



CentreVu[®] Call Management System

Release 3 Version 8

Disk-Mirrored Systems

585-210-940

Issue 2.2

October 2002

© 2001 Avaya Inc.
All Rights Reserved

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.

Preventing Toll Fraud

"Toll fraud" is the unauthorized use of your telecommunications system by an unauthorized party (for example, a person who is not a corporate employee, agent, subcontractor, or working on your company's behalf). Be aware that there may be a risk of toll fraud associated with your system and that, if toll fraud occurs, it can result in substantial additional charges for your telecommunications services.

Avaya 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.

Providing Telecommunications Security

Telecommunications security (of voice, data, and/or video communications) is the prevention of any type of intrusion to (that is, either unauthorized or malicious access to or use of your company's telecommunications equipment) by some party.

Your company's "telecommunications equipment" includes both this Avaya product and any other voice/data/video equipment that could be accessed via this Avaya product (that is, "networked equipment").

An "outside party" is anyone who is not a corporate employee, agent, subcontractor, or working on your company's behalf. Whereas, a "malicious party" is anyone (including someone who may be otherwise authorized) who accesses your telecommunications equipment with either malicious or mischievous intent.

Such intrusions may be either to/through synchronous (time-multiplexed and/or circuit-based) or asynchronous (character-, message-, or packet-based) equipment or interfaces for reasons of:

- Utilization (of capabilities special to the accessed equipment)
- Theft (such as, of intellectual property, financial assets, or toll-facility access)
- Eavesdropping (privacy invasions to humans)
- Mischief (troubling, but apparently innocuous, tampering)
- Harm (such as harmful tampering, data loss or alteration, regardless of motive or intent)

Be aware that there may be a risk of unauthorized intrusions associated with your system and/or its networked equipment. Also realize that, if such an intrusion should occur, it could result in a variety of losses to your company (including but not limited to, human/data privacy, intellectual property, material assets, financial resources, labor costs, and/or legal costs).

Your Responsibility for Your Company's Telecommunications Security

The final responsibility for securing both this system and its networked equipment rests with you - an Avaya customer's system administrator, your telecommunications peers, and your managers. Base the fulfillment of your responsibility on acquired knowledge and resources from a variety of sources including but not limited to:

- Installation documents
- System administration documents
- Security documents
- Hardware-/software-based security tools
- Shared information between you and your peers
- Telecommunications security experts

To prevent intrusions to your telecommunications equipment, you and your peers should carefully program and configure your:

- Avaya-provided telecommunications systems and their interfaces
- Avaya-provided software applications, as well as their underlying hardware/software platforms and interfaces
- Any other equipment networked to your Avaya products.

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.

European Union Declaration of Conformity

Avaya Business Communications Systems declares that equipment specified in this document conforms to the referenced European Union (EU) Directives and Harmonized Standards listed below:

EMC Directive 89/336/EEC

Low Voltage Directive 73/23/EEC



The "CE" mark affixed to the equipment means that it conforms to the above Directives.

Trademarks

See the Preface of this document for trademark information.

Avaya National Customer Care Center

Avaya provides a telephone number for you to use to report problems or to ask questions about your contact center. The support telephone number is 1-800-242-2121.

Avaya Web Page

<http://www.avaya.com>

Comments

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

Acknowledgment

This document was written by the CRM Development group.

Table of Contents

Preface	ix
Purpose	ix
Organization	ix
Reasons for reissue	ix
Conventions	ix
Related documents	xi
Hardware documents	xi
CMS software documents	xii
Switch documents	xii
Upgrade documents	xiii
Administration documents	xv
Other documents	xv
Documentation web sites	xvi
Helplines	xvii
Support web site	xvii
Customer support for the United States	xvii
Technician support for the United States	xvii
Customer and technician support outside the United States	xvii

Chapter 1: Introduction to Disk Mirroring

Contents	1
Disk mirroring defined	1
Metadevices	1
Concatenated metadevices	2
State databases	3
State database replicas	3
Mirrors	4
How CMS implements disk mirroring	5
Mirror layouts	7
Drive Device Name	7
Mirror layout on a Sun Enterprise 3500 computer	7
Mirror layout on a Sun Enterprise 3000 computer	8
Additional references	8

Chapter 2: Installing Mirrored Systems

Contents	9
Hardware Requirements	9
Overview	9
Contents	9

Required additional hardware	10
Required software	11
Hard drive slots	12
Enterprise 3000 Standard Configuration	12
Enterprise 3000 Diagram.	12
Enterprise 3500 Standard Configuration	13
Enterprise 3500 Diagram.	13
Increasing system capacity	13
How to upgrade an Enterprise platform to disk mirroring	14
Overview	14
Contents	14
Prerequisites	14
Identifying mirror devices	15
Identifying mirror devices on a E3500 system	15
Device/Metadevice Names for an Enterprise 3500 system table	15
Identifying mirror devices on a E3000 system	15
Device/Metadevice Names for an Enterprise 3000 system table	16
Installing the mirror disks	17
Installing the mirror disks on an Enterprise 3500 system.	17
Installing the mirror disks on an Enterprise 3000 system.	20
Partitioning the mirror disks.	20
Disk partitioning table.	21
CMS R3V8 Boot Disk Partition Values	21
CMS R3V8 Nonboot Disk Partition Values	21
.	22
Creating the state database replicas	23
Creating the md.tab table entries	25
Setting up the mirrors for the root file system.	26
Setting up the mirrors for the swap space.	28
Setting up the /cms mirror	31
Verify disk space.	33
Create the alternate boot device	33
Factory installation procedures for disk mirroring	34
Contents	34
Installing the hardware	34
Installing the operating system.	35
Disks table.	35
Disk partitioning tables	36
CMS R3V8 Boot Disk Partition Values	36
CMS R3V8 Nonboot Disk Partition Values	36
.	37
Installing the application software	37
Checking the disk.type file	37
Setting up the disk subsystem for Solstice DiskSuite	39
Running the setup scripts.	41
Verify disk space.	43
Install CMS	43
Create the alternate boot device	43
Verifying disk space	44
Setting up the alternate boot device.	46
Obtaining the correct disk device definition	46

Default device definitions table	46
Disk device definitions for Enterprise 3500 and 3501/3503 systems	46
Disk device definitions for Enterprise 3000 systems	46
Creating the alternate boot device	47

Chapter 3: Maintaining Mirrored Systems

Contents	51
How to tell whether your mirrored system is standard	52
Root metadvice device name table	52
Performing a CMSADM backup on a mirrored system	53
Overview	53
Contents	53
Printing the vfstab file	53
Running the backup	53
Labeling and storing the tapes	54
Adding new disks to a standard mirrored system	55
Overview	55
Boot-disk assignments for factory installed mirroring	55
Procedure	55
Adding new disks to a nonstandard mirrored system	59
Overview	59
Procedure	59
Identifying a faulty disk	62
Overview	62
Contents	62
Identifying the disk	62
Identifying the submirrors	64
Recovering a mirrored system after disk failure	65
Overview	65
Prerequisites	65
Contents	65
Determining the recovery procedure	66
Recovering a mirrored system that will boot into multiuser mode	68
Contents	68
Removing the failed disk	68
Re-creating the database replicas	70
Syncing the new disk	70
Recovering a mirrored system that will only boot into single user mode	73
Contents	73
Removing the failed disk	73
Re-creating the database replicas	75
Syncing the new disk	75
Partitioning disks with the format command	78
Disk partitioning tables	79
CMS R3V8 Boot Disk Partition Values	79
CMS R3V8 Nonboot Disk Partition Values	79
.	80

Restoring Data from a CMSADM Backup	82
Overview	82
Tape drive considerations.	82
Procedure	82
Repairing or rebuilding the /cms file system	84
Overview	84
Contents	84
Repairing /cms	84
Rebuilding /cms	86
Overview.	86
Prerequisites	86
Procedure	86
Restoring specific files from a CMSADM backup	89
Restoring a mirrored system from a CMSADM backup	90
Restoring the backup	90
Restoring the backup on a E3000 system	90
Restoring the backup on a E3500 system	93
Re-establishing the disk mirrors	96
Re-establishing the mirrors on the E3000 system.	96
Re-establishing the mirrors on a E3500 system.	104
State Database Replicas	114
Creating replicas.	114
Replacing replicas.	114
Maintaining the chkDisks crontab	115
Overview	115
Contents	115
Activating chkDisks	115
Verifying chkDisks	115
Changing the scheduled run time	116
Canceling chkDisks	116
Examples of valid system files	117
Valid vfstab files	117
Valid md.tab files.	118

Chapter 4: Troubleshooting

Contents	119
Solstice DiskSuite problems	119
The root filesystem runs out of space during a system install.	119
Excessively long resync.	120
The system fails to recognize all disk drives	121
Fixing “needs maintenance” messages	123
Boot problems.	125
Contents	125
Trying to boot from the wrong device	125
Primary boot device is not available	126

Excessive reboot time.	126
Log files	127
Error messages	128
Index	137

Preface

Purpose

This document is written for technicians and call center customers who install, provision, and maintain Release 3 Version 8 of the CentreVu Call Management System (CMS) using the “disk mirroring” feature.

Organization

This document includes the following chapters:

- Chapter 1 - [“Introduction to Disk Mirroring” on page 1](#)
- Chapter 2 - [“Installing Mirrored Systems” on page 9](#)
- Chapter 3 - [“Maintaining Mirrored Systems” on page 51](#)
- Chapter 4 - [“Troubleshooting” on page 119](#)

Reasons for reissue

The information in this document is being reissued for the following reasons:

- Updated the procedures [“Creating the state database replicas” on page 23](#), [“Setting up the mirrors for the root file system” on page 26](#), [“Setting up the mirrors for the swap space” on page 28](#), [“Setting up the /cms mirror” on page 31](#).
- Clarified the caution statements in [“Restoring a mirrored system from a CMSADM backup” on page 90](#).

Conventions

The following conventions are used in this document:

- Unless specified otherwise, all information and procedures in this document apply to the Sun Enterprise 3000 computer and the Sun Enterprise 3500 computer.
- The term CMS in this document always implies CentreVu CMS.
- Commands you enter from the console are shown in **bold courier** font.
- Keyboard commands are shown in **bold** text.
- Screens are shown to represent responses from the system.

Because of display constraints in this document, some screen representations are not identical to the screens on your system.

- *Italic* text represents variable information in commands or screen output.

Related documents

This section lists sources for related information about call center products and features. To order Avaya documentation, call the Avaya Publications Center at 1-800-457-1235 or +1-317-361-5353.

Hardware documents

Title	Document number
<i>CentreVu Call Management System Sun Enterprise 3500 Computer Hardware Installation, Maintenance, and Troubleshooting</i>	585-215-873
<i>CentreVu Call Management System Sun Enterprise 3500 Computer Connectivity Diagram</i>	585-215-877
<i>CentreVu Call Management System Sun Ultra 5 Computer Hardware Installation, Maintenance, and Troubleshooting</i>	585-215-871
<i>CentreVu Call Management System Sun Ultra 5 Computer Connectivity Diagram</i>	585-215-872
<i>CentreVu Call Management System Sun Enterprise 3000 and SPARCserver Computers Hardware Maintenance and Troubleshooting</i>	585-214-016
<i>CentreVu Call Management System Release 3 Version 6 Sun Enterprise 3000 Computer Hardware Installation</i>	585-215-867
<i>CentreVu Call Management System Release 3 Version 6 Sun Enterprise 3000 Computer Connectivity Diagram</i>	585-215-865
<i>CentreVu Call Management System Release 3 Version 6 Sun SPARCserver Computers Hardware Installation</i>	585-215-857
<i>CentreVu Call Management System Release 3 Version 6 Sun SPARCserver Computers Connectivity Diagram</i>	585-215-858
<i>CentreVu Call Management System Release 3 Version 5 Sun SPARCserver Installation and Maintenance</i>	585-215-827
<i>CentreVu Call Management System Release 3 Version 5 Sun SPARCserver Connectivity Diagram</i>	585-215-828
<i>CentreVu Call Management System Terminals, Printers, and Modems</i>	585-215-874

CMS software documents

Title	Document number
Installing Software on a CMS Computer	
<i>CentreVu Call Management System Release 3 Version 9 Software Installation, Maintenance, and Troubleshooting</i>	585-215-956
<i>CentreVu Call Management System Release 3 Version 8 Software Installation, Maintenance, and Troubleshooting</i>	585-210-941
<i>CentreVu Call Management System Software Installation and Setup (R3V6 and earlier)</i>	585-215-866
Setting Up a Disk-Mirrored System	
<i>CentreVu Call Management System Release 3 Version 8 Disk-Mirrored Systems</i>	585-210-940
<i>CentreVu Call Management System Disk-Mirrored Systems (R3V6)</i>	585-215-841

Switch documents

Title	Document number
<i>CentreVu Call Management System Switch Connections and Administration</i>	585-215-876

Upgrade documents

There are several upgrade paths supported with CMS. For each of these upgrades, there is a document designed to support that upgrade. Note that none of these documents are available from the publications center, but are available from the [Avaya CMS documentation](#) web site.

- Base load upgrades

A base load upgrade is used when upgrading CMS to a newer load of the same version (for example, R3V8 ak.g to R3V8 al.k). A specific set of instructions are written for the upgrade and are shipped to the customer site with the CMS software CD as part of a Quality Protection Plan Change Notice (QPPCN).

Title
<i>CentreVu Call Management System Release 3 Version 9 Base Load Upgrade Procedures</i>
<i>CentreVu Call Management System Release 3 Version 8 Base Load Upgrade Procedures</i>

- Platform upgrades and data migration

A platform upgrade is used when upgrading to a new hardware platform (for example, upgrading from a SPARCserver 5 to an Enterprise 3500). The new hardware platform ships from the Avaya factory with the latest CMS load. Therefore, as part of the upgrade, you will automatically upgrade to the latest CMS load (for example, R3V8 to R3V9, or a newer load of the same CMS version). A specific set of instructions are written for the upgrade and are shipped to the customer site with the new hardware.

Title
<i>CentreVu Call Management System Release 3 Version 9 Platform Upgrade and Data Migration Instructions</i>
<i>CentreVu Call Management System Release 3 Version 8 Platform Upgrade and Data Migration Instructions</i>

- CentreVu Upgrade Express (CVUE)

CVUE is used in the following conditions:

- CMS is being upgraded from an older version (for example, R3V5u or R3V6) to the latest version (for example, R3V8 or R3V9)
- The hardware platform is not changing.

A specific set of upgrade instructions are written for the upgrade and are shipped to the customer site with the CVUE kit.

Title
<i>CentreVu Call Management System Release 3 Version 9 Sun Ultra 5 Computer CVUE Instructions</i>
<i>CentreVu Call Management System Release 3 Version 9 Sun Enterprise 3000 Computer CVUE Instructions</i>
<i>CentreVu Call Management System Release 3 Version 9 Sun Enterprise 3000 Computer Mirrored System CVUE Instructions</i>
<i>CentreVu Call Management System Release 3 Version 9 Sun Enterprise 3500 Computer CVUE Instructions</i>
<i>CentreVu Call Management System Release 3 Version 9 Sun Enterprise 3500 Computer Mirrored System CVUE Instructions</i>
<i>CentreVu Call Management System Release 3 Version 8 Sun SPARCserver 5 Computer CVUE Instructions</i>
<i>CentreVu Call Management System Release 3 Version 8 Sun SPARCserver 20 Computer CVUE Instructions</i>
<i>CentreVu Call Management System Release 3 Version 8 Sun Ultra 5 Computer CVUE Instructions</i>
<i>CentreVu Call Management System Release 3 Version 8 Sun Enterprise 3000 Computer CVUE Instructions</i>
<i>CentreVu Call Management System Release 3 Version 8 Sun Enterprise 3000 Computer Mirrored System CVUE Instructions</i>
<i>CentreVu Call Management System Release 3 Version 8 Sun Enterprise 3500 Computer CVUE Instructions</i>
<i>CentreVu Call Management System Release 3 Version 8 Sun Enterprise 3500 Computer Mirrored System CVUE Instructions</i>

Administration documents

Title	Document number
<i>CentreVu Call Management System Release 3 Version 9 Administration</i>	585-214-015
<i>CentreVu Call Management System Release 3 Version 8 Administration</i>	585-210-910
<i>CentreVu Call Management System Release 3 Version 6 Administration (Volumes 1 and 2)</i>	585-215-850
<i>CentreVu Call Management System Release 3 Version 5 Administration (Volumes 1 and 2)</i>	585-215-820

Other documents

Title	Document number
<i>CentreVu CMS Open Database Connectivity</i>	585-210-951
<i>CentreVu CMS Release 3 Version 9 External Call History Interface</i>	585-215-952
<i>CentreVu CMS Release 3 Version 5 Real-Time and Historical Reports</i>	585-215-821
<i>CentreVu CMS Release 3 Version 5 Custom Reports</i>	585-215-822
<i>CentreVu CMS Release 3 Version 5 Forecast</i>	585-215-825

Documentation web sites

Use the following sites to view support documentation:

- Sun hardware documentation
<http://docs.sun.com>
- Okidata printer documentation
<http://www.okidata.com>

Helplines

If a problem arises that requires assistance, customers or Avaya technicians can:

- go to the support web site
 - call the appropriate support number
-

Support web site

Use the following site to:

- Register through Online Services. By registering you will be able to create new trouble tickets through the Avaya web server without having to call the 1-800-242-2121 United States helpline.
- Receive software downloads.
- View links to frequently asked questions about call center products and administration.

<http://www.avaya.com/support>

Customer support for the United States

1-800-242-2121

Customers can report problems and generate trouble tickets by calling this number.

The customer is prompted to identify the type of problem (that is, Automatic Call Distribution, hardware, or CentreVu CMS), and is connected to the appropriate service organization.

Technician support for the United States

1-800-248-1234

Avaya technicians can receive help by using this number.

Customer and technician support outside the United States

For customer and technician support outside the United States, contact your Avaya representative or distributor for more information.

Introduction to Disk Mirroring

This chapter introduces you to “disk mirroring,” an optional feature of CentreVu® Call Management System (CMS) that provides you with a completely redundant set of data, helping to ensure data security. To use disk mirroring, you must have a Sun®Enterprise 3000 or Sun Enterprise 3500 platform running CMS Release 3 Version 8 (R3V8).

Contents

“Introduction to Disk Mirroring” contains the following topics:

- [“Disk mirroring defined” on page 1](#)
- [“How CMS implements disk mirroring” on page 5](#)

Disk mirroring defined

“Mirrors” are a feature of the Solstice DiskSuite® software package. They allow you to build a hard disk system containing multiple 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, they are not meant to be a substitute for regular backups. Mirroring creates data redundancy but the data can still become corrupt. *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.

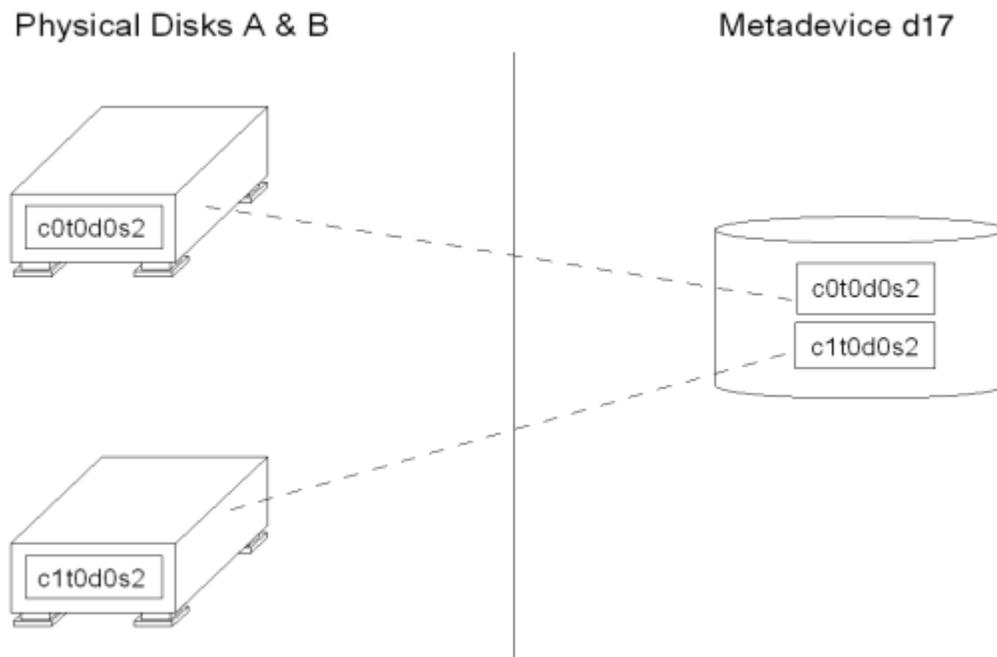
Solstice DiskSuite uses virtual disks to manage physical disks and their associated data. In Solstice DiskSuite, a virtual disk is called a metadevice. To a software application, a metadevice is identical to a physical disk drive. DiskSuite handles all I/O requests directed at a metadevice, converting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

Disk mirroring defined

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

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

The illustration below shows a simple, two-slice metadevice on a mirrored E3500 system.



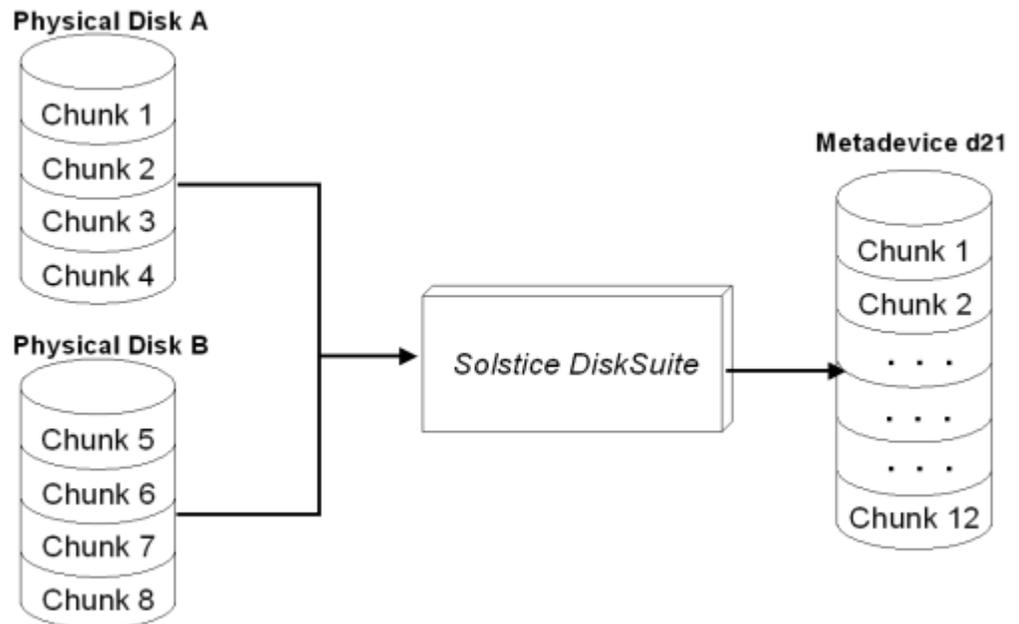
Concatenated metadevices

A metadevice can be configured as any one of three basic types:

- striped
- concatenated
- concatenated stripes.

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



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

State databases

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

State database replicas

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

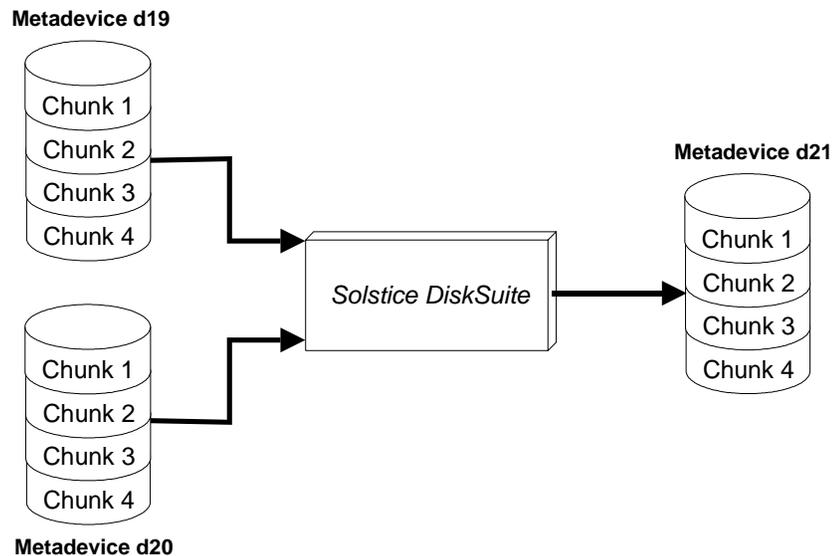
Mirrors

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

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

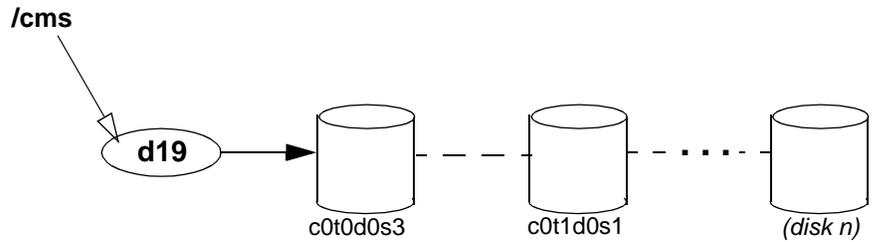
A mirror can be two-way or three-way. A two-way mirror is a mirror with two submirrors; a three-way mirror has three submirrors. *CMS supports only two-way mirrors.*

The illustration below shows a typical two-way mirror, d21, consisting of the submirrors d19 and d20.

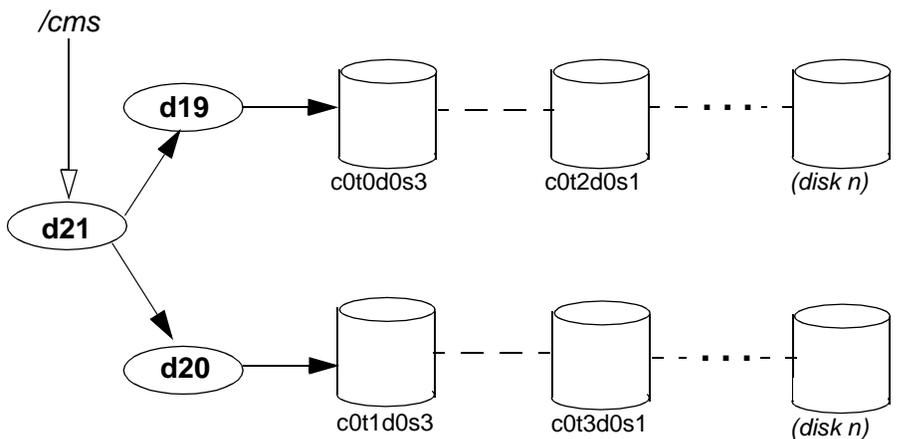


How CMS implements disk mirroring

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

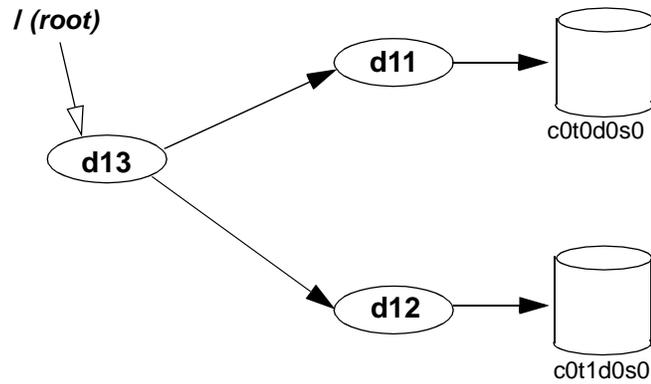


To implement disk mirroring, a metadvice d20 is created as a duplicate of d19, and the two metadvicees are configured as submirrors of d21, as shown in the following example:

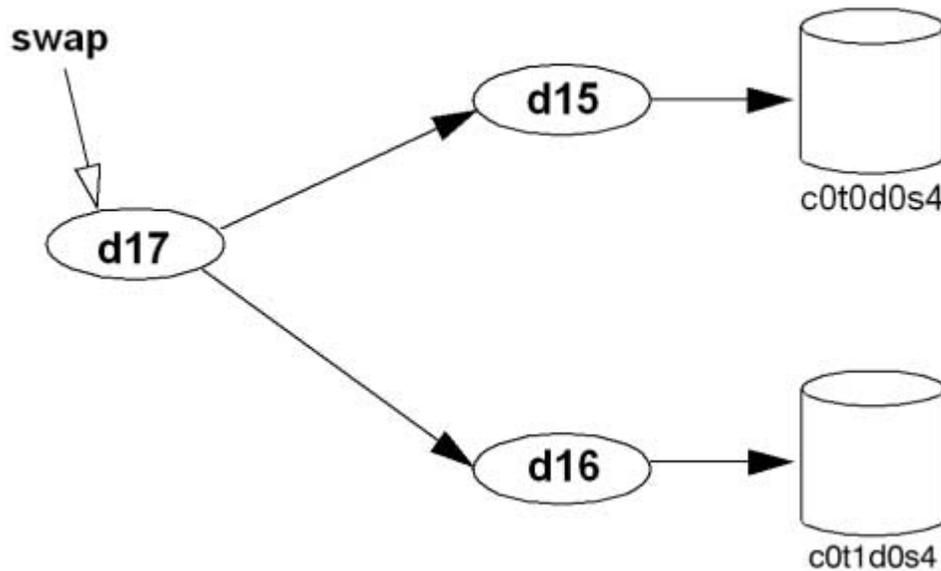


How CMS implements disk mirroring

In a CMS computer system with /cms mirrored, root is also mirrored. The root mirror is d13; its submirrors are d11 and d12. For example, if the system is an E3000:



If your computer system is running CMS R3V8, your swap partition is also mirrored. The swap mirror is d17; its submirrors are d15 and d16. For example:



Mirror layouts

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

Drive Device Name

A hard drive device name can be represented as cXtXdXsX.

- Where cX is the device controller number.
- Where tX is the device target number. The target number is the slot number for the hard drive.
- Where dX is the device disk list number. On CMS systems this is always 0.
- Where sX is the device partition slice number for the hard drive.

Example:

A drive in slot 0 has the device name c0t0d0. The drive in slot 0 is always the primary boot disk.

Mirror layout on a Sun Enterprise 3500 computer

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

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

How CMS implements disk mirroring

Mirror layout on a Sun Enterprise 3000 computer

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

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

Additional references

Sun Microsystems® documentation related to Solstice DiskSuite and disk maintenance and configuration includes:

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

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.

Contents

“Installing Mirrored Systems” contains the following topics:

- [“Hardware Requirements” on page 9](#)
- [“How to upgrade an Enterprise platform to disk mirroring” on page 14](#)
- [“Factory installation procedures for disk mirroring” on page 34](#)
- [“Verifying disk space” on page 44](#)
- [“Setting up the alternate boot device” on page 46](#)

Hardware Requirements

Overview

“Hardware Requirements” describes the drive slot arrangement in your system.

Contents

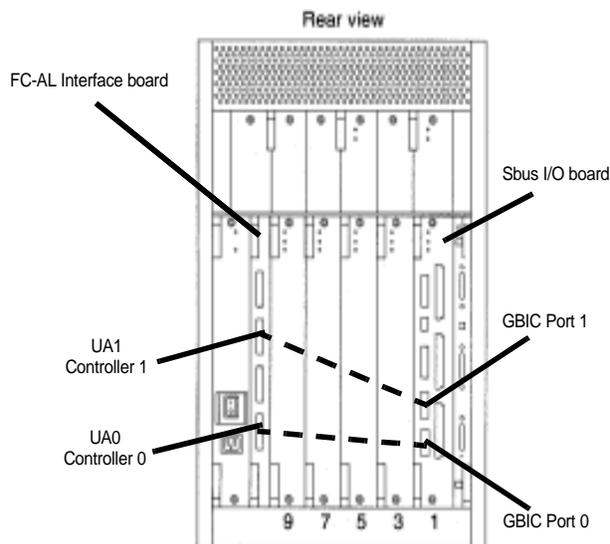
“Hardware Requirements” contains the following topics:

- [“Required additional hardware” on page 10](#)
- [“Required software” on page 11](#)
- [“Hard drive slots” on page 12](#)
- [“Increasing system capacity” on page 13](#)

Required additional hardware

A CMS system must have additional hardware installed in order to function as a mirrored system. You must have the following pieces of hardware in addition to what already came installed on your system:

- For a Sun Enterprise 3500 system, two GigaByte Interface Converter (GBIC) modules. A GBIC is a small hardware insert. One will be installed into the UA slot 0 on the FC-AL Interface board, and the other will be installed into GBIC Port 0 on the first I/O board. See the illustration below.
- For a Sun Enterprise 3500 system, a fiber cable to connect the UA port GBIC to GBIC Port 1 on the I/O board.
- Twice the number of disk drives needed for an unmirrored system. All the pairs of disks must be the same size. For example, slot 0 and slot 4 could be 9 GB disks. Slot 1 and slot 5 could be 18 GB disks.



Required software

The CMS system must have the following software installed in order to function as a mirrored system:

CMS R3V8
Solaris 7 3/99 or later
Solstice DiskSuite 4.2
any CMS R3V8 load

For a list of all required software for your system, see the *CentreVu Call Management System Release 3 Version 8 Software Installation, Maintenance, and Troubleshooting*, 585-210-941 book.

Hard drive slots

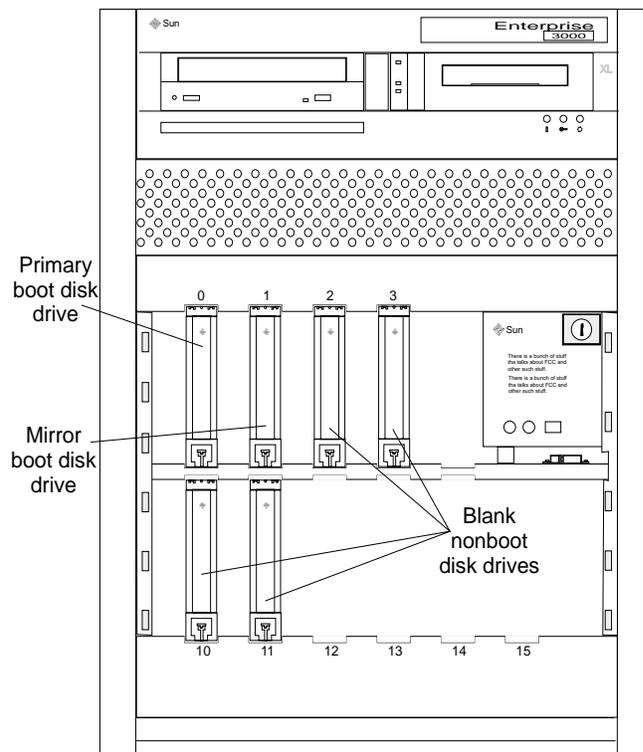
Enterprise 3000 Standard Configuration

In an Enterprise 3000 mirrored system, mirror disks are normally installed in target number order, beginning with the first empty slot.

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

It can be difficult to predict the mirror layout on an upgraded system.

Enterprise 3000 Diagram:

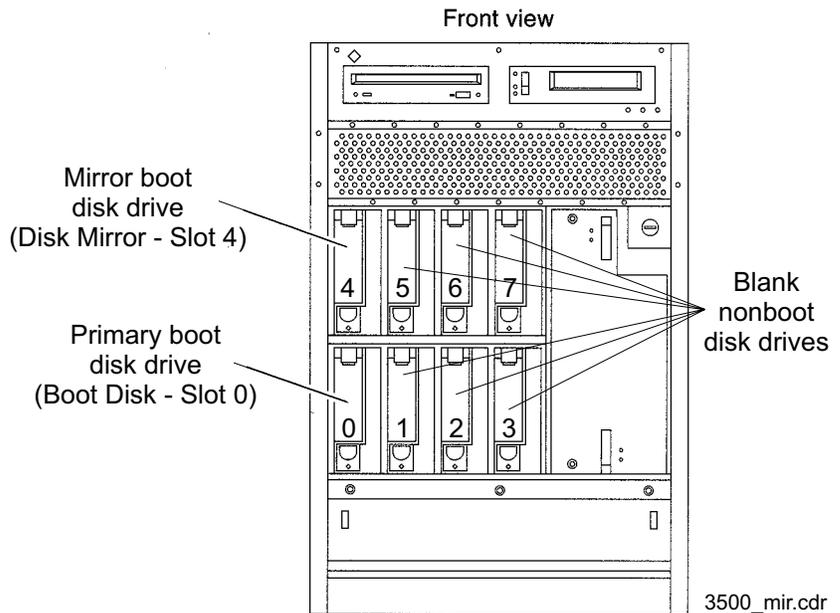


e3000_mirror.cdr

**Enterprise 3500
Standard
Configuration**

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

Enterprise 3500 Diagram:



In an Enterprise 3500 mirrored system, mirror disks always go in the lower bay. If you are mirroring a two-disk system, for example, the two original disks would already be installed in slots 0 and 1 in the lower bay, and you would install the two new disks in slots 4 and 5 in the upper bay.

Increasing system capacity

If you intend to increase your system capacity as well as mirror your system, first install the new disks needed to increase capacity. After the system capacity has been increased, you may then install disk mirroring.

Important!

You must make certain the basic system works before you try to mirror the disks.

For more information about hard drive installation, see the appropriate *CentreVu Call Management System Hardware Installation, Maintenance, and Troubleshooting* book for your hardware platform.

How to upgrade an Enterprise platform to disk mirroring

Overview

This section contains procedures for upgrading a Enterprise 3000 or Enterprise 3500 computer from a nonmirrored system to a mirrored system.

Contents

“How to upgrade an Enterprise platform to disk mirroring” contains the following procedures:

- [“Identifying mirror devices” on page 15](#)
 - [“Installing the mirror disks” on page 17](#)
 - [“Partitioning the mirror disks” on page 20](#)
 - [“Creating the state database replicas” on page 23](#)
 - [“Creating the md.tab table entries” on page 25](#)
 - [“Setting up the mirrors for the root file system” on page 26](#)
 - [“Setting up the mirrors for the swap space” on page 28](#)
 - [“Setting up the /cms mirror” on page 31](#)
 - [“Verify disk space” on page 33](#)
 - [“Create the alternate boot device” on page 33](#)
-

Prerequisites

Before you upgrade the system for mirroring, perform the following tasks:

- Verify that the system meets all of the hardware requirements for mirroring.
- Check the number and capacity of new disks to be installed on the system.
 - There should be as many new disks to be installed as there are old disks.
 - The capacity of the mirror disks should be the same as the old disks.

See [“Hardware Requirements” on page 9](#) for additional information.

Identifying mirror devices

Choose one of the following procedures depending on the hardware platform:

- [“Identifying mirror devices on a E3500 system” on page 15](#)
- [“Identifying mirror devices on a E3000 system” on page 15](#)

Identifying mirror devices on a E3500 system

If you have an Enterprise 3500 system, the device names have already been identified and recorded in the Enterprise 3500 table. See the [“Device/Metadevice Names for an Enterprise 3500 system table”](#).

Device/Metadevice Names for an Enterprise 3500 system table:

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

Identifying mirror devices on a E3000 system

To identify the device names for an Enterprise 3000 system:

1. Find the first empty slot in your system. That slot will hold the secondary boot disk.
2. Record in the [“Device/Metadevice Names for an Enterprise 3000 system table”](#) the following device names for that disk:
 - The slice 0 device name as submirror 12.
 - The slice 3 device name as disk 1 of submirror 20.
 - The slice 4 device name as submirror 16.

How to upgrade an Enterprise platform to disk mirroring

Example:

If the first empty drive slot in the system is 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>
Submirror d16 Device Names		
swap		<u>c0t2d0s4</u>

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

Example: If a system already has two disk drives in slots 0 and 1. You would reserve slot 2 for the second boot disk. The remaining new disk is a nonboot disk and will be installed in slot 3. So you would record the following values in the table:

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

Device/Metadevice Names for an Enterprise 3000 system table:

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

Installing the mirror disks

Choose one of the following procedures depending on the hardware platform:

- [“Installing the mirror disks on an Enterprise 3500 system” on page 17](#)
- [“Installing the mirror disks on an Enterprise 3000 system” on page 20.](#)

Installing the mirror disks on an Enterprise 3500 system

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

1. Enter:

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

The system displays a message similar to the following:

```
lrwxrwxrwx 1 root root 50 Nov 16 08:58
/dev/dsk/c1t6d0s0 -> ../../devices/sbus@3,0
/SUNW,fas@3,8800000/sd@6,0:a
```

2. In the preceding example output, notice that c1t6d0s0 is pointing to sd@6,0:a, which indicates that controller 1 is being used by the CD-ROM drive. To remove that association, enter the following commands:

```
rm /dev/dsk/c1*
```

```
rm /dev/rdisk/c1*
```

```
touch /reconfigure
```

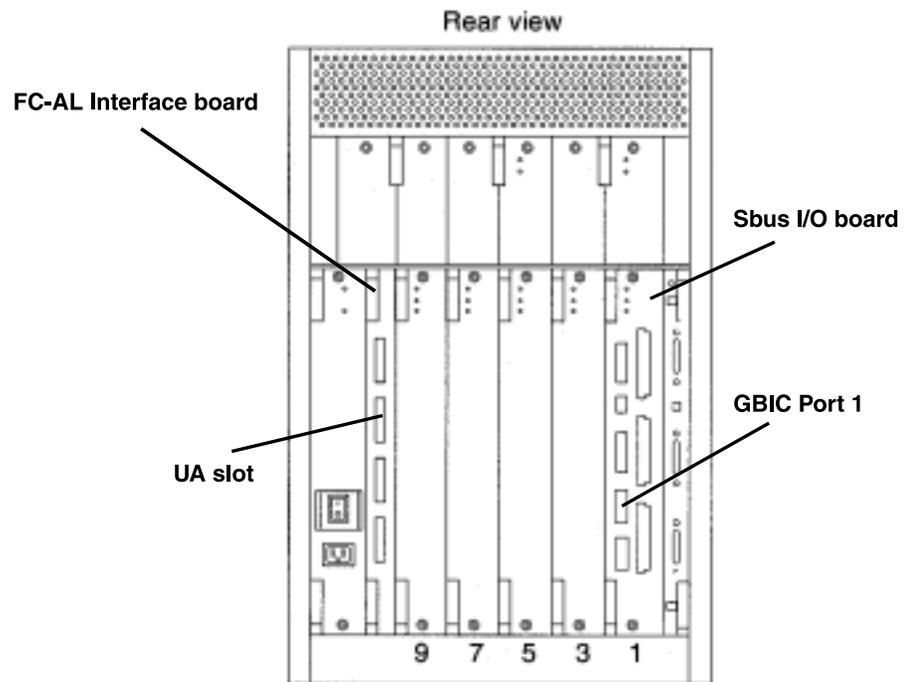
```
/usr/sbin/shutdown -y -g0 -i0
```

How to upgrade an Enterprise platform to disk mirroring

3. Install new hardware components:

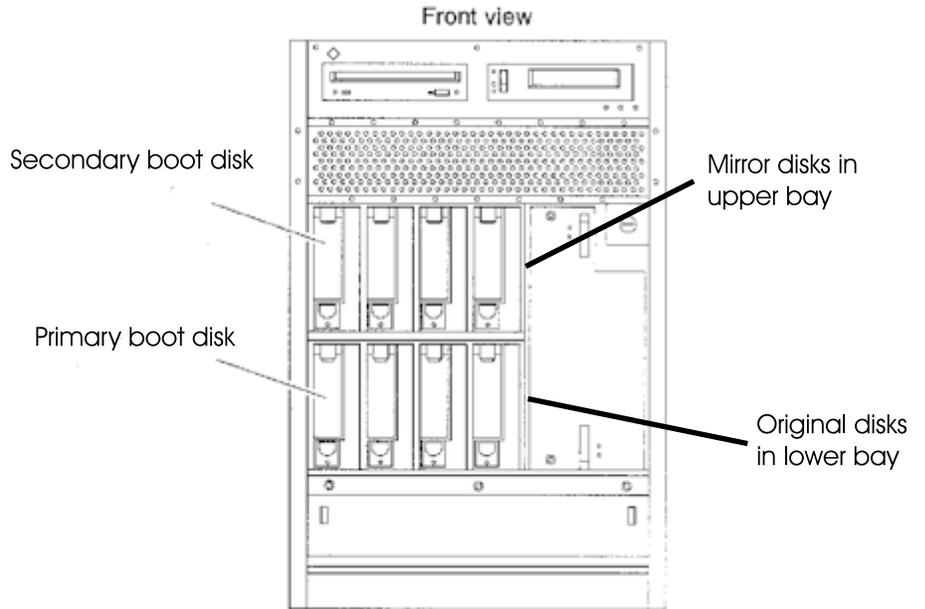
- Install a Gigabit Interface Converter (GBIC) into slot UA0 on the Interface board.
- Install another GBIC into port 0 on the SBus I/O board.
- Install a fiber cable between the two newly installed GBICs.

The illustration below shows where the GBICs are installed. For more detailed information, see *CentreVu Call Management System Sun Enterprise 3500 Computer Hardware Installation, Maintenance and Troubleshooting*, 585-215-873.



How to upgrade an Enterprise platform to disk mirroring

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



5. From the `ok` prompt, enter:

```
boot -r
```

6. Log in as **root**.

7. Enter:

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

The system displays a message similar to the following:

```
lrwxrwxrwx 1 root root 50 Nov 16 08:58
/dev/dsk/c2t6d0s0 -> ../../devices/sbus@3,0
/SUNW,fas@3,8800000/sd@6,0:a
```

In the example output, `c2t6d0s0` is pointing to `sd@6,0:a`, which indicates that the CD-ROM drive is now on controller 2.

How to upgrade an Enterprise platform to disk mirroring

Installing the mirror disks on an Enterprise 3000 system

Install the new disk drives, which will form the mirror, into the empty slots reserved in [“Identifying mirror devices” on page 15](#). Follow the procedures for adding a hard disk drive in the “Enterprise 3000 Maintenance” section of the *CentreVu Call Management System Sun Enterprise 3000 and SPARCserver Computers Hardware Maintenance and Troubleshooting*, 585-214-016 book.

Partitioning the mirror disks

1. Enter:

```
format
```

The system displays a menu of disks, and prompts you to select one.

2. Enter the menu number of the new disk to be partitioned.

The system displays the `format>` prompt.

NOTE:

If you have an Enterprise 3500, make certain that the new drives are named using “c1,” as `c1t4d0` or `c1t5d0`.

3. Enter:

```
partition
```

The system displays the `partition>` prompt.

4. Enter:

```
print
```

The system displays the partition table and the `partition>` prompt.

5. Compare the display to the [“Disk partitioning table” on page 21](#) to determine if you need to re-partition the disk.

If you do not need to re-partition the disk continue with [“Creating the state database replicas” on page 23](#). You will need to re-partition the disk in most scenarios.

How to upgrade an Enterprise platform to disk mirroring

6. Enter the number of the partition to be reconfigured.
The system displays a series of prompts for partitioning information.
7. Answer the prompts according to the CMS version and disk type you are partitioning. See the [“Disk partitioning table”](#) for partition information.
8. Repeat Steps 6 through 7 for each partition on the disk.

Disk partitioning table:

CMS R3V8 Boot Disk Partition Values:

Partition #		0	1	2 ^a	3	4
Flag		wm	wm	wm	wm	wu
Tag		boot disk 1: / or root <i>boot disk 2:</i> unassigned	unassigned	overlap or backup	unassigned	swap
4.2-GB Disk	Starting Cylinder	0	1023	0	1030	2909
	Size (cyl.)	1023	7	3880	1879	971
9.1-GB Disk SCSI (E3000) and FCAL (E3500)	Starting Cylinder	0	616	0	623	4339
	Size (cyl.)	616	7	4924	3716	585
18-GB Disk SCSI (E3000) and FCAL (E3500)	Starting Cylinder	0	471	0	478	7070
	Size (cyl.)	471	7	7506	6592	436

^a The default size of the overlap file system is always the size of the entire disk. Occasionally, the name *backup* will appear instead of *overlap*. Do not change the slice 2 value or name.

CMS R3V8 Nonboot Disk Partition Values:

Partition #	0	1	2 ^a
Flag	wm	wm	wm

How to upgrade an Enterprise platform to disk mirroring

Partition #		0	1	2 ^a
Tag		unassigned	unassigned	overlap or backup
4.2-GB Disk	Starting Cylinder	0	2	0
	Size (cyl.)	2	3878	3880
9.1-GB Disk SCSI (E3000) and FCAL (E3500)	Starting Cylinder	0	2	0
	Size (cyl.)	2	4922	4924
18-GB Disk SCSI (E3000) and FCAL (E3500)	Starting Cylinder	0	2	0
	Size (cyl.)	2	7504	7506
36-GB Disk FCAL (E3500)	Starting Cylinder	0	2	0
	Size (cyl.)	2	24618	24620

^a The default size of the overlap file system is always the size of the entire disk. Occasionally, the name *backup* will appear instead of *overlap*. Do not change the slice 2 value or name.

: 9. Enter:

```
print
```

The system displays the partition table and the `partition>` prompt.

10. Compare the display to the “[Disk partitioning table](#)” to make certain you have repartitioned the disk properly. If not, repeat Step 6.

11. Enter:

```
label
```

The system displays the following message:

```
Ready to label disk, continue?
```

12. Enter: `y`

The system displays the `partition>` prompt.

13. Enter: `q`

The system displays the `format>` prompt.

14. If there is another disk to be partitioned, perform the following procedure:

a. Enter:

```
disk
```

b. Repeat Steps 2 through 14.

15. After all of the new disks have been partitioned, enter: `q`

Creating the state database replicas

1. Enter the following commands:

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

2. Enter:

```
metadb -i
```

How to upgrade an Enterprise platform to disk mirroring

The system displays all of the database replicas in a message similar to the following:

```
      flags          first blk      block count
M    p  luo         16             1034        /dev/dsk/c0t0d0s1
M    p  luo        1050            1034        /dev/dsk/c0t0d0s1
M    p  luo        2084            1034        /dev/dsk/c0t0d0s1
M    p  luo        3118            1034        /dev/dsk/c0t0d0s1
.....
.....
.....
M    p  luo         16             1034        /dev/dsk/c0t12d0s0
M    p  luo        1050            1034        /dev/dsk/c0t12d0s0
M    p  luo        2084            1034        /dev/dsk/c0t12d0s0
o - replica active prior to last mddb configuration change
u - replica is up to date
l - locator for this replica was read successfully
c - replica's location was in /etc/opt/SUNWmd/mddb.cf
p - replica's location was patched in kernel
m - replica is master, this is replica selected as input
W - replica has device write errors
a - replica is active, commits are occurring to this replica
M - replica had problem with master blocks
D - replica had problem with data blocks
F - replica had format problems
S - replica is too small to hold current data base
R - replica had device read errors
```

3. Remove all of the existing state database replicas using the following command:

```
metadb -d -f /dev/dsk/disk_slice
```

where ***disk_slice*** is the physical device partition number containing the damaged database replica.

Example:

Using the output obtained in the previous step, the command would be entered as:

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

4. Enter:

```
metadb -a -c4 -f /dev/rdisk/device_location
```

where ***device_location*** is the boot device location on which the state database replicas will be created.

Example:

c0t0d0s1 on an Enterprise 3000 or 3500 system

5. For the alternate boot disk and any other non-boot disks on the system, enter:

```
metadb -a -c3 -f /dev/rdisk/device_location
```

Where ***device_location*** is the alternate boot device location.

6. Enter:

```
metadb -i
```

The system displays the state database replicas currently on the system.

7. Verify that all state database replicas have been created.

WARNING:

If the CMS mirrored system is installed only with a boot disk and an alternate boot disk and has no additional non-boot disks, the system will not automatically reboot after a boot disk failure that causes a subsequent system shutdown. If this does occur, see [“Recovering a mirrored system after disk failure”](#) on page 65.

Creating the md.tab table entries

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

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

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

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

The lines you add must be in the following format:

```
dx y 1 device
```

How to upgrade an Enterprise platform to disk mirroring

- **x** is the metadvice name.
- **y** denotes the total number of disks named in the line.
- **1 device** indicates one disk with the device name **device**. You should have recorded the device names in “[Identifying mirror devices](#)” on page 15.

The pattern **1 (space) device** repeats until all disks in the metadvice have been named. The **d12** entry must specify one disk; the **d20** entry must specify the same number of disks as the existing **d19** entry. If your system has two disks in slots 0 and 1 and you are adding a two-disk mirror in slots 2 and 3.

Example:

You would add the *md.tab* lines shown in boldface in the following example:

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

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

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

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

Setting up the mirrors for the root file system

This procedure sets up the mirrors for the *root* file system. Metadevices **d11** and **d13** are first initialized. Then the system is rebooted to test **d13**. After **d13** functionality is verified by rebooting the system, **d12** is also initialized and set up.

1. Enter:

```
metainit -f d11
```

The system activates metadvice **d11** and displays the following

How to upgrade an Enterprise platform to disk mirroring

message:

```
d11: Concat/stripes is set up
```

2. Enter the following commands:

```
metainit d12
```

```
metainit d13
```

```
metaroot d13
```

3. Enter:

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

The system reboots.

4. Log in as **root**.

5. Enter:

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

6. Enter:

```
metattach d13 d12
```

The system attaches d12 to d13 and displays the following message:

```
d13: submirror d12 is attached
```

7. Enter:

```
metastat d13
```

The system verifies that root is set up and resynching, and then displays the following message:

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

Check the "State:" lines in the output. If a status other than Okay or Resynching is displayed, a problem with the setup is indicated.

How to upgrade an Enterprise platform to disk mirroring

Discontinue the upgrade and telephone the Avaya National Customer Care Center (1.800.242.2121) or contact your Avaya representative or distributor.

Setting up the mirrors for the swap space

To set up the mirrors for the swap space:

1. Enter:

```
metainit -f d15
```

The system activates metadvice d15 and displays the following message:

```
d15: Concat/strip is set up
```

2. Enter the following commands:

```
metainit d16
```

```
metainit d17
```

3. Edit the `/etc/vfstab` file.

The relevant portion of the file looks like this:

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

Perform the following steps to modify the file:

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

How to upgrade an Enterprise platform to disk mirroring

the secondary boot disk. Comment out that line by inserting a pound sign at the beginning of the line.

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

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

e. Save the file and quit the editor.

4. Deactivate the current swap space by entering the following command:

```
swap -d /dev/dsk/c0t0d0s4
```

5. Enter:

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

The system sets up swap space.

6. Verify that swap has been set up correctly by entering the following command:

```
swap -l
```

7. Enter:

```
metastat | grep sync
```

8. Choose one of the following steps:

- If a syncing message is displayed, repeat step 7 after a few minutes.
- If *only* the command prompt is returned, then continue with step 9.

WARNING:

Do not reboot the system if any mirrors are syncing.

9. Enter:

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

The system reboots.

How to upgrade an Enterprise platform to disk mirroring

10. Log into the system as **root**.

11. Enter:

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

12. Enter:

```
metattach d17 d16
```

The system attaches d16 to d17 and displays the following message:

```
d17: submirror d16 is attached
```

13. Enter:

```
metastat d17
```

The system verifies that swap is set up and resynching, and displays the following message:

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

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

Setting up the /cms mirror

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

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  
-----
```

2. Save the **/etc/vfstab** file and exit the editor.

3. Enter:

```
metainit d20
```

The system activates metadvice d20 and displays the following message:

```
d20: Concat/stripes is set up
```

4. Enter:

```
metainit -f d21
```

The system initializes the new mirror d21 and displays the following message:

```
d21: Mirror is set up
```

5. Enter:

```
metattach d21 d20
```

The system attaches metadvice d20 to the mirror d21, and displays the following message:

```
d21: submirror d20 is attached
```

How to upgrade an Enterprise platform to disk mirroring

6. Enter:

```
metastat d21
```

The system verifies that d21 is active and displays the following message:

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

7. Enter:

```
metastat | grep sync
```

8. Choose one of the following steps:

- If a syncing message is displayed, repeat step 7 after a few minutes.
- If *only* the command prompt is returned, then continue with step 9.

WARNING:

Do not reboot the system if any mirrors are syncing.

9. Enter:

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

The system reboots.

10. Log into the system as **root**.

11. Enter:

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

Verify disk space

- Verify the existing disk space.
See [“Verifying disk space” on page 44](#) for more information.
-

Create the alternate boot device

1. Create the alternate boot device.
See [“Setting up the alternate boot device” on page 46](#) for more information.
2. Perform a CMSADM backup.
For more information see, *CentreVu CMS R3V8 Administration*, 585-210-910.

Factory installation procedures for disk mirroring

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

Contents

“Factory installation procedures for disk mirroring” contains the following procedures:

- [“Installing the hardware” on page 34](#)
 - [“Installing the operating system” on page 35](#)
 - [“Installing the application software” on page 37](#)
 - [“Checking the disk.type file” on page 37](#)
 - [“Setting up the disk subsystem for Solstice DiskSuite” on page 39](#)
 - [“Running the setup scripts” on page 41](#)
 - [“Verify disk space” on page 43](#)
 - [“Install CMS” on page 43](#)
 - [“Create the alternate boot device” on page 43](#)
-

Installing the hardware

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

Installing the operating system

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

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

- Use the “Disks table” on page 35 and the “Disk partitioning tables” on page 36 to determine disk partitioning details. The “Disks table” shows which disks are boot disks and which are not; the “Disk partitioning tables” shows partitioning details.

Disks table:

Disk	Enterprise 3000 Device Names	Enterprise 3500 Device Names
boot disk	c0t0d0	c0t0d0
alternate boot disk	c0t1d0 ^a	c1t4d0
other disks	c0t2d0 c0t3d0 c0t10d0 c0t11d0 c0t12d0 c0t13d0 c0t14d0 c0t15d0	c0t1d0 c0t2d0 c0t3d0 c1t5d0 c1t6d0 c1d7d0

^a True for factory-installed disk mirroring. Not necessarily true for systems upgraded to disk mirroring in the field.

Factory installation procedures for disk mirroring

Disk partitioning tables

CMS R3V8 Boot Disk Partition Values:

Partition #		0	1	2 ^a	3	4
Flag		wm	wm	wm	wm	wu
Tag		boot disk 1: / or root boot disk 2: unassigned	unassigned	overlap or backup	unassigned	swap
4.2-GB Disk	Starting Cylinder	0	1023	0	1030	2909
	Size (cyl.)	1023	7	3880	1879	971
9.1-GB Disk SCSI (E3000) and FCAL (E3500)	Starting Cylinder	0	616	0	623	4339
	Size (cyl.)	616	7	4924	3716	585
18-GB Disk SCSI (E3000) and FCAL (E3500)	Starting Cylinder	0	471	0	478	7070
	Size (cyl.)	471	7	7506	6592	436

^a The default size of the overlap file system is always the size of the entire disk. Occasionally, the name *backup* will appear instead of *overlap*. Do not change the slice 2 value or name.

CMS R3V8 Nonboot Disk Partition Values:

Partition #		0	1	2 ^a
Flag		wm	wm	wm
Tag		unassigned	unassigned	overlap or backup
4.2-GB Disk	Starting Cylinder	0	2	0
	Size (cyl.)	2	3878	3880

Partition #		0	1	2 ^a
9.1-GB Disk SCSI (E3000) and FCAL (E3500)	Starting Cylinder	0	2	0
	Size (cyl.)	2	4922	4924
18-GB Disk SCSI (E3000) and FCAL (E3500)	Starting Cylinder	0	2	0
	Size (cyl.)	2	7504	7506
36-GB Disk FCAL (E3500)	Starting Cylinder	0	2	0
	Size (cyl.)	2	24618	24620

^a The default size of the overlap file system is always the size of the entire disk. Occasionally, the name *backup* will appear instead of *overlap*. Do not change the slice 2 value or name.

: Installing the application software

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

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

Checking the disk.type file

The `olds` commands, are used to configure the Solstice DiskSuite software, and obtain information about your disk configuration from the `/olds/disk.type` file. If your boot disks are not as specified in the “Disks table” on page 35, check the `disk.type` file to verify that it names the proper boot disks for your system.

Factory installation procedures for disk mirroring

The relevant portion of the **disk.type** file looks similar to the following example:

```
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.

Setting up the disk subsystem for Solstice DiskSuite

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

```
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. On some systems **Ctrl+H** will not work. If this is the case, substitute `^H` for **Ctrl+H**.

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

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

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

4. Enter:

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

5. Check the disks by entering:

```
olds -mirrored -check_disks
```

The system checks the disks and displays the following message:

```
number of external scsi controllers with disks is = 0
number of disks is = 4
. . .
disk:c0t0d0 is partitioned ok
disk:c0t1d0 is partitioned ok
disk:c0t2d0 is partitioned ok
disk:c0t3d0 is partitioned ok
Success, checking disks for mirroring.
```

Factory installation procedures for disk mirroring

6. Enter:

```
olds -mirrored -mk_files
```

The system creates the system files for the Solstice DiskSuite software.

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

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.

Example:

The following example shows a total of four disk drives on the system, two in d19 and two in d20:

```
.  
. .  
#metaroot  
d11 1 1 c0t0d0s0  
d12 1 1 c0t2d0s0  
d13 -m d11  
  
#/cms  
d19 2 1 /dev/rdisk/c0t0d0s3 1 /dev/rdisk/c0t1d0s1  
d20 2 1 /dev/rdisk/c0t2d0s3 1 /dev/rdisk/c0t3d0s1  
d21 -m d19  
  
#metaswap  
d15 1 1 c0t0d0s4  
d16 1 1 c0t2d0s4  
d17 -m d15
```

If there is a discrepancy between what the system should have recognized and what it *did* recognize, complete the procedure in [“The system fails to recognize all disk drives” on page 121](#), and then return to Step 5.

Running the setup scripts

1. Enter:

```
olds -mirrored -metadbs
```

The system creates the state database replicas and displays the following message:

```
.
.
.
Success, setting up metadb replicas
```

2. Enter:

```
olds -mirrored -setroot
```

The system mirrors the root file system and displays the following message:

```
.
.
.
Success, root mirrored successfully.
```

3. Enter: **metastat**

Wait for all disks to be synched.

```
. . .
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
```

Factory installation procedures for disk mirroring

4. Enter:

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

The system displays the `ok` prompt.

⇒ NOTE:

Ignore any FSCK error messages which may be generated.

5. At the `ok` prompt, enter:

```
boot -r
```

6. Log in as **root**.

7. Setup the `/cms` metadevices by entering the following commands:

```
ksh -o vi
```

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

```
olds -mirrored -setup
```

The system displays the following message:

```
.  
. .  
3962848, 13997552, 14032256, 14066960, 14101664,  
14136368, 14171072,  
14205776, 14240480, 14275184, 14309888, 14344592,  
14376992, 14411696,  
14446400, 14481104, 14515808, 14550512, 14585216,  
14619920, 14654624,  
14689328, 14724032, 14758736, 14793440, 14828144,  
14862848, 14897552,  
14929952,  
ufs fsck: sanity check: /dev/md/rdisk/d21 okay  
Success, /cms mirrored successfully
```

8. Create and mount the `/cms` file system by entering the following commands:

```
mkdir /cms
```

```
mount /cms
```

Verify disk space

- Verify the existing disk space.
See [“Verifying disk space” on page 44](#) for more information.
-

Install CMS

Perform the following steps to install CMS:

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

Detailed installation instructions can be found in the *CentreVu CMS R3V8 Software Installation, Maintenance, and Troubleshooting*, 585-210-941 book.

Create the alternate boot device

1. Create the alternate boot device.
See [“Setting up the alternate boot device” on page 46](#) for more information.
2. Perform a CMSADM backup.
For more information see, *CentreVu CMS R3V8 Administration*, 585-210-910.

Verifying disk space

1. Enter: `df -k`

The system displays a file system table with the available disk space.:

Filesystem	kbytes	used	avail	capacity	M
ounted on					
/dev/md/dsk/d13	1039555	355366	580239	38%	/
/proc	0	0	0	0%	/
proc					
fd	0	0	0	0%	/

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

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

2. Enter:

`metastat`

The system displays the following message:

```
. . .
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
```

3. You can check connectivity by checking the output of the `metastat` command to see that all the disks are recognized

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 Avaya National Customer Care Center at 1.800.242.2121, or contact your Avaya representative or distributor.

Setting up the alternate boot device

When you set up your mirrored system, you create an alternate boot device named *bootdevice2*. The alternate boot device is a fail-safe device.

“Setting up the alternate boot device” includes the following procedures:

- [“Obtaining the correct disk device definition” on page 46](#)
- [“Creating the alternate boot device” on page 47](#)

Obtaining the correct disk device definition

Before the alternate boot disk can be created, you must determine the correct device definition for the disk to be used as the alternate boot. The “Default device definitions table” provides default device descriptions for the Enterprise 3000, 3500, 3501 and 3503 platforms.

Default device definitions table:

System	Device Definition
E3000	<code>sbus@3,0/SUNW,fas@3,8800000/sd@1,0</code>
E3500	<code>sbus@2,0/SUNW,socal@d,10000/sf@1,0/ssd@4,0</code>
E3501/3503	<code>sbus@3,0/SUNW,socal@d,10000/sf@1,0/ssd@4,0</code>

Disk device definitions for Enterprise 3500 and 3501/3503 systems

1. For Enterprise 3500, 3501 or 3503 systems, **always** use the appropriate device definition for your 3500 platform version, as shown in the [“Default device definitions table”](#).

2. If you are uncertain of the platform version, enter:

```
/usr/sbin/prtconf | grep ffb
```

- If the platform is a 3501 or 3503, the following output will be displayed:

```
SUNW,ffb, instance #0
```

- If the platform is a 3500, no output will be returned.

Disk device definitions for Enterprise 3000 systems

On an Enterprise 3000 system, the device definition will vary from the one shown in the [“Default device definitions table”](#) if a disk device other than `c0t1d0s0` is used as the alternate boot.

1. If a device other than c0t1d0s0 is selected, enter:

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

The system displays a message similar to the following:

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

2. Record the device definition (the part after “/devices/” and before “:a,raw”). For the example above, you would record the part beginning “sbus@3,0” and ending immediately after “sd@2,0” to excerpt the following string:

```
sbus@3,0/SUNW,fas@3,8800000/sd@2,0
```

This information is used in “Creating the alternate boot device”.

Creating the alternate boot device

1. Enter:

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

The system displays the ok prompt.

2. At the ok prompt, enter:

```
show-disks
```

The system displays a list of disk drives similar to the following example:

```
a) /sbus@3,0/SUNW,fas@3,8800000/sd  
b) /sbus@3,0/SUNW,socal@d,10000/sf@1,0/ssd  
c) /sbus@3,0/SUNW,socal@d,10000/sf@0,0/ssd  
d) NO SELECTION  
Enter selection, q to quit:
```

Setting up the alternate boot device

3. Select the letter that corresponds with the alternate boot disk drive. If you are unsure which disk is the alternate boot drive, see the [“Default device definitions table” on page 46](#) for more information.

The system saves the device string and in the edit clipboard, and displays a message similar to the following:

```
/sbus@3,0/SUNW,socal@d,10000/sf@1,0/ssd has been selected
Type ^Y (Control-Y) to insert it in the command line e.g. ok
nvalias mydev ^Y
      for creating devalias mydev for
/sbus@3,0/SUNW,socal@d,10000/sf@1,0/ssd
```

4. At the `ok` prompt, enter:

```
nvalias bootdevice2 Ctrl+Y
```

The system displays the disk alias saved in the clipboard and the cursor is at the end of the line.

```
/sbus@3,0/SUNW,socal@d,10000/sf@1,0/ssd
```

5. At the end of the line, enter: `@1,0`

6. Press **Enter**.

The system displays the `ok` prompt.

7. At the `ok` prompt, enter:

```
devalias
```

The output should include a line that is similar to the following example:

```
.....
.....
.....
Bootdevice2/sbus@3,0/SUNW,socal@d,10000/sf@1,0/ssd@1,0
```

8. Enter one of the following commands, depending upon whether you are at the boot prompt or in root:

- `setenv boot-device disk bootdevice2`
- `eeprom boot-device="disk bootdevice2"`

⇒ NOTE:

For more information on how to enable the alternate boot device on a E3501 or E3503 system, see "Resetting a Device Alias" in Chapter 4 of *CenterVu Call Management System Sun Enterprise 3500 Computer Hardware Installation Maintenance and Troubleshooting*, 585-215-873.

9. To test the alternate boot device, enter:

```
boot bootdevice2
```

When the computer restarts, login as **root** at the console login.

10. Reboot once again to return system control to the regular boot disk:

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

11. To set up a "cron job" for chkDisk, enter:

```
crontab -e
```

The **crontab** file is displayed in editor mode.

12. Look for the following line at the end of the **crontab** file:

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

13. Choose one of the following actions:

- If the line is not found, add it to the end of the file.
- If the line is present, verify that it is not commented out. If the line is commented out uncomment it.

14. Save and quit the file by pressing **Esc** and entering:

```
:wq!
```

15. Enter:

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

Setting up the alternate boot device

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.

Contents

“Maintaining Mirrored Systems” contains the following procedures:

- [“How to tell whether your mirrored system is standard” on page 52](#)
- [“Performing a CMSADM backup on a mirrored system” on page 53](#)
- [“Adding new disks to a standard mirrored system” on page 55](#)
- [“Adding new disks to a nonstandard mirrored system” on page 59](#)
- [“Identifying a faulty disk” on page 62](#)
- [“Recovering a mirrored system after disk failure” on page 65](#)
- [“Partitioning disks with the format command” on page 78](#)
- [“Restoring Data from a CMSADM Backup” on page 82](#)
- [“Repairing or rebuilding the /cms file system” on page 84](#)
- [“Restoring specific files from a CMSADM backup” on page 89](#)
- [“Restoring a mirrored system from a CMSADM backup” on page 90](#)
- [“State Database Replicas” on page 114](#)
- [“Maintaining the chkDisks crontab” on page 115](#)
- [“Examples of valid system files” on page 117](#)

How to tell whether your mirrored system is standard

Use this procedure to determine if you have a “standard” or “nonstandard” configuration.

1. Perform one of the following procedures:
 - Enter: `metastat d11 d12`
 - Edit the `/etc/opt/SUNWmd/md.tab` file.
2. Examine the command output or the contents of the file. You are looking for the d11 and d12 metadevices and their constituent disk drives.
3. Compare the d11 and d12 metadevices and their constituent disk drives to the [“Root metadevice device name table”](#) below.

Standard mirroring—mirroring that was installed at the factory or upgraded to a standard factory configuration will have the structure indicated in the “Root metadevice device name table” .

Root metadevice device name table:

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

Performing a CMSADM backup on a mirrored system

Overview

This section describes how to perform a 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. Mirroring creates data redundancy but the data can still become corrupt. Mirrored systems must be backed up just as often as unmirrored systems.

Contents

“Performing a CMSADM backup on a mirrored system” contains the following procedures:

- [“Printing the vfstab file” on page 53](#)
 - [“Running the backup” on page 53](#)
 - [“Labeling and storing the tapes” on page 54](#)
-

Printing 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:

1. Enter:

```
lp /etc/vfstab
```

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

Running 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: **who -r**

The system displays the run level and date.

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

3. If the system is not in a run level of 2 or 3, enter:

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

Performing a CMSADM backup on a mirrored system

4. Log in as **root**.
5. Enter the number for the `backup` option.
The system prompts for the tape drive.
6. Enter the appropriate option number for the tape drive on your system.
The system begins calculating the approximate number of tapes required and informs you of its calculation.

```
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:
```

7. Insert the tape, and press **Enter**.

NOTE:

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

The system begins the backup.

8. 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, the system will verify each tape. If you used more than one tape, you will have to reinsert each tape.

Labeling and storing the tapes

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

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

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

Adding new disks to a standard mirrored system

Overview

This procedure explains how to add 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. Disks are always added in pairs. Each disk in a pair must be the same size.

If your boot disk assignments differ from the standard factory arrangement shown in the [“Boot-disk assignments for factory installed mirroring”](#) table, you must add the disks using the procedure in [“Adding new disks to a nonstandard mirrored system”](#) on page 59.

Boot-disk assignments for factory installed mirroring:

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

Procedure

⇒ NOTE:

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

1. Enter:

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

The system displays the status of each metadevice.

2. Examine the output for the status of each metadevice. All metadevices must have a status of `okay` before you may proceed. Do not attempt this procedure while any metadevice is resynching or needs maintenance.
3. Do a CMSADM backup. For more information, see [“Performing a CMSADM backup on a mirrored system”](#) on page 53.
4. Choose one of the following commands:
 - `cmsadm`
 - `cmssvc`

The system displays the appropriate menu.

Adding new disks to a standard mirrored system

5. Select the `cms_run` option.
6. Select the `Turn off cms` option.
7. Enter:

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

The system reboots.

8. Install the new disk drives.

For more information about installation procedures, see the appropriate *Hardware Installation, Maintenance, and Troubleshooting* book for your platform.

9. As the console shows that the system is booting, press the **Stop + A** keys.

The system displays the `ok` prompt.

10. Enter: `boot -r`

The system boots into multi-user mode.

11. Log on as **root**.

12. Partition the new disks by following the instructions in [“Partitioning disks with the format command” on page 78](#).

13. Check disk space by entering: `df -k`

The system displays the following message:

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

14. Record the *avail* figure for `/cms` here: _____

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

```
cd /olds
```

```
olds -mirrored -mk_files newdisk1 newdisk2
```

Where, **newdisk1** and **newdisk2** are replaced with the device names of the disk drives added.

Example:

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

16. Load the file `/etc/opt/SUNWmd/md.tab.new` into an editor and verify that it names all of the disks in the system. This includes the new disks you just added.

Adding new disks to a standard mirrored system

For an example of the md.tab file format, see [“Examples of valid system files” on page 117](#).

17. Enter the following command on a single line at the command prompt:

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

The system saves a copy of the existing md.tab file.

18. Enter the following command on a single line at the command prompt:

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

The system replaces the existing md.tab file.

19. Enter:

```
olds -mirrored -setup newdisk1 newdisk2
```

Where, **newdisk1** and **newdisk2** are replaced with the device names of the disks you have added.

The system adds the new disks to the existing metadevices and displays a message similar to the following:

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

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

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

20. Choose one of the following commands:

- **cmsadm**
- **cmssvc**

The system displays the appropriate menu.

21. Select the `cms_run` option.
22. Select the `start cms` option.

Adding new disks to a standard mirrored system

23. Check the available disk space by entering: `df -k`

The system displays the following message.

```
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
```

24. Compare the `avail` figure for `/cms` to the figure you recorded in Step 13. The new figure should show an increase roughly equal to one-half of the total capacity of the drives added.

Adding new disks to a nonstandard mirrored system

Overview

In a “nonstandard” system, the assigned boot disks are not the same as those assigned during a factory installation of disk mirroring. A “nonstandard” disk drive addition cannot take advantage of advanced features of the olds setup scripts.

For more information about determining if your system is standard or nonstandard, see [“How to tell whether your mirrored system is standard” on page 52](#).

Procedure

To add new disks to a nonstandard mirrored system:

1. Choose one of the following commands:

- `cmsadm`
- `cmssvc`

The system displays the appropriate menu.

2. Select the `cms_run` option.
3. Select the `Turn off cms` option.
4. Enter:

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

The system reboots.

5. Install the new disk drives in the appropriate drive slots.

For more information about installation procedures, see the appropriate Hardware Installation, Maintenance, and Troubleshooting book for your platform.

6. As the console shows that the system is booting, press the **Stop + A** keys.

The system displays the `ok` prompt.

7. Enter: `boot -r`

The system boots into multi-user mode.

8. Log on as **root**.

Adding new disks to a nonstandard mirrored system

9. Enter:

```
df -k /cms
```

The system displays file system information for /cms.

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

10. Look for the “avail” figure for /cms and record that figure here:

/cms “avail” figure: _____

11. Partition the new disks. Follow the instructions in [“Partitioning disks with the format command” on page 78](#).

12. Add the olds directories to your path by entering the following commands:

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

```
export PATH
```

13. Create new versions of the setup files by entering:

```
olds -mirrored -mk_files newdisk1 newdisk2
```

Where, **newdisk1** and **newdisk2** are replaced with the device names of the disk drives added.

Example:

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

14. Load the file **/olds/md.tab.new** into an editor and check to make certain it names all of your disks. This includes the ones you just added.

For an example of the format of an md.tab file, see [“Examples of valid system files” on page 117](#).

15. Enter the following command on a single line at the command prompt:

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

The system saves a copy of the existing md.tab file.

16. Enter the following command on a single line at the command prompt:

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

The system replaces the existing md.tab file with the one you created in Step 13.

17. Attach one of the newly installed disk drives to each of the d21 submirrors (d19 and d20) by entering the following commands:

```
metattach d19 newdisk1
metattach d20 newdisk2
```

Where, *newdisk1* and *newdisk2* are replaced with the device names of the disk drives added.

Example:

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

18. Grow the /cms file system by entering:

```
growfs -M /cms /dev/md/rdsk/d21
```

⇒ NOTE:

If the operation should fail, do the following:

- a. Restore the original md.tab file by entering the following command on a single line at the command prompt:

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

- b. Look up the error message in Chapter 4, [“Error messages” on page 128](#).

19. Enter:

```
metastat
```

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

21. Enter:

```
df -k /cms
```

The system displays file system information for /cms.

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

22. Check the response. Look for the “avail” figure for /cms. That figure should be roughly twice the figure recorded in Step 9.

Identifying a faulty disk

Overview

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

Contents

“Identifying a faulty disk” contains the following procedures:

- [“Identifying the disk” on page 62](#)
 - [“Identifying the submirrors” on page 64](#)
-

Identifying the disk

A disk problem is usually indicated in the `/var/adm/messages` file.

1. Enter: `cd /var/adm`
2. Enter: `vi messages`

The system displays the following file:

```
Jun 12 16:27:08 leopard unix: WARNING:
Jun 12 16:27:08 leopard unix: Error for command
'read(10)' Error Level: R
Jun 12 16:27:09 leopard unix: retryable
Jun 12 16:27:09 leopard unix: Requested Block 0, Error
Block: 0
Jun 12 16:27:09 leopard unix: Sense Key: Media Error
Jun 12 16:27:09 leopard unix: Vendor 'SEAGATE':
Jun 12 16:27:09 leopard unix: ASC = 0x31 (medium format
corrupted), ASCQ
= 0x0, FRU = 0x9
Jun 12 16:27:09 leopard unix: WARNING:
/sbus@3,0/SUNW,fas@3,8800000/sd@2,0
```

NOTE:

The line shown in bold indicates a disk problem. The `sd@2,0` represents SCSI disk.

3. Enter: `:q!`

The system exits the file without making any changes.

4. To determine which disk has the problem, enter:

```
ls -l /dev/dsk/c*
```

The system displays the following message:

```
.
.
.
lrwxrwxrwx  1 root    root          50 Apr 24 15:21
/dev/dsk/c0t2d0s0 -> ../../devices/sbus@3,0/SUNW,fas
@3,8800000/sd@2,0:a
lrwxrwxrwx  1 root    root          50 Apr 24 15:21
/dev/dsk/c0t2d0s1 -> ../../devices/sbus@3,0/SUNW,fas
@3,8800000/sd@2,0:b
.....
.....
.....
#
```

5. Search the output for a device description matching that in the warning message.

If the "sd@2,0:X" information matches the same information in the warning message, then that disk is the faulty disk.

⇒ NOTE:

The *x* at the end of the device is the partition number. The partition information may not be displayed in the */var/adm/messages* file. The following table shows what letters correspond to which disk partition.

Letter	Disk partition
a	0
b	1
c	2
d	3
e	4
f	5
g	6
h	7

Identifying the submirrors

1. Enter the following commands:

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

The system adds the Solstice DiskSuite directories to your path.

2. Enter:

```
metastat
```

The system displays messages similar to the following:

```
d20: Submirror of d21
Size: 14960160 blocks
Stripe 0:
  Device  Start Block  Dbase  State  Hot Spare
  c0t2d0s1      0      No  Maintenance
Stripe 1:
  Device  Start Block  Dbase  State  Hot Spare
  c0t4d0s3      0      No    Okay
```

3. Search the output for the name of the faulty disk.

The faulty disk is usually indicated by a state of Maintenance.

Example:

The following lines indicate that the faulty c0t2d0 disk is in d20:

```
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.

4. Record the metadevices to which the disk belongs. For example, **d11** and **d19**.
5. After identifying all of the faulty disks and associated submirrors, see [“Recovering a mirrored system after disk failure” on page 65](#).

Recovering a mirrored system after disk failure

Overview

“Recovering a mirrored system after disk failure” contains procedures for the recovery of a mirrored system after the failure of single or multiple hard drives.

Important!

If the system loses the primary boot disk and the alternate boot disk, the system will need to be rebuilt to factory standards and any data restored. See [“Restoring a mirrored system from a CMSADM backup” on page 90](#) for more information.

Prerequisites

Before you recover a mirrored system, perform the following tasks:

- Verify that the alternate boot device is set up. See [“Setting up the alternate boot device” on page 46](#) for more information.
 - Identify the faulty disk or disks. See [“Identifying a faulty disk” on page 62](#) for more information.
-

Contents

“Recovering a mirrored system after disk failure” contains the following procedures:

- [“Determining the recovery procedure” on page 66](#)
- [“Recovering a mirrored system that will boot into multiuser mode” on page 68](#)
- [“Recovering a mirrored system that will only boot into single user mode” on page 73](#)

Determining the recovery procedure

1. Determine the number of physical disks in the system, and record this number in the formula at Step 2.
2. Use the following formula to calculate the total number of replica state databases.

$$3(\text{Number of physical disks}) + 1 = \text{Total number of replica state databases}$$

3. Record the total number of replica state databases in the formula at Step 9.
4. Determine if any of the failed disks was the primary boot disk. The primary boot disk will always be in slot 0.
5. Record the number of failed primary boot disk replicas in the formula at Step 8.
 - If the primary boot disk did fail, there are 4 primary boot disk replicas.
 - If the primary boot disk did *not* fail, there are 0 primary boot disk replicas.
6. Determine the number of failed non-primary boot disks, and enter this number in the formula for Step 7.
7. Use the following formula to calculate, and record the total number of failed non-primary boot disk replicas.

$$3(\text{Number of failed non-primary boot disks}) = \text{Total number of failed non-primary boot disk replicas}$$

8. Use the following formula to calculate, and record the total number of failed replicas.

$$(\text{Number of failed primary boot disk replicas}) + (\text{Total number of failed non-primary boot disk replicas}) = \text{Total number of failed replica state databases}$$

 **NOTE:**

You can verify the numbers calculated by looking at the output for `metadb -i`

The lines with lowercase letters in the far left column are active replica databases. Any lines with uppercase status letters indicate a problem with the replica state database. The number of lines with uppercase status letters should be equal to the total number of failed replica state databases in Step 8.

9. Use the following formula to calculate, and record half the total number of replica state databases.

$$\left(\frac{\text{Total number of replica state databases} - 1}{2} \right) + 1 = \text{Half the total number of replica state databases}$$

10. Compare the total number of failed replica state databases calculated in Step 8 to half the total number of replica state databases calculated in Step 9 and perform one of the following actions:
 - If the number from Step 8 is greater than or equal to the number from Step 9, see [“Recovering a mirrored system that will boot into multiuser mode”](#) on page 68.
 - If the number from Step 8 is less than the number from Step 9, see [“Recovering a mirrored system that will only boot into single user mode”](#) on page 73.

Recovering a mirrored system that will boot into multiuser mode

Contents

Use this procedure to recover a mirrored system that will boot into multiuser mode.

“Recovering a mirrored system that will boot into multiuser mode” contains the following procedures:

- [“Removing the failed disk” on page 68](#)
- [“Re-creating the database replicas” on page 70](#)
- [“Syncing the new disk” on page 70](#)

Removing the failed disk

1. Enter:

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

2. Turn off the system and any external devices.

3. Install the replacement disk.

NOTE:

For more information about installation of hard drives, see the appropriate Hardware Installation, Maintenance, and Troubleshooting book for your platform.

For more information about the physical set up of your system, see [“Hardware Requirements” on page 9](#).

4. Turn on any external devices and then the system.

5. As the console shows that the system is booting, press the **Stop + A** keys.

The system displays the `ok` prompt.

6. Enter: `boot -r`

The system boots into multi-user mode.

7. Log on as **root**.

8. Enter the following commands:

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

9. Enter: `metadb -i`

The system displays all of the database replicas.

```

      flags          first blk      block count
M    p  luo         16             1034        /dev/dsk/c0t0d0s1
M    p  luo        1050            1034        /dev/dsk/c0t0d0s1
M    p  luo        2084            1034        /dev/dsk/c0t0d0s1
M    p  luo        3118            1034        /dev/dsk/c0t0d0s1
.....
.....
.....
M    p  luo         16             1034        /dev/dsk/c0t12d0s0
M    p  luo        1050            1034        /dev/dsk/c0t12d0s0
M    p  luo        2084            1034        /dev/dsk/c0t12d0s0
o - replica active prior to last mddb configuration change
u - replica is up to date
l - locator for this replica was read successfully
c - replica's location was in /etc/opt/SUNWmd/mddb.cf
p - replica's location was patched in kernel
m - replica is master, this is replica selected as input
W - replica has device write errors
a - replica is active, commits are occurring to this replica
M - replica had problem with master blocks
D - replica had problem with data blocks
F - replica had format problems
S - replica is too small to hold current data base
R - replica had device read errors

```

10. Look for any replicas with a capital status letter in the first column. All replicas with a capital status letter need to be removed.
11. To remove any damaged database replicas, enter:

```
metadb -d -f /dev/dsk/disk_slice
```

Where ***disk_slice*** is the physical device partition number containing the damaged database replica.

Example:

Using the output obtained in Step 9, the command would be entered as:

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

12. Repeat Steps 9 through 11 until all lines with capital letters have been removed.

WARNING:

Do not reboot the system after removing the database replicas. The system may not boot.

13. Use the *format* command to partition the replacement disk exactly like the original disk. See [“Partitioning disks with the format command” on page 78](#) for more information.

Recovering a mirrored system after disk failure

Re-creating the database replicas

1. Re-create the database replicas by entering:

```
metadb -a -cX -f /dev/dsk/disk_slice
```

Where **X** is the number of database replicas per slice.

- If the hard drive is in slot 0, **X** will be 4.
- If the hard drive is in any other slot except 0, **X** will be 3.

Where **disk_slice** is the physical device partition number that the database replica resides on.

- The replica database will be on slice 1, for the primary and alternate boot drives.
- The replica database will be on slice 0, for all other hard drives.

Syncing the new disk

1. Enter:

```
metastat | egrep metareplace
```

The system displays a message similar to the following:

```
Invoke: metareplace d13 c0t0d0s0 <new device>  
Invoke: metareplace d21 c0t0d0s3 <new device>  
Invoke: metareplace d17 c0t0d0s4 <new device>
```

2. Enter the following commands:

```
metareplace -e d13 new_device
```

```
metareplace -e d21 new_device
```

```
metareplace -e d17 new_device
```

Where **new_device** is the new metadvice number.

Example:

If the system displays the following message when the `metastat | egrep metareplace` command is entered:

```
Invoke: metareplace d13 c0t0d0s0 <new device>  
Invoke: metareplace d21 c0t0d0s3 <new device>  
Invoke: metareplace d17 c0t0d0s4 <new device>
```

The following commands would be entered:

```
metareplace -e d13 c0t0d0s0
```

```
metareplace -e d21 c0t0d0s3
```

```
metareplace -e d17 c0t0d0s4
```

3. Enter:

```
metastat |egrep Resync
```

The system displays a message similar to the following:

```
State: Resyncing
Resync in progress: 5% done
State: Resyncing
c0t0d0s0      0  No  Resyncing
State: Resyncing
Resync in progress: 0% done
State: Resyncing
c0t0d0s3      0  No  Resyncing
State: Resyncing
Resync in progress: 0% done
State: Resyncing
c0t0d0s4      0  No  Resyncing
```

4. Look for the system resyncing.

5. When the system is done resyncing, enter:

```
metastat |pg
```

The system displays a message similar to the following:

```
d20: Submirror of d21
State: Okay
Size: 20811600 blocks
Stripe 0:
  Device          Start Block  Dbase State      Hot Spare
  c0t1d0s3         0           No  Okay
Stripe 1:
  Device          Start Block  Dbase State      Hot Spare
  c0t10d0s1        0           No  Okay
Stripe 2:
  Device          Start Block  Dbase State      Hot Spare
  c0t12d0s1        0           No  Okay
```

6. Verify that none of the metadevices show a state of maintenance.

⇒ NOTE:

If any of the metadevices show a state of maintenance repeat the procedure for [“Syncing the new disk” on page 70](#) until none of the metadevices show a state of maintenance. These steps are necessary if multiple disk stripes have failed on a metadvice. DiskSuite will only sync one stripe at a time.

Recovering a mirrored system after disk failure

7. Enter:

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

8. Log in to the system as **root**.

9. Enter:

```
metastat
```

The system displays messages similar to the following:

```
d20: Submirror of d21
State: Okay
Size: 20811600 blocks
Stripe 0:
  Device          Start Block  Dbase State      Hot Spare
  c0t1d0s3         0           No   Okay
Stripe 1:
  Device          Start Block  Dbase State      Hot Spare
  c0t10d0s1        0           No   Okay
Stripe 2:
  Device          Start Block  Dbase State      Hot Spare
  c0t12d0s1        0           No   Okay
```

10. Verify that all metadevices display a state of Okay.

Recovering a mirrored system that will only boot into single user mode

Contents

Removing the failed disk

Use this procedure to recover a mirrored system that will only boot into maintenance mode.

“Recovering a mirrored system that will only boot into single user mode” contains the following procedures:

- [“Removing the failed disk” on page 73](#)
- [“Re-creating the database replicas” on page 75](#)
- [“Syncing the new disk” on page 75](#)

1. Turn off the system and any external devices.
2. Install the replacement disk.

⇒ NOTE:

For more information about installation of hard drives, see the appropriate Hardware Installation, Maintenance, and Troubleshooting book for your platform.

For more information about the physical set up of your system, see [“Hardware Requirements” on page 9](#).

3. Turn on any external devices and then the system.

The system displays error messages about stale databases and a prompt to log in to maintenance mode.

⇒ NOTE:

If replacement disks are not available, you may be able to boot off of the alternate boot drive by performing Steps [4](#) through [10](#). Once the disks have been installed, reboot the system and start at Step [11](#).

4. Enter the root password for system maintenance.
5. Enter the following commands:

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

⇒ NOTE:

Ignore any read only file system messages.

Recovering a mirrored system after disk failure

6. Enter: `metadb -i`

The system displays all of the database replicas.

```
      flags          first blk      block count
M    p  luo         16              1034      /dev/dsk/c0t0d0s1
M    p  luo        1050              1034      /dev/dsk/c0t0d0s1
M    p  luo        2084              1034      /dev/dsk/c0t0d0s1
M    p  luo        3118              1034      /dev/dsk/c0t0d0s1
.....
.....
.....
M    p  luo         16              1034      /dev/dsk/c0t12d0s0
M    p  luo        1050              1034      /dev/dsk/c0t12d0s0
M    p  luo        2084              1034      /dev/dsk/c0t12d0s0
o - replica active prior to last mddb configuration change
u - replica is up to date
l - locator for this replica was read successfully
c - replica's location was in /etc/opt/SUNWmd/mddb.cf
p - replica's location was patched in kernel
m - replica is master, this is replica selected as input
W - replica has device write errors
a - replica is active, commits are occurring to this replica
M - replica had problem with master blocks
D - replica had problem with data blocks
F - replica had format problems
S - replica is too small to hold current data base
R - replica had device read errors
```

7. Look any replicas with a capital status letter in the first column.

8. All replicas with a capital status letter need to be removed. Enter:

```
metadb -d -f /dev/dsk/disk_slice
```

Where *disk_slice* is the physical device partition number containing the damaged database replica.

Example:

Using the output obtained in Step 6, the command would be entered as:

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

9. Repeat Steps 6 through 7 until all lines with capital letters have been removed.

10. Enter:

```
reboot -- -r
```

The system reboots.

11. Use the *format* command to partition the replacement disk exactly like the original disk. See [“Partitioning disks with the format command” on page 78](#) for more information.

Re-creating the database replicas

1. Re-create the database replicas by entering:

```
metadb -a -cX /dev/dsk/disk_slice
```

Where **X** is the number of database replicas per slice.

- If the hard drive is in slot 0, **X** will be 4.
- If the hard drive is in any other slot except 0, **X** will be 3.

Where **disk_slice** is the physical device partition number that the database replica resides on.

- The replica database will be on slice 1, for the primary and alternate boot drives.
- The replica database will be on slice 0, for all other hard drives.

Syncing the new disk

1. Enter:

```
metastat | egrep metareplace
```

The system displays a message similar to the following:

```
Invoke: metareplace d13 c0t0d0s0 <new device>  
Invoke: metareplace d21 c0t0d0s3 <new device>  
Invoke: metareplace d17 c0t0d0s4 <new device>
```

2. Enter the following commands:

```
metareplace -e d13 new_device
```

```
metareplace -e d21 new_device
```

```
metareplace -e d17 new_device
```

Where **new_device** is the new metadvice number.

Example:

If the system displays the following message when the `metastat | egrep metareplace` command is entered:

```
Invoke: metareplace d13 c0t0d0s0 <new device>  
Invoke: metareplace d21 c0t0d0s3 <new device>  
Invoke: metareplace d17 c0t0d0s4 <new device>
```

The following commands would be entered:

```
metareplace -e d13 c0t0d0s0
```

```
metareplace -e d21 c0t0d0s3
```

```
metareplace -e d17 c0t0d0s4
```

Recovering a mirrored system after disk failure

3. Enter:

```
metastat |egrep Resync
```

The system displays a message similar to the following:

```
State: Resyncing
Resync in progress: 5% done
State: Resyncing
c0t0d0s0      0   No   Resyncing
State: Resyncing
Resync in progress: 0% done
State: Resyncing
c0t0d0s3      0   No   Resyncing
State: Resyncing
Resync in progress: 0% done
State: Resyncing
c0t0d0s4      0   No   Resyncing
```

4. Look for the system resyncing.

5. When the system is done resyncing, enter:

```
metastat |pg
```

The system displays a message similar to the following:

```
d20: Submirror of d21
State: Okay
Size: 20811600 blocks
Stripe 0:
  Device                Start Block  Dbase State      Hot Spare
  c0t1d0s3                0             Okay
Stripe 1:
  Device                Start Block  Dbase State      Hot Spare
  c0t10d0s1              0             Okay
Stripe 2:
  Device                Start Block  Dbase State      Hot Spare
  c0t12d0s1              0             Okay
```

6. Verify that none of the metadevices show a state of maintenance.

⇒ NOTE:

If any of the metadevices show a state of maintenance repeat the procedure for [“Syncing the new disk” on page 75](#) until none of the metadevices show a state of maintenance. These steps are necessary if multiple disk stripes have failed on a metadevice. DiskSuite will only sync one stripe at a time.

7. Enter:

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

8. Log in to the system as **root**.

9. Enter:

```
metastat
```

The system displays messages similar to the following:

```
d20: Submirror of d21
State: Okay
Size: 20811600 blocks
Stripe 0:
  Device          Start Block  Dbase State      Hot Spare
  c0t1d0s3        0           Okay
Stripe 1:
  Device          Start Block  Dbase State      Hot Spare
  c0t10d0s1       0           Okay
Stripe 2:
  Device          Start Block  Dbase State      Hot Spare
  c0t12d0s1       0           Okay
```

10. Verify that all metadevices display a state of Okay.

Partitioning disks with the format command

To use the format command to partition a disk:

1. Enter: **format**

The system displays a menu of disks. The system prompts you to select one.

2. Enter the menu number of the disk to be partitioned.

The system displays a list of commands and the `format>` prompt.

 **NOTE:**

If you have an Enterprise 3500, make certain that the new drives are named using “c1,” as c1t4d0 or c1t5d0.

3. Enter: **partition**

The system displays the `partition>` prompt.

4. Enter: **print**

The system displays the partition table and the `partition>` prompt.

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

The system displays a series of prompts for partitioning information.

6. Answer the prompts according to the CMS version and disk type you are partitioning. See the [“Disk partitioning tables” on page 79](#) for partition information.

Disk partitioning tables **CMS R3V8 Boot Disk Partition Values:**

Partition #		0	1	2 ^a	3	4
Flag		wm	wm	wm	wm	wu
Tag		boot disk 1: / or root <i>boot disk 2:</i> unassigned	unassigned	overlap <i>or</i> backup	unassigned	swap
4.2-GB Disk	Starting Cylinder	0	1023	0	1030	2909
	Size (cyl.)	1023	7	3880	1879	971
9.1-GB Disk SCSI (E3000) and FCAL (E3500)	Starting Cylinder	0	616	0	623	4339
	Size (cyl.)	616	7	4924	3716	585
18-GB Disk SCSI (E3000) and FCAL (E3500)	Starting Cylinder	0	471	0	478	7070
	Size (cyl.)	471	7	7506	6592	436

^a The default size of the overlap file system is always the size of the entire disk. Occasionally, the name *backup* will appear instead of *overlap*. Do not change the slice 2 value or name.

CMS R3V8 Nonboot Disk Partition Values:

Partition #		0	1	2 ^a
Flag		wm	wm	wm
Tag		unassigned	unassigned	overlap <i>or</i> backup
4.2-GB Disk	Starting Cylinder	0	2	0
	Size (cyl.)	2	3878	3880

Partitioning disks with the format command

Partition #		0	1	2 ^a
9.1-GB Disk SCSI (E3000) and FCAL (E3500)	Starting Cylinder	0	2	0
	Size (cyl.)	2	4922	4924
18-GB Disk SCSI (E3000) and FCAL (E3500)	Starting Cylinder	0	2	0
	Size (cyl.)	2	7504	7506
36-GB Disk FCAL (E3500)	Starting Cylinder	0	2	0
	Size (cyl.)	2	24618	24620

^a The default size of the overlap file system is always the size of the entire disk. Occasionally, the name *backup* will appear instead of *overlap*. Do not change the slice 2 value or name.

: 7. Enter: **print**

The system displays the partition table and the `partition>` prompt.

8. Compare the display to the “[Disk partitioning tables](#)” on page 79 to make certain you have repartitioned the disk properly. If not, repeat Steps 5 through 8.

9. Enter: **label**

The system displays a “Ready to label disk, continue?” message.

10. Enter: `y`
The system displays the `partition>` prompt.
11. Enter: `q`
The system displays the `format>` prompt.
12. If there is another disk to be partitioned, perform the following procedure:
 - a. Enter: `disk`
 - b. Repeat Steps 2 through 11.
13. After all of the disks have been partitioned, enter: `q`

Restoring Data from a CMSADM Backup

Overview

This section contains the procedure for restoring /cms data files from a 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.

Tape drive considerations

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

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

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

Procedure

To restore data from a CMSADM backup:

1. Obtain the latest CMSADM backup tape.
2. Load the backup tape into the tape drive.

3. Enter the following command on a single line at the command prompt:

```
cpio -icmudv -C 10240 -I /dev/rmt/device -M  
"Insert Tape %d" "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 or 40-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 *CentreVu CMS R3V8 Administration*, 585-210-941 document for details on restoring a maintenance backup.

Repairing or rebuilding the /cms file system

Overview

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

Contents

“Repairing or rebuilding the /cms file system” contains the following procedures:

- [“Repairing /cms” on page 84](#)
 - [“Rebuilding /cms” on page 86](#)
-

Repairing /cms

To repair /cms:

1. Turn off CMS.
 - a. Enter `cmsadm` or `cmssvc`.
 - b. Select the option number for `run_cms`.
 - c. Select the option number for `Turn off CMS`.
2. Enter:

```
umount /cms
```

The system unmounts /cms, and then displays the # prompt.
3. **Do this step only if the `umount` command returned the response `/cms: device busy`.**
 - a. Load the `/etc/vfstab` file into your editor.
 - b. Find the `/cms` and `/cms/swap` lines. Comment out both lines. (To comment out a line, insert a pound sign at the beginning.)
 - c. Save the file and exit the editor.
 - d. Reboot by entering the following command:

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

4. Attempt to repair /cms by entering:

```
fsck -y /dev/md/rdisk/d21
```

The system displays the following message:

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

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

If you get an error message, however, repeat the `fsck` command. If the command still returns errors after 10 repetitions, stop this procedure and skip ahead to [“Rebuilding /cms” on page 86](#).

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

5. This step remounts /cms and reallocates the swap file. How that is done depends upon whether you altered the `vfstab` file earlier in this procedure.
- If you altered the `vfstab` file earlier (Step 3), do the following:
 - a. Load the `/etc/vfstab` file into your editor.
 - b. Find the `#/cms` line. Uncomment the line. (To uncomment a line, delete the initial pound sign.)
 - c. Save the file and exit the editor.
 - d. Reboot by entering:

```
/usr/sbin/shutdown -i6 -g0 -y
```
 - If you did NOT alter the `vfstab` file, do the following:
 - a. Enter:

```
mount /cms
```

The system remounts /cms.

Repairing or rebuilding the /cms file system

6. Turn on CMS:
 - a. Enter either `cmsadm` or `cmssvc`.
 - b. Select the option number for `run_cms`.
 - c. Select the option number for `Turn on CMS`.

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

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

Rebuilding /cms

Overview

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

If you are running a Sun Enterprise 3500 platform, do not begin this procedure until you have checked the `/kernel/drv/st.conf` file to make sure the 3500’s 8mm tape drive is properly defined. See [“Tape drive considerations” on page 82](#) for details.

Prerequisites

Before you rebuild /cms, perform the following tasks:

- Obtain the latest CMSADM backup.
- Verify that /cms is unmounted. If it is not, repeat Steps 2 through 4 of [“Repairing /cms” on page 84](#).

Procedure

1. Add the Solstice DiskSuite directories to your path:

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

2. Enter:

```
metadetach d21 d20
```

The system detaches the d20 submirror and displays the following message:

```
d21: submirror d20 is detached
```

3. Enter:

```
metaclear d20
```

The system clears the d20 submirror and displays the following message:

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

4. Enter:

```
metaclear d21
```

The system clears the cms mirror and displays the following message:

```
mirror is cleared
```

5. Enter:

```
metaclear d19
```

The system clears the d19 submirror and displays the following message:

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

6. Enter:

```
metainit d19
```

The system reinitializes the d19 submirror and displays the following message:

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

7. Enter:

```
metainit d20
```

The system reinitializes the d20 submirror and displays the following message:

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

8. Enter:

```
metainit d21
```

The system reinitializes the cms mirror and displays the following message:

```
mirror is set up
```

Repairing or rebuilding the /cms file system

9. Prepare the cms mirror for a new file system by entering:

```
newfs -m 0 /dev/md/rdisk/d21
```

The system displays the following message:

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

⇒ NOTE:

This will take 15 to 30 minutes to complete.

10. How you remount /cms depends upon whether you had to alter the *vfstab* file when you attempted to repair the file system.

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

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

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

The system reboots.

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

- a. Enter:

```
mount /cms
```

The system remounts /cms.

11. Enter: `cd /`

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

13. Restore CMS data from the tape by entering the following command on a single line at the command prompt:

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

14. Restore the latest full maintenance backup.

Restoring specific files from a CMSADM backup

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

1. Enter:

```
cd /
```

2. Enter:

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

Where *full_path_name* is the full path name of the file to be restored.

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

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

Restoring a mirrored system from a CMSADM backup

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

Restoring the backup

Restore the backup according to your specific platform. This procedure are for the following platforms groups:

- [“Restoring the backup on a E3000 system” on page 90](#)
- [“Restoring the backup on a E3500 system” on page 93](#)

Restoring the backup on a E3000 system

1. Perform the following procedures, see Chapter 2 “Installing the Solaris Operating System” of the *CentreVu CMS R3V8 Software Installation, Maintenance, and Troubleshooting*, 585-210-941 book for additional information.
 - a. “Booting from the Solaris 7 Software 3/99 CD”
 - b. “Identifying the System”
 - c. “Setting the Date and Time”
 - d. “Selecting the Solaris 7 System Files”
 - e. “Partitioning the Hard Disks”
 - f. “Assigning a Root Password”
2. Insert the first cmsadm backup tape into the tape drive.
3. To verify the tape device path, enter:

```
mt -f /dev/rmt/dev# status
```

Where **dev#** is the device name.

If you enter an incorrect device path, the system displays a “No such file or directory” message; the correct device path displays information similar to the following:

```
Tandberg 2.5 Gig QIC tape drive:  
sense key(0x0)= No Additional Sense    residual= 0  
retries= 0      file no= 0    block no= 0
```

NOTE:

The device name is usually `/dev/rmt/0c`. However, the device name used depends on the drive's SCSI ID. Possible device names are:

- `/dev/rmt/0` Indicates the first noncompressing tape drive with the lowest target address
- `/dev/rmt/1` Indicates the second noncompressing tape drive with the second lowest target address
- `/dev/rmt/0c` Indicates the first compressed-mode tape drive with the lowest target address
- `/dev/rmt/1c` Indicates the second compressed-mode tape drive with the second lowest target address

4. Enter:

```
cd /
```

5. To restore all files excluding CMS, perform one of the following actions:

- If you have only one backup tape, enter the following command:

```
cpio -icmudvf -C 10240 -I /dev/rmt/dev# "cms/*"
```

Where **dev#** is substituted for the device name.

- If you have more than one backup tape, enter the following command on a single line at the command prompt:

```
cpio -icmudvf -C 10240 -I /dev/rmt/dev# -M  
"Remove current tape, insert tape number %d,  
and press Enter" "cms/*"
```

Where **dev#** is the device name.

CAUTION:

Do not reboot the system after the tape restore. If the system is shutdown you will have to repeat the entire procedure.

Restoring a mirrored system from a CMSADM backup

As the restore proceeds, the light-emitting diode (LED) on the tape drive will alternately flash and light steadily. You may see some error messages about the `/home` and `/xfn` directories. These errors are displayed when the directories are already present and can be ignored.

```
cpio: Cannot create temporary file, errno 18, Cross-device link
(There is a time delay between the first message and the rest of the
messages)
cpio: Cannot chmod() "/home", errno 89, Operation not applicable
cpio: Unable to reset modification time for "/home", errno 89,
Operation not applicable
cpio: Cannot chmod() "/home", errno 89, Operation not applicable
cpio: Cannot chown() "/home", errno 89, Operation not applicable
cpio: Cannot chmod() "/xfn", errno 89, Operation not applicable
cpio: Unable to reset modification time for "/xfn", errno 89, Operation
not applicable
cpio: Cannot chmod() "/xfn", errno 89, Operation not applicable
cpio: Cannot chown() "/xfn", errno 89, Operation not applicable
602780 blocks
9 error(s)
```

6. Enter:

```
vi /etc/inittab
```

The system displays the following message:

```
ap::sysinit:/sbin/autopush -f /etc/iu.ap
ap::sysinit:/sbin/soconfig -f /etc/sock2path
.....
.....
.....
od:234:respawn:/cms/dc/odbc/rqb_start
as:0236:respawn:/opt/cc/aas/bin/faasdb.sh </dev/null >/dev/null 2>&1
cm:0236:respawn:/cms/bin/cms_mon /cms </dev/null >/dev/null 2>&1
```

7. Look for a `cm:` entry in the last line of the file. If a `cm:` entry is present, remove it.

8. Press **Esc**, and then enter **:wq!** to save and exit the file.

Restoring the backup on a E3500 system

1. Perform the following procedures, see Chapter 2 "Installing the Solaris Operating System" of the *CentreVu CMS R3V8 Software Installation, Maintenance, and Troubleshooting*, 585-210-941 book for procedure descriptions.
 - a. "Booting from the Solaris 7 Software 3/99 CD"
 - b. "Identifying the System"
 - c. "Setting the Date and Time"
 - d. "Selecting the Solaris 7 System Files"
 - e. "Partitioning the Hard Disks"
 - f. "Assigning a Root Password"



CAUTION:

Do not reboot the system until you are directed to do so in this procedure. If the system is shutdown prematurely, you will have to repeat the entire procedure.

2. For *Enterprise 3500* platforms, you must alter the `/kernel/drv/st.conf` file so the operating system recognizes the 8mm tape drive. Do the following steps:
 - a. Enter:

```
vi /kernel/drv/st.conf
```
 - b. Find the `#tape-config-list=` line. It should be in the vicinity of line number 40.
 - c. Insert, immediately before that line, the following content:

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

- d. Enter:

```
:wq!
```

3. Insert the first cmsadm backup tape into the tape drive.

Restoring a mirrored system from a CMSADM backup

4. To verify the tape device path, enter:

```
mt -f /dev/rmt/dev# status
```

Where, **dev#** is the device name.

If you enter an incorrect device path, the system displays a “No such file or directory” message; the correct device path displays information similar to the following:

```
Tandberg 2.5 Gig QIC tape drive:
sense key(0x0)= No Additional Sense   residual= 0
retries= 0      file no= 0   block no= 0
```

⇒ NOTE:

The device name is usually `/dev/rmt/0c`. However, the device name used depends on the drive's SCSI ID. Possible device names are:

<code>/dev/rmt/0</code>	Indicates the first noncompressing tape drive with the lowest target address
<code>/dev/rmt/1</code>	Indicates the second noncompressing tape drive with the second lowest target address
<code>/dev/rmt/0c</code>	Indicates the first compressed-mode tape drive with the lowest target address
<code>/dev/rmt/1c</code>	Indicates the second compressed-mode tape drive with the second lowest target address

5. Enter:

```
cd /
```

6. To restore all files excluding CMS, `dev/dsk` and `dev/rdisk` files perform one of the following actions:

- If you have only one backup tape, enter the following command on a single line at the command prompt:

```
cpio -icmudvf -C 10240 -I /dev/rmt/dev# "cms/*"
"dev/dsk" "dev/rdisk" "dev/dsk/*" "dev/rdisk/*"
"etc/path_to_inst"
```

Where **dev#** is substituted for the tape device name.

Restoring a mirrored system from a CMSADM backup

- If you have more than one backup tape, enter the following command on a single line at the command prompt:

```
cpio -icmudvf -C 10240 -I /dev/rmt/dev#  
-M "Remove current tape, insert tape number %d,  
and press Enter" "cms/*" "dev/dsk"  
"dev/rdsk" "dev/dsk/*" "dev/rdsk/*"  
"etc/path_to_inst"
```

Where **dev#** is the device name.

As the restore proceeds, the light-emitting diode (LED) on the tape drive will alternately flash and light steadily.

CAUTION:

Do not reboot the system after the tape restore. The system will not successfully boot until `olds -cleanup` is run. If the system is shutdown you will have to repeat the entire procedure.

You may see some error messages about the `/home` and `/xfn` directories. These errors are displayed when the directories are already present and can be ignored.

```
cpio: Cannot create temporary file, errno 18, Cross-device link  
(There is a time delay between the first message and the rest of the  
messages)  
cpio: Cannot chmod() "/home", errno 89, Operation not applicable  
cpio: Unable to reset modification time for "/home", errno 89,  
Operation not applicable  
cpio: Cannot chmod() "/home", errno 89, Operation not applicable  
cpio: Cannot chown() "/home", errno 89, Operation not applicable  
cpio: Cannot chmod() "/xfn", errno 89, Operation not applicable  
cpio: Unable to reset modification time for "/xfn", errno 89, Operation  
not applicable  
cpio: Cannot chmod() "/xfn", errno 89, Operation not applicable  
cpio: Cannot chown() "/xfn", errno 89, Operation not applicable  
602780 blocks  
9 error(s)
```

Restoring a mirrored system from a CMSADM backup

7. Enter:

```
vi /etc/inittab
```

The system displays the `inittab` file.

```
ap::sysinit:/sbin/autopush -f /etc/iu.ap
ap::sysinit:/sbin/soconfig -f /etc/sock2path
.....
.....
.....
od:234:respawn:/cms/dc/odbc/rqb_start
as:0236:respawn:/opt/cc/aas/bin/faasdb.sh </dev/null >/dev/null 2>&1
cm:0236:respawn:/cms/bin/cms_mon /cms </dev/null >/dev/null 2>&1
```

8. Look for a `cm:` entry in the last line of the file. If a `cm:` entry is present, remove it.

9. Enter `:wq!` to save and exit the file.

Re-establishing the disk mirrors

This section explains how to re-establish a disk mirroring system according to your specific platform. This section contains the following procedures:

- [“Re-establishing the mirrors on the E3000 system” on page 96](#)
- [“Re-establishing the mirrors on a E3500 system” on page 104](#)

Re-establishing the mirrors on the E3000 system

To re-establish the mirrors on a E3000 system:

1. Enter: `cd /dev/md/dsk`

2. Enter:

```
ln -s ../../../../devices/pseudo/md@0:0,17,blk d17
```

3. Add the Solstice DiskSuite directories to your path by entering the following commands:

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

```
export PATH
```

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

```
olds -mirrored -cleanup
```

The system displays a message similar to the following:

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

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

5. Enter: **y**

6. Enter: **sync; sync; reboot -- -r**

⇒ NOTE:

Ignore any FCK errors displayed during the boot process.

7. Add the Solstice DiskSuite directories to your path by entering the following commands:

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

```
export PATH
```

8. Check the disk partitioning by entering:

```
olds -mirrored -check_disks
```

The system displays the following message:

```
.  
.  
Success, checking md.tab.new
```

9. Create DiskSuite files by entering:

```
olds -mirrored -mk_files
```

The system displays the following message:

```
.  
.  
Success, creating md.tab.new
```

10. Activate the state database by entering:

```
olds -mirrored -metadbs
```

The system displays the following message:

```
.  
.  
Success, setting up metadb replicas
```

Restoring a mirrored system from a CMSADM backup

11. Set up the root mirror by entering:

```
olds -mirrored -setroot
```

The system displays the following message:

```
.  
.
Success, root mirrored successfully.
```

12. Enter:

```
metastat | grep %
```

The system will return to the prompt without displaying any output *only* after the syncing process is complete.

CAUTION:

Wait until the metadevices are finished synchronizing, before continuing with Step 13. Do **NOT** proceed until synchronizing is finished.

13. Enter:

```
sync; sync; reboot -- -r
```

The system reboots.

14. Log in as **root**.

15. Add the DiskSuite directories to your path: by entering the following commands:

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

```
export PATH
```

16. Activate the `/cms` metadevice by entering:

```
olds -mirrored -setup
```

The system displays the following message:

```
.  
.
14550512, 14585216, 14619920, 14654624, 14689328, 147240
14758736, 14793440, 14828144, 14862848, 14897552, 149299
ufs fsck: sanity check: /dev/md/rdisk/d21 okay
Success, /cms mirrored successfully
```

17. Mount `/cms` by entering:

```
mount /cms
```

Restoring a mirrored system from a CMSADM backup

18. Enter:

```
cd /
```

19. To restore CMS, perform one of the following actions:

- If you have only one backup tape, enter the following command:

```
cpio -icmudv -C 10240 -I /dev/rmt/dev# "cms/*"
```

Where **dev#** is the device name.

- If you have more than one backup tape, enter the following command on a single line at the command prompt:

```
cpio -icmudv -C 10240 -I /dev/rmt/dev# -M  
"Remove current tape, insert tape number %d,  
and press Enter" "cms/*"
```

Where **dev#** is the device name.

20. Enter:

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

The system reboots.

21. Choose one of the following procedures:

- Set up CMS interactively. See "Setting Up CMS Interactively from a Terminal" in Chapter 2, of *CentreVu CMS R3V8 Software Installation, Maintenance and Troubleshooting* 585-210-941.
- Set up CMS from a Unix flat file (below)

⇒ NOTE:

CMS setup information was created when CMS was initially setup. The Unix flat file will not contain any updates made after the initial system setup.

- a. To setup CMS from the flat file, enter the following command on a single line at the command prompt:

```
cp /cms/install/logdir/cms.inst.out  
/cms/install/cms_install/cms.install
```

- b. Enter:

```
chmod 644 /cms/install/cms_install/cms.install
```

- c. Enter:

```
cmssvc
```

Restoring a mirrored system from a CMSADM backup

The system displays the CMS Services Menu.

```
Failure in sql 'database' (-329, -2)
>>> CMS is already off <<<

CentreVu(R) Call Management System Services Menu

Select a command from the list below.
 1) auth_display Display feature authorizations
 2) auth_set     Authorize capabilities/capacities
 3) run_cms      Turn CentreVu CMS on or off
 4) setup        Set up the initial configuration
 5) swinfo       Display switch information
 6) swsetup      Change switch information
 7) patch_inst   Install a single CMS patch from CD
 8) patch_rmv    Backout an installed CMS patch
 9) load_all     Install all CMS patches found on CD
10) back_all     Backout all installed CMS patches from machine
Enter choice (1-10) or q to quit:
```

d. Enter **4** to select the setup option.

The system displays the following message:

```
Select the language for this server:

All languages are ISO Latin except Japanese. Selection of the
server language assumes that existing customer data is
compatible. (Upgrade from any ISO Latin language to any ISO
Latin language or from Japanese to Japanese is supported).

 1) English
 2) Dutch
 3) French
 4) German
 5) Italian
 6) Portuguese
 7) Spanish
 8) Japanese
Enter choice (1-8): (default: 1)
```

Restoring a mirrored system from a CMSADM backup

- e. Enter the number for the language used on the system.
The system displays the following message:

```
The input will be read from
 1) the terminal
 2) a flat file
Enter choice (1-2):
```

- f. Enter **2** to select the flat file option.

The system displays the following message:

```
*** The rest of this command is running in the background ***
```

- g. Enter:

```
tail -f /cms/install/logdir/admin.log
```

⇒ NOTE:

The `-f` option in the `tail` command updates the console as messages are written to the `admin.log` file. All failure messages are logged in this file.

You may edit the CMS flat file and add comments about the packages that were installed or authorized.

Restoring a mirrored system from a CMSADM backup

The system displays the following message:

```
01350 Mon Nov 6 12:19:24 2000 SRC_ERR_NUM=-00329
PROCESS=pre_cms_env PID=000482 Sql.c:00071
SEVERITY=INFO ACD=-01 startdb

01308 Mon Nov 6 12:19:24 2000 SRC_ERR_NUM=-00329
PROCESS=pre_cms_env PID=000482 getlang.c:00027
SEVERITY=INFO ACD=-01 Select::execute

CMS setup started Mon Nov 6 12:19:26 MST 2000
Failure in sql 'database' (-349, -2)
Failure in sql 'drop database' (-329, -2)
Failure in sql 'drop database' (-329, -2)
Schema upgrade started Mon Nov 6 12:19:28 MST 2000
Executing schema files from /cms/install/dbinit/cms:
r3v5ai.a.1 r3v5ai.a.2 r3v6as.a.2 r3v6as.a.3 r3v8aa.c.1
r3v8aa.d.0 r3v8aa.d.1 r3v8aa.i.0 r3v8ab.d.0 r3v8ab.d.1
r3v8ab.e.0 r3v8ae.a.0 r3v8ae.a.1 r3v8ae.a.2 r3v8ae.a.3
r3v8ae.f.0 r3v8af.c.0 r3v8af.e.0 r3v8ag.f.0 r3v8ag.f.1
r3v8ai.b.0
Customer CMS data successfully initialized. Mon Nov 6
12:22:29 MST 2000
Executing schema files from /cms/install/dbinit/vec:
r3v5ai.a.5 r3v5ai.a.6 r3v6as.a.5 r3v6as.a.6 r3v6as.a.7
r3v8ae.d.1
Approximate space requirements:
  Data Type                Allocation
  -----
  VDNs                      1425509
  Splits                    888602
  Trunks                    623270
  Agents                    540396
  Trunk groups              328199
  Vectors                   309940
  Agent login/logout records 6390
  Agents trace records      4900
  Exceptions records        250
  Call work codes           26
  Call records              8
  Forecasting               0

File systems/space available:
  /cms      12196974
File systems/current blocks free:
  /cms      12196974
/cms:
VDN,SPLIT,TRUNK,AGENT,TKGRP,VECTOR,AGENT_LOG_REC,AGENT_TRAC
E_REC,EXCEPTIONS_REC,WORKCODE,CALL_REC,
Number of calls to fill_fs():12
Setup completed successfully Mon Nov 6 12:24:20 MST 2000
```

h. Press **Delete** to break out of the `tail -f` command.

22. Enter:

```
cp /cms/install/cms_install/fp.install /cms/db
```

The system will display one of two possible messages:

- If feature packages were installed on the system, the system displays the command prompt (#).
- If no feature packages were installed on the system, the system displays the following message:

```
cp: cannot access /cms/install/cms_install/fp.install
```

The **fp.install** file will be present on the tape backup if feature packages were installed. If the system cannot access the file, you will have to manually set up the feature packages that were previously authorized and installed on the system.

23. Choose one of the following procedures:

- If *no* feature packages were installed on the system, go to Step 25.
- If feature packages were installed on the system, perform the following procedure:

a. Enter:

```
. /opt/informix/bin/setenv
```

b. Enter:

```
export PATH=$PATH:/cms/install/bin
```

c. Enter the following command on a single line at the command prompt:

```
cp -p /cms/install/cms_install/fp.install  
/cms/db/fp.install
```

d. Enter the following command to install the feature packages that were present on the system before the restore:

```
/cms/install/bin/ins_proc -f
```

Restoring a mirrored system from a CMSADM backup

24. The system displays a message similar to the following example:

```
Feature package installation begun Mon Apr  8 13:45:40 MDT
2002
Feature package installation begun.
Forecasting package begin installation Mon Apr  8 13:45:43 MDT
2002

Installation was successful
```

25. Enter:

```
cmssvc
```

The system displays the CMS Services menu.

26. Enter **3** to select the `run_cms` option.

27. Enter **1** to turn on CMS

28. Perform a Maintenance Restore using the most recent CMS Maintenance backup. See the “Maintenance” chapter in *CentreVu CMS R3V8 Administration*, 585-210-910.

Re-establishing the mirrors on a E3500 system

To re-establish the mirrors on a E3500 system:

1. Enter:

```
cd /dev/md/dsk
```

2. Enter:

```
ln -s ../../../../devices/pseudo/md@0,17,blk d17
```

3. Add the Solstice DiskSuite directories to your path by entering the following commands:

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

```
export PATH
```

4. Check the disk partitioning by entering:

```
olds -mirrored -check_disks
```

The system displays the following message:

```
.
.
Success, checking md.tab.new
```

5. Create DiskSuite files by entering:

```
olds - mirrored -mk_files
```

The system displays the following message:

```
.  
.br/>Success, creating md.tab.new
```

6. Activate the state database by entering:

```
olds -mirrored -metadbs
```

The system displays the following message:

```
.  
.br/>Success, setting up metadb replicas
```

⇒ NOTE:

If the system does not indicate that the metadbs commands executed successfully, continue with Steps 7 through 10. Otherwise, if the command did execute correctly, proceed to Step 12.

7. If the command executed in the preceding step failed, enter:

```
metadb -i
```

The system displays all of the unique physical disk partition numbers.

```
flags      first blk  block count  
a m p luo  16        1034       /dev/dsk/c0t0d0s0  
a   p luo  16        1034       /dev/dsk/c0t0d0s1  
a   p luo  16        1034       /dev/dsk/c0t1d0s0  
.br/>.
```

8. Record all of the unique physical disk partition numbers displayed in the output for use in the next step.

Example:

In the output shown in Step 7, you would record c0t0d0s0, c0t0d0s1, and c0t1d0s0.

9. Enter:

```
metadb -d -f /dev/dsk/boot_disk_slice next_disk_slice...
```

Where ***boot_disk_slice*** and ***next_disk_slice*** are the physical device partition numbers you recorded in Step 8.

Example:

Using the output obtained in Step 7, the command would be entered

Restoring a mirrored system from a CMSADM backup

as:

```
metadb -d -f /dev/dsk/c0t0d0s0 c0t0d0s1 c0t1d0s0
```

10. Enter the following commands:

```
metadetach d13 d12
```

```
metadetach d17 d15
```

```
metadetach d12 d19
```

```
metaclear d13 d11 d12 d17 d15 d16 d21 d19 d20
```

Enter metadvice numbers in the exact sequence shown above. You may safely ignore any error messages that may be displayed.

11. Repeat Steps 4 through 7 as needed.
12. After Steps 4 through 7, have been successfully completed, set up the root mirror. Enter:

```
olds -mirrored -setroot
```

The system displays the following message:

```
.\n.\nSuccess, root mirrored successfully.
```

13. Enter:

```
metastat | grep %
```

The system will return to the prompt without displaying any output only after the syncing process is complete. Wait until the metadvice are finished synchronizing, then continue with Step 14. Do **NOT** proceed until synchronizing is finished.

14. Comment out the d21 entry for /cms from the /etc/vfstab file with the text editor:

a. Enter: **vi /etc/vfstab**

The vfstab file displays.

```
#device      device      mount      FS      fsck      mount      mount
#to mount    to fsck     point      type     pass     at boot  options
#
#/dev/dsk/c1d0s2 /dev/rdisk/c1d0s2 /usr      ufs      1        yes      -
fd      -      /dev/fd fd      -      no      -
/proc   -      /proc  proc   -      no      -
/dev/dsk/c0t0d0s4 -      -      swap   -      no      -
/dev/md/dsk/d13 /dev/md/rdisk/d13 /         ufs     1       no      -
/dev/md/dsk/d21 /dev/md/rdisk/d21 /cms     ufs     2       yes     -
```

b. After commenting out the d21 entry, enter: **:wq!**

15. Enter:

```
sync; sync; reboot -- -r
```

⇒ NOTE:

Enter the command exactly as shown above. Also, ignore any FSCK errors displayed during the boot process.

16. Add the Solstice DiskSuite directories to your path. Enter the following commands:

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

17. Activate the */cms* metadvice. Enter:

```
olds -mirrored -setup
```

The system displays the following message:

```
.
.
14550512, 14585216, 14619920, 14654624, 14689328, 147240
14758736, 14793440, 14828144, 14862848, 14897552, 149299
ufs fsck: sanity check: /dev/md/rdisk/d21 okay
Success, /cms mirrored successfully
```

18. Mount */cms* by entering:

```
mount /cms
```

19. Enter:

```
cd /
```

20. To restore CMS, perform one of the following actions:

- If you have only one backup tape, enter the following command:

```
cpio -icmudv -C 10240 -I /dev/rmt/dev# "cms/*"
```

Where **dev#** is substituted for the device name.

- If you have more than one backup tape, enter the following command on a single line at the command prompt:

```
cpio -icmudv -C 10240 -I /dev/rmt/dev# -M
"Remove current tape, insert tape number %d,
and press Enter" "cms/*"
```

Where **dev#** is the device name.

Restoring a mirrored system from a CMSADM backup

21. Enter:

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

The system reboots.

22. Choose one of the following procedures:

- Set up CMS interactively. See “Setting Up CMS Interactively from a Terminal” in Chapter 2, of *CentreVu CMS R3V8 Software Installation, Maintenance and Troubleshooting* 585-210-941.

⇒ NOTE:

CMS setup information was created when CMS was initially setup. The Unix flat file will not contain any updates made after the initial system setup.

- Set up CMS from a Unix flat file.
 - a. To setup CMS from the flat file, enter the following command on a single line at the command prompt:

```
cp /cms/install/logdir/cms.inst.out  
/cms/install/cms_install/cms.install
```

b. Enter:

```
chmod 644 /cms/install/cms_install/cms.install
```

c. Enter:

```
cmssvc
```

Restoring a mirrored system from a CMSADM backup

The system displays the CMS Services Menu.

```
Failure in sql 'database' (-329, -2)
```

```
>>> CMS is already off <<<
```

```
CentreVu(R) Call Management System Services Menu
```

```
Select a command from the list below.
```

```
 1) auth_display Display feature authorizations
 2) auth_set     Authorize capabilities/capacities
 3) run_cms     Turn CentreVu CMS on or off
 4) setup       Set up the initial configuration
 5) swinfo      Display switch information
 6) swsetup     Change switch information
 7) patch_inst  Install a single CMS patch from CD
 8) patch_rmv   Backout an installed CMS patch
 9) load_all    Install all CMS patches found on CD
10) back_all    Backout all installed CMS patches from machine
Enter choice (1-10) or q to quit:
```

d. Enter 4 to select the setup option.

The system displays the following message:

```
Select the language for this server:
```

```
All languages are ISO Latin except Japanese. Selection of the
server language assumes that existing customer data is
compatible. (Upgrade from any ISO Latin language to any ISO
Latin language or from Japanese to Japanese is supported).
```

```
1) English
2) Dutch
3) French
4) German
5) Italian
6) Portuguese
7) Spanish
8) Japanese
```

```
Enter choice (1-8): (default: 1)
```

Restoring a mirrored system from a CMSADM backup

e. Enter the number for the language used on the system.

The system displays the following message:

```
The input will be read from
 1) the terminal
 2) a flat file
Enter choice (1-2):
```

f. Enter **2** to select the flat file option.

The system displays the following message:

```
*** The rest of this command is running in the background ***
```

g. Enter:

```
tail -f /cms/install/logdir/admin.log
```

⇒ NOTE:

The `-f` option in the `tail` command updates the console as messages are written to the `admin.log` file. All failure messages are logged in this file.

You may edit the CMS flat file and add comments about the packages that were installed or authorized.

Restoring a mirrored system from a CMSADM backup

The system displays the following message:

```
01350 Mon Nov 6 12:19:24 2000 SRC_ERR_NUM=-00329
PROCESS=pre_cms_env PID=000482 Sql.c:00071
SEVERITY=INFO ACD=-01 startdb
```

```
01308 Mon Nov 6 12:19:24 2000 SRC_ERR_NUM=-00329
PROCESS=pre_cms_env PID=000482 getlang.c:00027
SEVERITY=INFO ACD=-01 Select::execute
```

```
CMS setup started Mon Nov 6 12:19:26 MST 2000
Failure in sql 'database' (-349, -2)
Failure in sql 'drop database' (-329, -2)
Failure in sql 'drop database' (-329, -2)
Schema upgrade started Mon Nov 6 12:19:28 MST 2000
Executing schema files from /cms/install/dbinit/cms:
r3v5ai.a.1 r3v5ai.a.2 r3v6as.a.2 r3v6as.a.3 r3v8aa.c.1
r3v8aa.d.0 r3v8aa.d.1 r3v8aa.i.0 r3v8ab.d.0 r3v8ab.d.1
r3v8ab.e.0 r3v8ae.a.0 r3v8ae.a.1 r3v8ae.a.2 r3v8ae.a.3
r3v8ae.f.0 r3v8af.c.0 r3v8af.e.0 r3v8ag.f.0 r3v8ag.f.1
r3v8ai.b.0
```

```
Customer CMS data successfully initialized. Mon Nov 6
12:22:29 MST 2000
```

```
Executing schema files from /cms/install/dbinit/vec:
r3v5ai.a.5 r3v5ai.a.6 r3v6as.a.5 r3v6as.a.6 r3v6as.a.7
r3v8ae.d.1
```

Approximate space requirements:

Data Type	Allocation
VDNs	1425509
Splits	888602
Trunks	623270
Agents	540396
Trunk groups	328199
Vectors	309940
Agent login/logout records	6390
Agents trace records	4900
Exceptions records	250
Call work codes	26
Call records	8
Forecasting	0

File systems/space available:

```
/cms 12196974
```

File systems/current blocks free:

```
/cms 12196974
```

/cms:

```
VDN,SPLIT,TRUNK,AGENT,TKGRP,VECTOR,AGENT_LOG_REC,AGENT_TRAC
E_REC,EXCEPTIONS_REC,WORKCODE,CALL_REC,
```

```
Number of calls to fill_fs():12
```

```
Setup completed successfully Mon Nov 6 12:24:20 MST 2000
```

h. Press **Delete** to break out of the `tail -f` command.

Restoring a mirrored system from a CMSADM backup

23. Enter:

```
cp /cms/install/cms_install/fp.install /cms/db
```

The system will display one of two possible messages:

- If feature packages were installed on the system, the system displays the command prompt (#).
- If no feature packages were installed on the system, the system displays the following message:

```
cp: cannot access /cms/install/cms_install/fp.install
```

The **fp.install** file will be present on the tape backup if feature packages were installed. If the system cannot access the file, you will have to manually set up the feature packages that were previously authorized and installed on the system.

24. Choose one of the following procedures:

- If *no* feature packages were installed on the system, go to Step 26.
- If feature packages were installed on the system, perform the following procedure:

a. Enter:

```
. /opt/informix/bin/setenv
```

b. Enter:

```
export PATH=$PATH:/cms/install/bin
```

c. Enter the following command on a single line at the command prompt:

```
cp -p /cms/install/cms_install/fp.install  
/cms/db/fp.install
```

d. Enter the following command to install the feature packages that were present on the system before the restore:

```
/cms/install/bin/ins_proc -f
```

25. The system displays a message similar to the following example:

```
Feature package installation begun Mon Apr  8 13:45:40 MDT
2002
Feature package installation begun.
Forecasting package begin installation Mon Apr  8 13:45:43 MDT
2002

Installation was successful
```

26. Enter:

```
cmssvc
```

The system displays the CMS Services menu.

27. Enter **3** to select the `run_cms` option.

28. Enter **1** to turn on CMS

29. Perform a Maintenance Restore using the most recent CMS Maintenance backup. For details, see the “Maintenance” chapter in *CentreVu CMS R3V8 Administration*, 585-210-910.

State Database Replicas

Creating replicas

1. Verify that the `/etc/opt/SUNWmd/md.tab` file is correct (see [“Examples of valid system files” on page 117](#)).
2. Enter the `metadb` command:

```
metadb -a -c4 -f /dev/dsk/c0t0d0s1
```

```
metadb -a -c3 -f /dev/dsk/c0t1d0s1
```

Replacing replicas

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

1. Remove all existing replicas with the following command:

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

 **WARNING:**

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

2. Recreate two replicas on the same partition:

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

Maintaining the chkDisks crontab

Overview

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.

Contents

“Maintaining the chkDisks crontab” contains the following procedures:

- [“Activating chkDisks” on page 115](#)
 - [“Verifying chkDisks” on page 115](#)
 - [“Changing the scheduled run time” on page 116](#)
 - [“Canceling chkDisks” on page 116](#)
-

Activating chkDisks

Activate chkDisks by doing the following steps:

1. Enter: `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

1. To verify that `cron` is running, enter:

```
crontab -l
```

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

Changing the scheduled run time

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

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

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

Canceling chkDisks

1. To stop cron from running, enter:

```
crontab -e
```

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

Examples of valid system files

Valid *vfstab* files

Valid *vfstab* file for a non mirrored R3V8 CMS system:

```
#device          device          mount   FS   fsck   mount
mount
#to mount        to
fsck             point   type   pass  at boot  options
#
#/dev/dsk/c1d0s2 /dev/rdisk/c1d0s2 /usr    ufs   1     yes   -
fd              -          /dev/fd fd     -     no    -
/proc           -          /proc   proc  -     no    -
/dev/dsk/c0t0d0s4 -          -       swap  -     no    -
/dev/dsk/c0t0d0s0 /dev/rdisk/c0t0d0s0 /        ufs   1     no    -
swap            -          /tmp    tmpfs  -     yes   -
/dev/md/dsk/d19 /dev/md/rdisk/d19 /cms    ufs   2     yes
```

Valid *vfstab* file for an R3V8 CMS system with mirrored disks:

```
#device          device          mount   FS   fsck   mount
mount
#to mount        to
fsck             point   type   pass  at boot  options
#
fd              -          /dev/fd fd     -     no    -
/proc           -          /proc   proc  -     no    -
/dev/md/dsk/d17 -          -       swap  -     no    -
/dev/md/dsk/d13 /dev/md/rdisk/d13 /        ufs   1     no    -
/dev/md/dsk/d21 /dev/md/rdisk/d21 /cms    ufs   2     yes
```

Valid *md.tab* files

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

```
-----  
#state database replicas  
mddb00 c0t0d0s1  
mddb01 c0t1d0s0  
  
#/cms  
d19 3 1 /dev/rdisk/c0t0d0s3 1 /dev/rdisk/c0t1d0s1 1 /dev/rdisk/c1t0d0s1  
d21 -m d19  
-----
```

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

```
-----  
#state database replicas  
mddb00 c0t0d0s1  
mddb01 c1t4d0s1  
mddb02 c0t1d0s0  
mddb03 c1t5d0s0  
  
#metaroot  
d11 1 1 c0t0d0s0  
d12 1 1 c1t4d0s0  
d13 -m d11  
  
#/cms  
d19 2 1 /dev/rdisk/c0t0d0s3 1 /dev/rdisk/c0t1d0s1  
d20 2 1 /dev/rdisk/c1t4d0s3 1 /dev/rdisk/c1t5d0s1  
d21 -m d19  
-----
```

Troubleshooting

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

Contents

“Troubleshooting” contains the following procedures:

- “Solstice DiskSuite problems” on page 119
- “The system fails to recognize all disk drives” on page 121
- “Fixing “needs maintenance” messages” on page 123
- “Boot problems” on page 125
- “Log files” on page 127
- “Error messages” on page 128

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. Enter: `umount /cms`

The system unmounts */cms*.

2. Enter the command: `ls -l /cms`

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

- a. Enter: `mount /cms`

The system remounts */cms*.

- b. Enter: `df -k`

The system verifies that space is available.

Excessively long resync

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

The system fails to recognize all disk drives

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

1. Reboot the system by entering the following command:

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

The system reboots and displays the `ok` prompt:

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

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

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

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

- a. Enter: `n` (to prevent the probe from continuing).
- b. Enter: `setenv auto-boot? false`
(to prevent a reboot when you do the next command)
- c. Enter: `reset-all` (it may take a minute to complete).
- d. Depending on your system, enter one of the following commands:
 - `ok probe-scsi-all <- Enterprise 3000`
 - `ok probe-fcal-all <- Enterprise 3500`

The system fails to recognize all disk drives

The system displays a message similar to the following:

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

The specific response varies between platforms.

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

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



WARNING:

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

- Verify that the system is recognizing all its disk drives.
- Enter the following command from the `ok` prompt:

```
boot -r
```

The system reboots and displays the following message:

```
.
.
.
console login: root
Password:
```

- Log in as **root**.

Fixing “needs maintenance” messages

On rare occasions mirrors may get severely out of sync.

1. Verify that there are no faulty disks in the system before continuing with this procedure. A failed disk could also generate a “Needs maintenance” message. For more information, see [“Identifying a faulty disk” on page 62](#).
2. Use the `metareplace` command to force the recovery of the failed slice.

`metareplace -e mirror device`

Where ***mirror*** is the disk mirror and ***device*** is the metadevice number.

Example:

An extreme example is when a disk is removed from a powered system without first being detached with a `metadetach` command. The situation is indicated by the appearance of a *Needs maintenance* message in a *State* line of a `metastat` response.

```
# 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
    . . .

#
```

The `metareplace` command can force the recovery of the failed slice. The *Invoke* line of the `metastat` response leads the way.

Fixing “needs maintenance” messages

Using the current example, you would enter the following:

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

Boot problems

“Boot problems” contains procedures for some common problems experienced when booting a system. Most boot problems are the result of either trying to boot from the wrong device or the primary boot device is not available.

Contents

“Boot problems” contains the following procedures:

- [“Trying to boot from the wrong device” on page 125](#)
 - [“Primary boot device is not available” on page 126](#)
 - [“Excessive reboot time” on page 126](#)
-

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. Enter:

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

The system reboots.

2. From the `ok` prompt enter:

```
printenv boot-device
```

The system displays the following message:

```
boot-device device
```

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

3. To change the boot device, enter one of the following commands:
 - `setenv boot-device bootdevice2`
 - `setenv boot-device disk`

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

Primary boot device is not available

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

- To change to the secondary boot device, enter the following command from the `ok` prompt:

```
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
```

The startup scripts are modified to indicate that this system should never be configured as a router.

Log files

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

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

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

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

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

```
ls /olds/chkDisks
```

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

```
cp /cms/rdonly/olds_install/chkDisks /olds
```

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

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

```
/olds/chkDisks
```

Error messages

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

Error Messages

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

Message:	Attempt an operation on a submirror that has erred components.
Explanation:	Component needs maintenance before the operation can be carried out.
Response:	<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 style="margin-left: 40px;"># metareplace -e d13 c0t0d0s0</pre> 3. Reboot system.

Message:	CDROM must not be present, remove CD
Explanation:	In order for the command to execute, the CDROM drive must be empty.
Response:	<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.

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>Solstice DiskSuite</i> setup.
Response:	Verify that the device should not be used.

Message:	device: <i>devicename</i> will not be used
Explanation:	The device named is a CD-ROM drive or is not available and will not be used in the <i>Solstice DiskSuite</i> setup.
Response:	If the device is <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 re-added during a boot disk replacement.
Response:	Reboot the system.

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

Message:	Failed: disk 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 or is not in the execution path.
Response:	Use the entire path—as <code>pathname/command</code> —or add it to the execution 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>metadetach</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 Avaya 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> <p>The "/cms" entry has probably been mistyped.</p>
Response:	Remove the /cms directory and all files and try again.

Message:	<i>command</i> : not found
Explanation:	Command does not exist or is not in the execution path.
Response:	Use the entire path—as <code>pathname/command</code> —or add it to the execution 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>metadetach</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:	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>Solstice 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: <code>/usr/sbin/shutdown -y -g0 -i6</code> g. If the system still refuses to reboot, reload <i>Solaris</i>.

Message:	<i>filename</i> restored from <i>filebackup</i>
Explanation:	The action failed, and the md.tab file was restored from the previous version. Consequently, the configuration files reflect the previous system setup.
Response:	Determine the cause of the problem and try again.
Message:	resync in progress
Response:	The command entered cannot be executed because the mirrors are currently being resynched. Try again later. You can tell whether a resync is in progress by entering a <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>Solstice DiskSuite</i> before doing this task.
Message:	Submirror too small to attach
Explanation:	The submirrors d11 and d12 (or d19 and d20) do not match. You have forgotten to label one or more disks, or you have partitioned one or more disks incorrectly.
Response:	<ol style="list-style-type: none"> 1. Determine the specific cause of the message by examining, via the <i>format</i> and <i>partition</i> commands, the disks involved. 2. If the message was a response to a <code>metattach d13 d12</code> command, enter the following commands: <pre># metaclear d12 # metadb -d -f /dev/dsk/c0t<x>d0s1</pre> (where <x> is the second boot disk) 3. If the message was a response to a <code>metattach d21 d20</code> command, enter the following command: <pre># metaclear d20</pre> 4. Label or repartition the disks in error, as needed. 5. Restart the mirror upgrade process. Note that you need not recreate the <code>md.tab</code> entries the second time through the procedure.
Message:	syntax error
Response:	Check the syntax and usage of the command. Reenter the command, correcting syntax errors you have made.

Error Messages (Contd)

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

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

Message:	The file <i>filename</i> could not be restored.
Explanation:	The previous action failed, and the <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 <i>/cms</i> filesystem needs to be mounted
Explanation:	<i>/cms</i> must be mounted for the command to work.
Response:	Mount <i>/cms</i> 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 <code>reset-all</code> to reset the system before executing this command. Do you wish to continue?
Response:	<ol style="list-style-type: none"> 1. Enter: <code>N</code> (to prevent the probe from continuing). 2. Enter: <code>setenv auto-boot? false</code> (to keep the system from rebooting) 3. Enter: <code>reset-all</code> The reset may take a minute to complete. Once it does, you may do the <i>probe-scsi</i> or <i>probe-scsi-all</i> and perform any other boot prom level diagnostics. 4. IMPORTANT: Before you reboot again, enter: <code>setenv auto-boot? true</code> Failure to do so will cause the reboot to stop at the boot prompt.
Message:	<code>touch: /cms/db/unix_start cannot create</code>
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:	<code>trying to add two new disks, but a mirrored setup needs to be in place.</code>
Explanation:	You have tried to do an <code>olds -mirrored -setup</code> on an unmirrored system.
Response:	If you want a mirrored system, set it up as one. See the appropriate section in Chapter 2 for instructions on how to set up a mirrored system.
Message:	Unbalanced configurations not supported.
Explanation:	You tried to set up a system with an odd number of disk drives, or you tried to add disks without having rebooted via a "boot -r" command.
Response:	If necessary, add another disk drive to make the drive count even. Then reboot the system with "boot -r".

Error Messages (Contd)

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

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

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

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

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

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

Index

Symbols

/cms mirror 5

A

alternate boot device
 using 46

B

backing up
 CMSADM backup 53–54
 CMSADM restore 82–83
boot device
 creating alternate 33, 43
 using the alternate 46
boot disk designations 52
boot problems 125

C

capacity, system 13
chkDisks 115, 127
chunks 2
CMSADM backup 53–54
CMSADM restore 82–83
concatenated metadevice 2
 capacity 3
customer support xvii

D

d11 metadevice 6
d12 metadevice 6
d17 metadevice 6
d19 metadevice 5
d20 metadevice 5
d21 metadevice 5
disk drives
 adding 55
 recognition problem 121
 restoring 90
disk failure
 recovery 65
disk mirroring
 definition 1
 factory installation 34
 upgrade procedure 14
disk partitions 2
disks.type file 37
DiskSuite
 see *Solstice DiskSuite*
Documents xi
drive slots 12

E

error messages 128–136

F

factory installation procedures 34
faulty disk
 identifying 62
formatting disks 78

H

Hardware requirements 9

I

increasing system capacity 13
installation procedures, factory 34
installing mirrored systems 9

K

Korn shell 39

L

layout, mirror
 Enterprise 3000 8
 Enterprise 3500 7
log files 127

M

Mammoth 8mm tape drive 82
md.tab file 25
metadevices
 concatenated 2
 definition 1
mirror disks
 partitioning 20
mirrors
 and metadevices 4
 definition 1
 identifying 15
 installing 9
 layout 7
 maintaining 51
 required software 11

N

needs maintenance messages 123

P

partitioning	
procedure	20
printing <i>vfstab</i>	53
publications center	. xi

R

recovering a system	65
replicas	
replacing	114
setup	114
restoring a CMSADM backup	82–83
resync problem	120
root mirror	6

S

slices	2
software, required	11
<i>Solstice DiskSuite</i>	1
installation problems	119
setup scripts	39
<i>st.conf</i> file	82
standard configuration	52
state database	
definition	3
state database replicas	3
state database replicas	
replacing	114
setup	114
submirrors	4

T

tape drives	
<i>st.conf</i> file	82
technician support	xvii
three-way mirror	4
two-way mirror	4

U

upgrading	
and increasing system capacity	13
to disk mirroring	14

V

<i>vfstab</i> file	28
printing	53

