
<i>cron</i> file	14, 21

D	
d11 metadevice	4
d12 metadevice	4
d19 metadevice	4
d20 metadevice	4
d21 metadevice	4
disk drives	
adding	36–41
recognition problem	54
replacing	42–46
restoring	34
disk mirroring	
definition	1
factory installation	15–21
upgrading an <i>Enterprise 3000</i> system	7–14
disk partitioning	49
disk partitions	1
<i>disks.type</i> file	17
<i>DiskSuite</i>	
see <i>Solstice DiskSuite</i>	
drive slots	7

E	
<i>Enterprise 3000</i>	
disk partitioning	10, 16, 49
upgrading	7–14
error messages	60–68

F	
factory installation procedures	15–21
faulty disk	
identifying	42
replacing	42–46

I	
installation procedures, factory	15–21
installing mirrored systems	7

K	
Korn shell	17

L	
log files	58–59

M	
<i>md.tab</i> file	11
metadevices	
concatenated	2
definition	1
mirror disks	
partitioning	9
mirrors	
and metadevices	3
definition	1
identifying	8
installing	7
layout	5
maintaining	23
required software	6

N	
<i>needs maintenance</i> messages	56

O	
<i>olds</i> scripts	23–24
commands	24
examples	24
format	23
parameters	23

P

partitioning disks	49
printing <i>vfstab</i>	27

R

replacing a faulty disk	42–46
replicas	
adding	47
replacing	47
setup	47–49
restoring a CMSADM backup	29–35
resync problem	53

S

slices	1
<i>Solstice DiskSuite</i>	1
installation problems	53
setup scripts	17
state database	
definition	3
state database replicas	3
state database replicas	
adding	47
replacing	47
setup	47–49
submirrors	3

T

three-way mirror	3
two-way mirror	3

U

upgrading	
and increasing system capacity	7
<i>Enterprise 3000</i> system	7–14

V

<i>vfstab</i> file	12
printing	27
sample	48