Sun Fire™ T2000 Server Disk Volume Management Guide

Sun Microsystems, Inc.
www.sun.com

Part No. 819-3801-10
February 2006, Revision A

Submit comments about this document at: http://www.sun.com/hwdocs/feedback

Please
Recycle

Contents

Preface

The *Sun Fire T2000 Server Disk Volume Management Guide* is for experienced system administrators. The guide includes general descriptive information about the Sun Fire™ T2000 server, and detailed instructions for configuring and administering disk volumes on the server. To use the information in this manual, you must have working knowledge of computer network concepts and terms, and advanced familiarity with the Solaris™ Operating System (Solaris OS).

Note – For information about changing the hardware configuration of your server, or about running diagnostics, see the *Sun Fire T2000 Server Service Manual* (819-2548-10)

How This Book Is Organized

The *Sun Fire T2000 Server Disk Volume Management Guide* provides a conceptual description of redundant arrays of independent disks (RAID), how to manage disk volumes, and how to configure hardware mirroring and striping.

Using UNIX Commands

This document might not contain information on basic UNIX® commands and procedures such as shutting down the system, booting the system, and configuring devices. See the following for this information:

- Software documentation that you received with your system

- Solaris OS documentation, which is at:

  http://docs.sun.com

Shell Prompts

Shell	Prompt
C shell	*machine-name*%
C shell superuser	*machine-name*#
Bourne shell and Korn shell	$
Bourne shell and Korn shell superuser	#

Typographic Conventions

Typeface*	Meaning	Examples
`AaBbCc123`	The names of commands, files, and directories; on-screen computer output	Edit your `.login` file. Use `ls -a` to list all files. `% You have mail.`
`AaBbCc123`	What you type, when contrasted with on-screen computer output	`% su` `Password:`
AaBbCc123	Book titles, new words or terms, words to be emphasized. Replace command-line variables with real names or values.	Read Chapter 6 in the *User's Guide*. These are called *class* options. You *must* be superuser to do this. To delete a file, type `rm` *filename*.

* The settings on your browser might differ from these settings.

Sun Fire T2000 Server Documentation

You can view and print the following manuals from the Sun documentation web site at http://www.sun.com/documentation

Title	Description	Part Number
Sun Fire T2000 Server Site Planning Data Guide	Site planning information for the Sun Fire T2000 server	819-2545-xx
Sun Fire T12000 Server Product Notes	Late breaking information about the server. The latest notes are posted at: `http://www.sun.com/documentation`	819-2544-xx
Sun Fire T2000 Server Getting Started Guide	Information about where to find documentation to get your system installed and running quickly	819-2542-xx
Sun Fire T2000 Server Installation Guide	Detailed rack mounting, cabling, power-on, and configuration information	819-2546-xx

Title	Description	Part Number
Sun Fire T2000 Server Administration Guide	How to perform administrative tasks that are specific to the Sun Fire T2000 server	819-2549-xx
Sun Fire T2000 Server Service Manual	How to run diagnostics to troubleshoot your server and how to remove and replace parts in the server	819-2548-xx
Advanced Lights Out Manager (ALOM) CMT v1.1 Guide	How to use the Advanced Lights Out Manager (ALOM) software on the Sun Fire T2000 server	819-3250-xx

Accessing Sun Documentation

You can view, print, or purchase a broad selection of Sun documentation, including localized versions, at:

http://www.sun.com/documentation

Third-Party Web Sites

Sun is not responsible for the availability of third-party web sites mentioned in this document. Sun does not endorse and is not responsible or liable for any content, advertising, products, or other materials that are available on or through such sites or resources. Sun will not be responsible or liable for any actual or alleged damage or loss caused by or in connection with the use of or reliance on any such content, goods, or services that are available on or through such sites or resources.

Contacting Sun Technical Support

If you have technical questions about this product that are not answered in this document, go to:

http://www.sun.com/support

Sun Welcomes Your Comments

Sun is interested in improving its documentation and welcomes your comments and suggestions. You can submit your comments by going to:

http://www.sun.com/hwdocs/feedback

Please include the title and part number of your document with your feedback:

Sun Fire T2000 Server Disk Volume Management Guide, part number 819-3801-10

1

Managing Disk Volumes

This document describes redundant array of independent disks (RAID) concepts, and how to configure and manage RAID disk volumes using the Sun Fire T2000 server's on-board serial attached SCSI (SAS) disk controller.

This chapter contains the following sections:

Requirements

To configure and use RAID disk volumes on the Sun Fire T2000 server, you must install the appropriate patches. For the latest information on patches for the Sun Fire T2000 server, see the *Sun Fire T2000 Server Product Notes* (part number 819-2544). Patches are available for download from `http://www.sun.com/sunsolve`. Installation procedures for patches are included in text README files that accompany the patches.

Disk Volumes

From the perspective of the Sun Fire T2000 server's on-boardon-board disk controller, *disk volumes* are logical disk devices comprising one or more complete physical disks.

Once you create a volume, the operating system uses and maintains the volume as if it were a single disk. By providing this logical volume management layer, the software overcomes the restrictions imposed by physical disk devices.

The on-board disk controller of the Sun Fire T2000 server provides for the creation of as many as two hardware RAID volumes. The controller supports either two-disk RAID 1 (integrated mirror, or IM) volumes, or two-, three- or four-disk RAID 0 (integrated stripe, or IS) volumes.

Note – Due to the volume initialization that occurs on the disk controller when a new volume is created, properties of the volume such as geometry and size are unknown. RAID volumes created using the hardware controller must be configured and labeled using `format`(1M) prior to use with the Solaris Operating System. See "To Configure and Label a Hardware RAID Volume for Use in the Solaris Operating System" on page 11, or the format(1M) man page for further details.

Volume migration (relocating all RAID volume disk members from one Sun Fire T2000 chassis to another) is not supported. If this operation must be performed, please contact Sun Service.

RAID Technology

RAID technology allows for the construction of a logical volume, made up of several physical disks, in order to provide data redundancy, increased performance, or both. The Sun Fire T2000 server's on-board disk controller supports both RAID 0 and RAID 1 volumes.

This section describes the RAID configurations supported by the on-board disk controller:

- Integrated stripe, or IS volumes (RAID 0)
- Integrated mirror, or IM volumes (RAID 1)

Integrated Stripe Volumes (RAID 0)

Integrated stripe volumes are configured by initializing the volume across two or more physical disks, and sharing the data written to the volume across each physical disk in turn, or *striping* the data across the disks.

Integrated stripe volumes provide for a logical unit (LUN) that is equal in capacity to the sum of all its member disks. For example, a three-disk IS volume configured on 72 GB drives will have a 216 GB capacity.

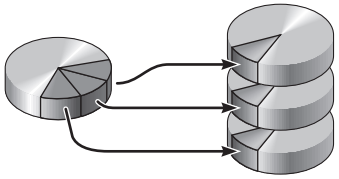

FIGURE 1-1 Graphical Representation of Disk Striping

Caution – There is no data redundancy in an IS volume configuration. Thus, if a single disk fails, the entire volume fails, and all data is lost. If an IS volume is manually deleted, all data on the volume is lost.

IS volumes are likely to provide better performance than IM volumes or single disks. Under certain workloads, particularly some write or mixed read-write workloads, I/O operations complete faster because the I/O operations are being handled in a round-robin fashion, with each sequential block being written to each member disk in turn.

Integrated Mirror Volumes (RAID 1)

Disk mirroring (RAID 1) is a technique that uses data redundancy – two complete copies of all data stored on two separate disks – to protect against loss of data due to disk failure. One logical volume is duplicated on two separate disks.

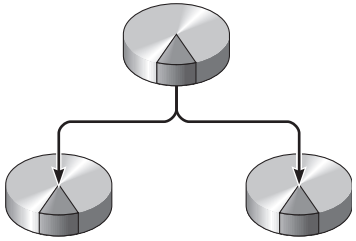

FIGURE 1-2 Graphical Representation of Disk Mirroring

Whenever the operating system needs to write to a mirrored volume, both disks are updated. The disks are maintained at all times with exactly the same information. When the operating system needs to read from the mirrored volume, it reads from whichever disk is more readily accessible at the moment, which can result in enhanced performance for read operations.

Caution – Creating RAID volumes using the on-board disk controller destroys all data on the member disks. The disk controller's volume initialization procedure reserves a portion of each physical disk for metadata and other internal information used by the controller. Once the volume initialization is complete, you can configure the volume and label it using `format`(1M). You can then use the volume in the Solaris operating system.

Hardware Raid Operations

On the Sun Fire T2000 server, the SAS controller supports mirroring and striping using the Solaris OS `raidctl` utility.

A hardware RAID volume created under the `raidctl` utility behaves slightly differently than one created using volume management software. Under a software volume, each device has its own entry in the virtual device tree, and read-write operations are performed to both virtual devices. Under hardware RAID volumes, only one device appears in the device tree. Member disk devices are invisible to the operating system, and are accessed only by the SAS controller.

Physical Disk Slot Numbers, Physical Device Names, and Logical Device Names for Non-RAID Disks

To perform a disk hot-swap procedure, you must know the physical or logical device name for the drive that you want to install or remove. If your system encounters a disk error, often you can find messages about failing or failed disks in the system console. This information is also logged in the `/var/adm/messages` files.

These error messages typically refer to a failed hard drive by its physical device name (such as `/devices/pci@1f,700000/scsi@2/sd@1,0`) or by its logical device name (such as `c0t1d0`). In addition, some applications might report a disk slot number (0 through 3).

You can use TABLE 1-1 to associate internal disk slot numbers with the logical and physical device names for each hard drive.

TABLE 1-1 Disk Slot Numbers, Logical Device Names, and Physical Device Names

Disk Slot Number	Logical Device Name*	Physical Device Name
Slot 0	c0t0d0	/devices/pci@780/pci@0/pci@9/scsi@0/sd@0,0
Slot 1	c0t1d0	/devices/pci@780/pci@0/pci@9/scsi@0/sd@1,0
Slot 2	c0t2d0	/devices/pci@780/pci@0/pci@9/scsi@0/sd@2,0
Slot 3	c0t3d0	/devices/pci@780/pci@0/pci@9/scsi@0/sd@3,0

* The logical device names might appear differently on your system, depending on the number and type of add-on disk controllers installed.

▼ To Create a Hardware Mirrored Volume

1. **Verify which hard drive corresponds with which logical device name and physical device name. using the `raidctl` command:**

```
# raidctl
No RAID volumes found.
```

See "Physical Disk Slot Numbers, Physical Device Names, and Logical Device Names for Non-RAID Disks" on page 4.

The preceding example indicates that no RAID volume exists. In another case:

```
# raidctl
RAID    Volume  RAID            RAID            Disk
Volume  Type    Status          Disk            Status
-----------------------------------------------------------
c0t0d0  IM      OK              c0t0d0          OK
                                c0t1d0          OK
```

In this example, a single IM volume has been enabled. It is fully synchronized and is online.

The Sun Fire T2000 server's on-board SAS controller can configure as many as two RAID volumes. Prior to volume creation, ensure that the member disks are available and that there are not two volumes already created.

The RAID status might be `OK`, indicating that the RAID volume is online and fully synchronized, but also might be `RESYNCING` in the event that the data between the primary and secondary member disks in an IM are still synchronizing. The RAID status might also be `DEGRADED`, if a member disk is failed or otherwise offline. Finally, it might be `FAILED`, indicating that volume should be deleted and reinitialized. This failure can occur when any member disk in an IS volume is lost, or when both disks are lost in an IM volume.

The Disk Status column displays the status of each physical disk. Each member disk might be `OK`, indicating that it is online and functioning properly, or it might be `FAILED`, `MISSING`, or otherwise `OFFLINE`, indicating that the disk has hardware or configuration issues that need to be addressed.

For example, an IM with a secondary disk that has been removed from the chassis appears as:

```
# raidctl
RAID    Volume  RAID            RAID            Disk
Volume  Type    Status          Disk            Status
-----------------------------------------------------
c0t0d0  IM      DEGRADED        c0t0d0          OK
                                c0t1d0          MISSING
```

See the `raidctl`(1M) man page for additional details regarding volume and disk status.

Note – The logical device names might appear differently on your system, depending on the number and type of add-on disk controllers installed.

---

2. **Type the following command:**

```
# raidctl -c primary secondary
```

The creation of the RAID volume is interactive, by default. For example:

```
# raidctl -c c0t0d0 c0t1d0
Creating RAID volume c0t0d0 will destroy all data on member disks,
proceed
(yes/no)? yes
Volume 'c0t0d0' created
#
```

As an alternative, you can use the `-f` option to force the creation if you are sure of the member disks, and sure that the data on both member disks can be lost. For example:

```
# raidctl -f -c c0t0d0 c0t1d0
Volume 'c0t0d0' created
#
```

When you create a RAID mirror, the secondary drive (in this case, `c0t1d0`) disappears from the Solaris device tree.

3. **To check the status of a RAID mirror, type the following command:**

```
# raidctl
RAID    Volume  RAID          RAID          Disk
Volume  Type    Status        Disk          Status
-----------------------------------------------------------
c0t0d0  1M      RESYNCING     c0t0d0        OK
                              c0t1d0        OK
```

The preceding example indicates that the RAID mirror is still re-synchronizing with the backup drive.

The following example shows that the RAID mirror is synchronized and online.

```
# raidctl
RAID    Volume  RAID            RAID            Disk
Volume  Type    Status          Disk            Status
-----------------------------------------------------
c0t0d0  IM      OK              c0t0d0          OK
                                c0t1d0          OK
```

The disk controller synchronizes IM volumes one at a time. If you create a second IM volume before the first IM volume completes its synchronization, the first volume's RAID status will indicate `RESYNCING`, and the second volume's RAID status will indicate `OK`. Once the first volume has completed, its RAID status changes to `OK`, and the second volume automatically starts synchronizing, with a RAID status of `RESYNCING`.

Under RAID 1 (disk mirroring), all data is duplicated on both drives. If a disk fails, replace it with a working drive and restore the mirror. For instructions, see "To Perform a Mirrored Disk Hot-Plug Operation" on page 16.

For more information about the `raidctl` utility, see the `raidctl`(1M) man page.

▼ To Create a Hardware Mirrored Volume of the Default Boot Device

Due to the volume initialization that occurs on the disk controller when a new volume is created, the volume must be configured and labeled using the `format`(`1M`) utility prior to use with the Solaris Operating System (see "To Configure and Label a Hardware RAID Volume for Use in the Solaris Operating System" on page 11). Because of this limitation, `raidctl`(`1M`) blocks the creation of a hardware RAID volume if any of the member disks currently have a file system mounted.

This section describes the procedure required to create a hardware RAID volume containing the default boot device. Since the boot device always has a mounted file system when booted, an alternate boot medium must be employed, and the volume created in that environment. One alternate medium is a network installation image in single-user mode (refer to the *Solaris 10 Installation Guide* for information about configuring and using network-based installations).

1. **Determine which disk is the default boot device**

   From the OpenBoot `ok` prompt, type the `printenv` command, and if necessary the `devalias` command, to identify the default boot device. For example:

```
ok printenv boot-device
boot-device =         disk

ok devalias disk
disk                  /pci@780/pci@0/pci@9/scsi@0/disk@0,0
```

2. **Type the `boot net -s` command**

```
ok boot net -s
```

3. **Once the system has booted, use the `raidctl(1M)` utility to create a hardware mirrored volume, using the default boot device as the primary disk.**

   See "To Create a Hardware Mirrored Volume" on page 5. For example:

```
# raidctl -c c0t0d0 c0t1d0
Creating RAID volume c0t0d0 will destroy all data on member disks,
proceed
(yes/no)? yes
Volume  c0t0d0  created
#
```

4. **Install the volume with the Solaris Operating System using any supported method.**

   The hardware RAID volume `c0t0d0` appears as a disk to the Solaris installation program.

Note – The logical device names might appear differently on your system, depending on the number and type of add-on disk controllers installed.

▼ To Create a Hardware Striped Volume

1. **Verify which hard drive corresponds with which logical device name and physical device name.**

   See "Disk Slot Numbers, Logical Device Names, and Physical Device Names" on page 5.

To verify the current RAID configuration, type:

```
# raidctl
No RAID volumes found.
```

The preceding example indicates that no RAID volume exists.

Note – The logical device names might appear differently on your system, depending on the number and type of add-on disk controllers installed.

---

2. **Type the following command:**

```
# raidctl -c -r 0 disk1 disk2 ...
```

The creation of the RAID volume is interactive, by default. For example:

```
# raidctl -c -r 0 c0t1d0 c0t2d0 c0t3d0
Creating RAID volume c0t1d0 will destroy all data on member disks,
proceed
(yes/no)? yes
Volume 'c0t1d0' created
#
```

When you create a RAID striped volume, the other member drives (in this case, `c0t2d0` and `c0t3d0`) disappear from the Solaris device tree.

As an alternative, you can use the `-f` option to force the creation if you are sure of the member disks, and sure that the data on all other member disks can be lost. For example:

```
# raidctl -f -c -r 0 c0t1d0 c0t2d0 c0t3d0
Volume 'c0t1d0' created
#
```

3. **To check the status of a RAID striped volume, type the following command:**

```
# raidctl
RAID    Volume  RAID            RAID            Disk
Volume  Type    Status          Disk            Status
------------------------------------------------------------
c0t1d0  IS      OK              c0t1d0          OK
                                c0t2d0          OK
                                c0t3d0          OK
```

The example shows that the RAID striped volume is online and functioning.

Under RAID 0 (disk striping),there is no replication of data across drives. The data is written to the RAID volume across all member disks in a round-robin fashion. If any one disk is lost, all data on the volume is lost. For this reason, RAID 0 cannot be used to ensure data integrity or availability, but can be used to increase write performance in some scenarios.

For more information about the `raidctl` utility, see the `raidctl`(1M) man page.

▼ To Configure and Label a Hardware RAID Volume for Use in the Solaris Operating System

After a creating a RAID volume using `raidctl`, use `format`(`1M`) to configure and label the volume before attempting to use it in the Solaris operating system.

1. **Start the `format` utility**

```
# format
```

The `format` utility might generate messages about corruption of the current label on the volume, which you are going to change. You can safely ignore these messages.

2. **Select the disk name that represents the RAID volume that you have configured.**
   In this example, `c0t2d0` is the logical name of the volume.

```
# format
Searching for disks...done
AVAILABLE DISK SELECTIONS:
        0. c0t0d0 <SUN72G cyl 14084 alt 2 hd 24 sec 424>
           /pci@780/pci@0/pci@9/scsi@0/sd@0,0
        1. c0t1d0 <SUN72G cyl 14084 alt 2 hd 24 sec 424>
           /pci@780/pci@0/pci@9/scsi@0/sd@1,0
        2. c0t2d0 <SUN72G cyl 14084 alt 2 hd 24 sec 424>
           /pci@780/pci@0/pci@9/scsi@0/sd@2,0
Specify disk (enter its number): 2
selecting c0t2d0
[disk formatted]
FORMAT MENU:
         disk       - select a disk
         type       - select (define) a disk type
         partition  - select (define) a partition table
         current    - describe the current disk
         format     - format and analyze the disk
         fdisk      - run the fdisk program
         repair     - repair a defective sector
         label      - write label to the disk
         analyze    - surface analysis
         defect     - defect list management
         backup     - search for backup labels
         verify     - read and display labels
         save       - save new disk/partition definitions
         inquiry    - show vendor, product and revision
         volname    - set 8-character volume name
         !<cmd>     - execute <cmd>, then return
         quit
```

3. **Type the `type` command at the `format>` prompt, then select 0 (zero) to auto configure the volume.**

   For example:

```
format> type

AVAILABLE DRIVE TYPES:
         0. Auto configure
         1. DEFAULT
         2. SUN72G
         3. SUN72G
         4. other
Specify disk type (enter its number)[3]: 0
c0t2d0: configured with capacity of 68.23GB
<LSILOGIC-LogicalVolume-3000 cyl 69866 alt 2 hd 16 sec 128>
selecting c0t2d0
[disk formatted]
```

4. **Use the `partition` command to partition, or *slice*, the volume according to your desired configuration.**

   See the `format`(`1M`) man page for additional details.

5. **Write the new label to the disk using the `label` command.**

```
format> label
Ready to label disk, continue? yes
```

6. **Verify that the new label has been written by printing the disk list using the `disk` command.**

```
format> disk

AVAILABLE DISK SELECTIONS:
        0. c0t0d0 <SUN72G cyl 14084 alt 2 hd 24 sec 424>
           /pci@780/pci@0/pci@9/scsi@0/sd@0,0
        1. c0t1d0 <SUN72G cyl 14084 alt 2 hd 24 sec 424>
           /pci@780/pci@0/pci@9/scsi@0/sd@1,0
       2. c0t2d0 <LSILOGIC-LogicalVolume-3000 cyl 69866 alt 2 hd
16 sec 128>
           /pci@780/pci@0/pci@9/scsi@0/sd@2,0
Specify disk (enter its number)[2]:
```

   Note that `c0t2d0` now has a type indicating it is an `LSILOGIC-LogicalVolume`.

7. **Exit the `format` utility.**

   The volume can now be used in the Solaris Operating System.

Note – The logical device names might appear differently on your system, depending on the number and type of add-on disk controllers installed.

▼ To Delete a Hardware RAID Volume

1. **Verify which hard drive corresponds with which logical device name and physical device name.**

   See "Disk Slot Numbers, Logical Device Names, and Physical Device Names" on page 5.

2. **Determine the name of the RAID volume, type:**

```
# raidctl
RAID    Volume  RAID            RAID            Disk
Volume  Type    Status          Disk            Status
-----------------------------------------------------
c0t0d0  IM      OK              c0t0d0          OK
                                c0t1d0          OK
```

   In this example, the RAID volume is `c0t1d0`.

Note – The logical device names might appear differently on your system, depending on the number and type of add-on disk controllers installed.

---

3. **To delete the volume, type the following command:**

```
# raidctl -d mirrored-volume
```

For example:

```
# raidctl -d c0t0d0
RAID Volume ‘c0t0d0’ deleted
```

In the event that the RAID volume is an IS volume, the deletion of the RAID volume is interactive, for example:

```
# raidctl -d c0t0d0
Deleting volume c0t0d0 will destroy all data it contains, proceed
(yes/no)? yes
Volume ’c0t0d0’ deleted.
#
```

The deletion of an IS volume results in the loss of all data that it contains. As an alternative, you can use the `-f` option to force the deletion if you are sure that you no longer need the IS volume, or the data it contains. For example:

```
# raidctl -f -d c0t0d0
Volume ’c0t0d0’ deleted.
#
```

4. **To confirm that you have deleted the RAID array, type the following command:**

```
# raidctl
```

For example:

```
# raidctl
No RAID volumes found
```

For more information, see the `raidctl`(1M) man page.

▼ To Perform a Mirrored Disk Hot-Plug Operation

1. **Verify which hard drive corresponds with which logical device name and physical device name.**

   See "Disk Slot Numbers, Logical Device Names, and Physical Device Names" on page 5.

2. **To confirm a failed disk, type the following command:**

```
# raidctl
```

   If the Disk Status is FAILED, then the drive can be removed and a new drive inserted. Upon insertion, the new disk should be OK and the volume should be RESYNCING.

   For example:

```
# raidctl
RAID    Volume  RAID            RAID            Disk
Volume  Type    Status          Disk            Status
-----------------------------------------------------------
c0t1d0  IM      DEGRADED        c0t1d0          OK
                                c0t2d0          FAILED
```

   This example indicates that the disk mirror has degraded due to a failure in disk `c0t2d0`.

Note – The logical device names might appear differently on your system, depending on the number and type of add-on disk controllers installed.

---

3. **Remove the hard drive, as described in the** *Sun Fire T2000 Server Service Manual.*

   There is no need to issue a software command to bring the drive offline when the drive has failed.

4. **Install a new hard drive, as described in the** *Sun Fire T2000 Server Service Manual.*

   The RAID utility automatically restores the data to the disk.

5. **To check the status of a RAID rebuild, type the following command:**

```
# raidctl
```

For example:

```
# raidctl
RAID    Volume  RAID            RAID            Disk
Volume  Type    Status          Disk            Status
-----------------------------------------------------------
c0t1d0  IM      RESYNCING       c0t1d0          OK
                                c0t2d0          OK
```

This example indicates that RAID volume `c0t1d0` is resynchronizing.

If you issue the command again once synchronization has completed, it indicates that the RAID mirror is finished resynchronizing and is back online:

```
# raidctl
RAID    Volume  RAID            RAID            Disk
Volume  Type    Status          Disk            Status
-----------------------------------------------------------
c0t1d0  IM      OK              c0t1d0          OK
                                c0t2d0          OK
```

For more information, see the `raidctl`(1M) man page.

▼ To Perform a Nonmirrored Disk Hot-Swap Operation

1. **Verify which hard drive corresponds with which logical device name and physical device name.**

   See "Disk Slot Numbers, Logical Device Names, and Physical Device Names" on page 5.

   Ensure that no applications or processes are accessing the hard drive.

2. **Type the following command:**

```
# cfgadm -al
```

For example:

```
# cfgadm -al
Ap_Id            Type         Receptacle   Occupant      Condition
c0               scsi-bus     connected    configured    unknown
c0::dsk/c0t0d0   disk         connected    configured    unknown
c0::dsk/c0t1d0   disk         connected    configured    unknown
c0::dsk/c0t2d0   disk         connected    configured    unknown
c0::dsk/c0t3d0   disk         connected    configured    unknown
c1               scsi-bus     connected    configured    unknown
c1::dsk/c1t0d0   CD-ROM       connected    configured    unknown
usb0/1           unknown      empty        unconfigured  ok
usb0/2           unknown      empty        unconfigured  ok
usb1/1.1         unknown      empty        unconfigured  ok
usb1/1.2         unknown      empty        unconfigured  ok
usb1/1.3         unknown      empty        unconfigured  ok
usb1/1.4         unknown      empty        unconfigured  ok
usb1/2           unknown      empty        unconfigured  ok
#
```

Note – The logical device names might appear differently on your system, depending on the number and type of add-on disk controllers installed.

The `-al` options return the status of all SCSI devices, including buses and USB devices. In this example, no USB devices are connected to the system.

Note that while you can use the Solaris OS `cfgadm install_device` and `cfgadm remove_device` commands to perform a hard drive hot-swap procedure, these commands issue the following warning message when you invoke these commands on a bus containing the system disk:

```
# cfgadm -x remove_device c0::dsk/c0t1d0
Removing SCSI device: /devices/pci@1f,4000/scsi@3/sd@1,0
This operation will suspend activity on SCSI bus: c0
Continue (yes/no)? y
dev = /devices/pci@780/pci@0/pci@9/scsi@0/sd@1,0
cfgadm: Hardware specific failure: failed to suspend:
     Resource              Information
------------------  -------------------------
/dev/dsk/c0t0d0s0   mounted filesystem "/"
/dev/dsk/c0t0d0s6   mounted filesystem "/usr"
```

This warning is issued because these commands attempt to quiesce the (SAS) SCSI bus, but the Sun Fire T2000 server firmware prevents it. This warning message can be safely ignored in the Sun Fire T2000 server, but the following step avoids this warning message altogether.

3. **Remove the hard drive from the device tree.**

   To remove the hard drive from the device tree, type the following command:

```
# cfgadm -c unconfigure Ap-Id
```

   For example:

```
# cfgadm -c unconfigure c0::dsk/c0t3d0
```

   This example removes `c0t3d0` from the device tree. The blue OK-to-Remove LED lights.

4. **Verify that the device has been removed from the device tree.**

   Type the following command:

```
# cfgadm -al
Ap_Id           Type         Receptacle   Occupant      Condition
c0              scsi-bus     connected    configured    unknown
c0::dsk/c0t0d0  disk         connected    configured    unknown
c0::dsk/c0t1d0  disk         connected    configured    unknown
c0::dsk/c0t2d0  disk         connected    configured    unknown
c0::dsk/c0t3d0  unavailable  connected    configured    unknown
c1              scsi-bus     connected    unconfigured  unknown
c1::dsk/c1t0d0  CD-ROM       connected    configured    unknown
usb0/1          unknown      empty        unconfigured  ok
usb0/2          unknown      empty        unconfigured  ok
usb1/1.1        unknown      empty        unconfigured  ok
usb1/1.2        unknown      empty        unconfigured  ok
usb1/1.3        unknown      empty        unconfigured  ok
usb1/1.4        unknown      empty        unconfigured  ok
usb1/2          unknown      empty        unconfigured  ok
#
```

   Note that `c0t3d0` is now `unavailable` and `unconfigured`. The corresponding hard drive OK-to-Remove LED is lit.

5. **Remove the hard drive, as described in the** *Sun Fire T2000 Server Service Manual***.**

   The blue OK-to-Remove LED goes out when you remove the hard drive.

6. **Install a new hard drive, as described in the** *Sun Fire T2000 Server Service Manual***.**

7. **Configure the new hard drive.**

   Type the following command:

```
# cfgadm -c configure Ap-Id
```

   For example:

```
# cfgadm -c configure c1::dsk/c0t3d0
```

   The green Activity LED flashes as the new disk at `c1t3d0` is added to the device tree.

8. **Verify that the new hard drive is in the device tree.**

   Type the following command:

```
# cfgadm -al
Ap_Id           Type         Receptacle   Occupant      Condition
c0              scsi-bus     connected    configured    unknown
c0::dsk/c0t0d0  disk         connected    configured    unknown
c0::dsk/c0t1d0  disk         connected    configured    unknown
c0::dsk/c0t2d0  disk         connected    configured    unknown
c0::dsk/c0t3d0  disk         connected    configured    unknown
c1              scsi-bus     connected    configured    unknown
c1::dsk/c1t0d0  CD-ROM       connected    configured    unknown
usb0/1          unknown      empty        unconfigured  ok
usb0/2          unknown      empty        unconfigured  ok
usb1/1.1        unknown      empty        unconfigured  ok
usb1/1.2        unknown      empty        unconfigured  ok
usb1/1.3        unknown      empty        unconfigured  ok
usb1/1.4        unknown      empty        unconfigured  ok
usb1/2          unknown      empty        unconfigured  ok
#
```

   Note that `c0t3d0` is now listed as `configured`.

Index

S