Platform Notes: Using luxadm Software

Send comments about this document to: docfeedback@sun.com

Please
Recycle

Adobe PostScript™

Contents

Tables

Preface

Platform Notes: Using luxadm Software describes how to use the `luxadm` administrative program for the Sun StorEdge™ A5x00 disk array (formerly named the Sun Enterprise Network Array™), SPARCstorage™ disk array, and Sun Fire™ 880 internal storage subsystem. These instructions are designed for an experienced system administrator.

Note – The `ssaadm` command is now linked to the `luxadm` command; the `luxadm` command has incorporated all the features of the `ssaadm` command.

Using UNIX Commands

This document may not contain information on basic UNIX® commands and procedures such as shutting down the system, booting the system, and configuring devices.

See one or more of the following for this information:

- *Solaris Handbook for Sun Peripherals*
- AnswerBook2™ online documentation for the Solaris™ operating environment
- Other software documentation that you received with your system

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	Read Chapter 6 in the *User's Guide*. These are called *class* options. You *must* be superuser to do this.
	Command-line variable; replace with a real name or value	To delete a file, type `rm` *filename*.

Shell Prompts

Shell	Prompt
C shell	*machine_name*%
C shell superuser	*machine_name*#
Bourne shell and Korn shell	$
Bourne shell and Korn shell superuser	#

Related Documentation

Title	Part Number
Sun StorEdge A5000 Installation and Service Manual	802-7573
Sun Fire 880 Server Owner's Guide	806-6592
Sun Fire 880 Service Manual	806-6597

Accessing Sun Documentation Online

The `docs.sun.com`[SM] web site enables you to access a select group of Sun technical documentation on the Web. You can browse the `docs.sun.com` archive or search for a specific book title or subject at:

`http://docs.sun.com`

Ordering Sun Documentation

Fatbrain.com, an Internet professional bookstore, stocks select product documentation from Sun Microsystems, Inc.

For a list of documents and how to order them, visit the Sun Documentation Center on Fatbrain.com at:

`http://www.fatbrain.com/documentation/sun`

Sun Welcomes Your Comments

Sun is interested in improving its documentation and welcomes your comments and suggestions. You can email your comments to Sun at:

`docfeedback@sun.com`

Please include the part number (816-1466-10) of your document in the subject line of your email.

`luxadm` Command Basics

The `luxadm` command is an administrative command for managing the Sun Enterprise Network Array (SENA) specifically the Sun StorEdge A5x00 disk array, the SPARCstorage Array (SSA), and the Sun Fire 880 internal disk arrays. The `luxadm` command performs a variety of control and query tasks, depending on the command-line arguments and options used.

This chapter describes the command syntax for `luxadm` and explains how to specify or address an individual disk or disk array.

Topics covered in this chapter include:

- "About the `luxadm` Command Syntax" on page 2
- "About Addressing a Disk or Disk Array" on page 3
    - "Addressing a Sun StorEdge A5x00 Disk or Array" on page 3
    - "Addressing a Sun Fire 880 Disk or Internal Array" on page 6
    - "Addressing a SPARCstorage Array" on page 9
    - "Addressing a SPARCstorage RSM Tray" on page 9

About the `luxadm` Command Syntax

The command line for `luxadm` must contain a subcommand. You can enter options, usually at least one enclosure name or path name, and other parameters depending on the subcommand. The basic syntax is as follows:

```
/usr/sbin/luxadm [options] subcommand [options] {enclosure[,dev]| pathname...}
```

When specifying a subcommand, you need to enter only as many characters as are required to uniquely identify the subcommand. For example, to run the `display` subcommand on an enclosure named `box1`, you could enter:

```
#luxadm disp box1
```

You specify the device with which a subcommand interacts by entering a path name. The path name is the logical or physical path of a Sun StorEdge A5x00 SENA Interface Board (IB), Sun Fire 880 enclosure services device (SES), SPARCstorage Array or RSM controller, or individual Fibre Channel-Arbitrated Loop (FC-AL) device. A path name can also be the World Wide Name (WWN) of the Sun StorEdge A5x00 IB, an individual FC-AL disk, or the Sun Fire 880 SES.

For a Sun StorEdge A5x00 SENA subsystem IB or individual disk, or a Sun Fire 880 SES or individual disk, you can also specify a device by enclosure name and an optional identifier for a particular disk in the enclosure. See "About Addressing a Disk or Disk Array" on page 3 for more information.

About Addressing a Disk or Disk Array

This section explains how to specify, or address, a disk or disk array.

Addressing a Sun StorEdge A5x00 Disk or Array

There are several ways to specify an individual disk or an array to `luxadm`. You can specify the path name, the World Wide Name (WWN), or the enclosure name.

Path Name

The path name is the complete physical path name or logical path name to a device or array. A logical link to the physical path to a Sun StorEdge A5x00 array (and all SENA IBs on the system) is kept in the directory `/dev/es`.

World Wide Name

The WWN is a unique 16-digit hexadecimal value that is programmed into each device during manufacturing. You can use the WWN in place of a path name to select an individual FC-AL disk or an array.

Enclosure Name

You can specify a device by its enclosure name and slot number, as follows:

box_name `,[f|r]`*slot_number*

TABLE 1-1   Addressing a Disk in a Sun StorEdge A5x00 Disk Array by Enclosure Name

Options/Arguments	Description
box_name	The name of the Sun StorEdge A5x00 enclosure you assign with the `enclosure_name` subcommand or the front panel module. Use the `box_name` without an optional *slot_number* to identify the Sun StorEdge A5x00 IB.
`f` or `r`	Specify the front or rear slots in the Sun StorEdge A5x00 enclosure.
slot_number	The slot number of the device in the Sun StorEdge A5x00 enclosure is either 0 to 6 or 0 to 10.

Addressing Examples for a Sun StorEdge A5x00

Specifying an Individual Disk

You can specify an individual disk in a Sun StorEdge A5x00 disk array in any of the following ways:

- Using the disk's complete physical path name. A typical physical path name for a Sun StorEdge A5x00 disk is:

```
/devices/pci@8,700000/pci@2/SUNW,qlc@4/fp@0,0/ssd@w21000020372028d0,0:c,raw
```

- Using the disk's complete logical path name. A typical logical path name for a Sun StorEdge A5x00 disk is:

```
/dev/rdsk/c2t0d4s2
```

- Using the disk's unique WWN. A typical WWN for a Sun StorEdge A5x00 disk is:

```
2200002037000f96
```

- Using the Sun StorEdge A5x00 enclosure name (*box_name*,[f|r*slot_number*]) followed by a disk slot identifier. A typical example for a Sun StorEdge A5x00 is:

```
macs,f4
```

Specifying an Entire Disk Array

> **Note –** When addressing an entire array, the path names, WWN, and enclosure name (*box_name*) specify the Sun StorEdge A5x00 IB.

You can address an entire Sun StorEdge A5x00 disk array in any of the following ways:

- Using the complete physical path name of a Sun StorEdge A5x00 IB. A typical physical path name of an IB is:

```
/devices/pci@8,700000/pci@2/SUNW,qlc@4/fp@0,0/ses@w5080020000061899,0:0
```

- Using the complete logical path name of a Sun StorEdge A5x00 IB. A typical logical path name for a Sun StorEdge A5x00 IB is:

```
/dev/es/ses1
```

- Using the unique WWN for the Sun StorEdge A5x00 IB. A typical WWN for a Sun StorEdge A5x00 IB is:

```
5080020000000599
```

- Using only the enclosure name (*box_name*) for the Sun StorEdge A5x00 IB. A typical enclosure name for a Sun StorEdge A5x00 IB is:

```
macs1
```

Addressing a Sun Fire 880 Disk or Internal Array

There are several ways to specify an individual disk or an array to `luxadm`. You can specify a path name, the World Wide Name (WWN), or the enclosure name. You can determine these values for an array or an individual disk by using the `luxadm` `probe`, `enclosure_name`, and `display` commands.

Path Name

A path name is either the complete physical or logical path name of an individual internal disk or the complete physical or logical path name to an array's SES controller. To determine a logical path name, use the `probe` subcommand. To determine a physical path name, use the `probe -p` subcommand. Logical links to the physical paths for Sun Fire 880 devices are kept in `/dev/rdsk` and `/dev/dsk`.

World Wide Name

The WWN for an array is a unique 16-digit hexadecimal value that is programmed into the firmware on the FC-AL backplane during manufacturing. Each array has its own unique WWN. When two backplanes are joined as a single array, only the WWN for the base backplane is used.

The WWN for an individual FC-AL disk is a unique 16-digit hexadecimal value that specifies either the port used to access a device or the device itself. Both ports on a disk share a single WWN. The WWN for each disk is programmed into the firmware of each device during manufacturing.

Use the `probe` command to determine the WWN of the SES. Use the `display` command to determine the WWN of each FC-AL disk in the enclosure.

Enclosure Name

You can specify a device by its enclosure name and slot number, as follows:

```
box_name,[s]slot_number
```

TABLE 1-2 Addressing a Disk in a Sun Fire 880 Disk Array by Enclosure Name

Options/Arguments	Description
box_name	The name you assign to the SES using the `luxadm` `enclosure_name` subcommand. See "Renaming a Sun StorEdge A5x00 Disk Array or Sun Fire 880 Internal Storage Array (enclosure_name)" on page 31 for more information.
s	Specifies a slot in a Sun Fire 880 enclosure.
slot_number	The slot number of the device in the Sun Fire 880 enclosure. Each slot is numbered 0 to 11 on the Sun Fire 880 enclosure.

Addressing Examples for Sun Fire 880

Specifying an Individual Disk

You can specify an individual internal disk in a Sun Fire 880 enclosure in any of the following ways:

- Using the disk's complete physical path name. A typical physical path name for a Sun Fire 880 internal disk is:

```
/devices/pci@8,600000/SUNW,qlc@2/fp@0,0/ssd@w210000203717c1e1,0:c,raw
```

- Using the disk's complete logical path name. A typical logical path name for a Sun Fire 880 internal disk is:

```
/dev/rdsk/c1t0d2s2
```

- Using the disk's unique WWN. A typical WWN for a Sun Fire 880 internal disk is:

```
2100002033842637
```

- Using the enclosure name followed by `,s` followed by a disk slot number (*box_name,sslot_number*), for example:

```
dak,s4
```

Specifying an Entire Disk Array

Note – When addressing an entire array, the path names, WWN, and enclosure name (*box_name*) specify the SES.

You can address an entire Sun Fire 880 disk array in any of the following ways:

- Using the complete physical path name to the array's SES controller. A typical physical path name is:

```
/devices/pci@8,600000/SUNW,qlc@2/fp@0,0/ses@w5080020010adbabf,0:0
```

- Using the complete logical path name to the array's SES controller. A typical logical path name is:

```
/dev/es/ses1
```

- Using the unique WWN for the FC-AL backplane. A typical WWN is:

```
5080020000000599
```

- Using only the enclosure name (*box_name*) for a Sun Fire 880 internal array, for example:

```
dak1
```

Addressing a SPARCstorage Array

When addressing the SPARCstorage Array, the path name specifies the SPARCstorage Array controller or a disk in the SPARCstorage Array. The controller name is specified by its physical name. For example:

```
/devices/.../.../SUNW,soc@3,0/SUNW,pln@axxxxxxx,xxxxxxxx:ctlr
```

You can also specify the controller name by a name of the form cN, where N is the logical controller number. The luxadm command uses the cN name to find an entry in the /dev/rdsk directory of a disk that is attached to the SPARCstorage Array controller. The /dev/rdsk entry is then used to determine the physical name of the SPARCstorage Array controller.

A disk in the SPARCstorage Array is specified by its logical or physical device name. For example:

```
/dev/rdsk/c1t0d0s2
```

or

```
/devices/.../.../SUNW,soc@3,0/SUNW,pln@axxxxxxx,xxxxxxxx/ssd@0,0:c,raw
```

See the disks(1M) and devlinks(1M) man pages for more information on logical names for disks and subsystems.

Addressing a SPARCstorage RSM Tray

When addressing the SPARCstorage RSM tray, the path name specifies the controller or a disk in the SPARCstorage RSM tray. The controller name is specified by its physical name. For example:

```
/devices/pci@8,600000/QLGC,isp@1,10000/sd@8,0:c,raw
```

You can also specify the controller name by a name of the form cN, where N is the logical controller number. The luxadm command uses the cN name to find an entry in the /dev/rdsk directory of a disk that is attached to the SPARCstorage Array controller. The /dev/rdsk entry is then used to determine the physical name of the controller.

A disk in the SPARCstorage RSM tray is specified by its logical or physical device name. For example:

```
/dev/rdsk/c2t8d0s2
```

See the disks(1M) and devlinks(1M) man pages for mdore information on logical names for disks and subsystems.

`luxadm` Subcommands

This chapter discusses the `luxadm` subcommands and is divided into the following sections:

Subcommand Support Matrix

The following table lists the basic luxadm subcommands and indicates which commands are supported on the Sun StorEdge A5x00 disk array, the SPARCstorage Array, SPARCstorage RSM trays, and the Sun Fire 880 internal storage array. See Appendix A for information about expert mode subcommands and their supported platforms.

TABLE 2-1 Subcommand Support Matrix

Subcommand	Sun StorEdge A5x00 Array	SPARCstorage Array	SPARCstorage RSM	Sun Fire 880 Internal Storage Subsystem
alarm_off		yes	yes	
alarm_on		yes	yes	
alarm_set		yes	yes	
display	yes	yes		yes
download	yes	yes		yes
enclosure_name	yes			yes
env_display			yes	
fast_write		yes		
fc_s_download		yes		
fcal_s_download	yes			
fcode_download	yes			
inquiry	yes	yes	yes	yes
insert_device	yes		yes	yes
led	yes	yes	yes	yes
led_blink	yes			yes
led_off	yes	yes		yes
led_on		yes	yes	
nvram_data		yes		
perf_statistics		yes		
power_off	yes	yes		
power_on	yes			

TABLE 2-1   Subcommand Support Matrix *(Continued)*

Subcommand	Sun StorEdge A5x00 Array	SPARCstorage Array	SPARCstorage RSM	Sun Fire 880 Internal Storage Subsystem
probe	yes			yes
purge		yes		
qlgc_s_download	yes			yes
release	yes	yes	yes	
remove_device	yes		yes	yes
replace_device			yes	
reserve	yes	yes		
set_boot_dev	n/a	n/a	n/a	
start		yes		
stop		yes		
sync_cache		yes		

Subcommand Options

The following options are supported by all subcommands:

-e   Run in expert mode. Use only if you are a qualified system
administrator who is knowledgeable about the systems you are
managing.

-v   Run in verbose mode.

Checking and Setting Disk LEDs

Checking the Current State of a Disk LED (`led`)

Use the `led` subcommand to check the current state of the yellow LED associated with a specific disk. For a Sun Fire 880 internal disk, this is the OK-to-Remove LED.

Supported on:

- Sun StorEdge A5x00
- Sun Fire 880 disks
- SPARCstorage Array
- SPARCstorage RSM

```
luxadm [ -v ] led { enclosure,dev... | pathname...}
```

TABLE 2-2 `led` Options and Arguments

Option/Argument	Description
enclosure	The *box_name* of a Sun StorEdge A5x00 IB or Sun Fire 880 SES that you assign with the `luxadm enclosure_name` subcommand. See "About Addressing a Disk or Disk Array" on page 3 for more information.
dev	The slot number of a specific disk in an enclosure. For a Sun StorEdge A5x00 device this is [,f] or [,r] and the slot number; for a Sun Fire 880 device this is [,s] and the slot number. See "About Addressing a Disk or Disk Array" on page 3 for more information.
pathname	The physical path name, logical path name, or WWN of a specific disk in an array.

Example:

```
# luxadm led /devices/pci@8,600000/SUNW,qlc@2/fp@0,0/ssd@w210000203717802c,0:c,raw
LED state is OFF for device in location: front,slot 3
#
```

Setting a Disk LED to the Blink Mode (`led_blink`)

Use the `led_blink` subcommand to make the LED associated with a specific disk blink or flash. For a Sun Fire 880 internal disk, this is the OK-to-Remove LED.

Supported on:

- Sun StorEdge A5x00
- Sun Fire 880 internal storage array

```
luxadm [ -v ] led_blink { enclosure,dev... | pathname...}
```

TABLE 2-3　　`led_blink` Options and Arguments

Option/Argument	Description
enclosure	The *box_name* of a Sun StorEdge A5x00 IB or Sun Fire 880 SES that you assign with the `luxadm enclosure_name` subcommand. See "About Addressing a Disk or Disk Array" on page 3 for more information.
dev	The slot number of a specific disk in an enclosure. For a Sun StorEdge A5x00 device this is [,f] or [,r] and the slot number; for a Sun Fire 880 device this is [,s] and the slot number. See "About Addressing a Disk or Disk Array" on page 3 for more information.
pathname	The physical path name, logical path name, or WWN of a specific disk in an array.

Example:

```
# luxadm led_blink dak,s3
LED state is BLINKING for device in location: slot 3
#
```

Turning Off a Disk LED (`led_off`)

Use the `led_off` subcommand to turn off the yellow LED associated with a specific disk. For a Sun Fire 880 internal disk, this is the OK-to-Remove LED.

Supported on:

- Sun StorEdge A5x00
- Sun Fire 880 internal storage array
- SPARCstorage Array
- SPARCstorage RSM

```
luxadm [ -v ] led_off { enclosure,dev... | pathname...}
```

TABLE 2-4   `led_off` Options and Arguments

Option/Argument	Description
enclosure	The *box_name* of a Sun StorEdge A5x00 IB or Sun Fire 880 SES that you assign with the `luxadm enclosure_name` subcommand. See "About Addressing a Disk or Disk Array" on page 3 for more information.
dev	The slot number of a specific disk in an enclosure. For a Sun StorEdge A5x00 device this is [,f] or [,r] and the slot number; for a Sun Fire 880 device this is [,s] and the slot number. See "About Addressing a Disk or Disk Array" on page 3 for more information.
pathname	The physical path name, logical path name, or WWN of a specific disk in an array.

Note – On a Sun StorEdge A5x00 disk array this may or may not cause the yellow OK-to-Remove LED to turn off or stop blinking, depending on the state of the StorEdge A5x00 disk array. Refer to the *Sun* StorEdge A5x00 *Installation and Service Manual* for details.

Example:

```
# luxadm led_off
/devices/pci@8,600000/SUNW,qlc@2/fp@0,0/ssd@w210000203717802c,0:c,raw
LED state is OFF for device in location: front,slot 0
#
```

Turning On a Disk LED (`led_on`)

Use the `led_on` subcommand to turn on the OK-to-Remove LED associated with a specific disk.

Supported on:

- SPARCstorage Array
- SPARCstorage RSM

```
luxadm [ -v ] led_on pathname
```

TABLE 2-5 `led_on` Options and Arguments

Option/Argument	Description
pathname	The physical or logical path name of a SPARCStorage Array or a RSM controller (`cN  name`)

Displaying Enclosure and Disk Information

Probing for Sun StorEdge A5x00 Disk Arrays and Sun Fire 880 Internal Storage Arrays (`probe`)

Use the `probe` subcommand to display information about all attached Sun StorEdge A5x00 disk arrays, Sun Fire 880 internal storage arrays, and individual FC-AL devices. The information displayed includes the logical path names, the World Wide Names (WWNs), and the enclosure names.

Supported on:

- Sun StorEdge A5x00
- Sun Fire 880 internal storage array
- Individual FC-AL devices

```
luxadm [ -v ] probe [-p]
```

TABLE 2-6   `probe` Options and Arguments

Option	Description
-p	Displays the physical path name

The following example shows the enclosure information displayed by the `probe` subcommand issued to a Sun Fire 880 system (DAKTEST) with an attached Sun StorEdge A5x00 disk array (DRIVEBAY).

Example:

```
# luxadm probe
Found Enclosure(s):
SUNWGS INT FCBPL    Name: DAKTEST    Node WWN:50000800208618f7
  Logical Path:/dev/es/ses0
  Logical Path:/dev/es/ses1

SENA              Name: DRIVEBAY   Node WWN:500008002000000eda0
  Logical Path:/dev/es/ses4
  Logical Path:/dev/es/ses5
#
```

Displaying Enclosure or Device Specific Data (`display`)

Use the `display` subcommand to display enclosure specific or device specific data.

Enclosure data consists of enclosure environmental sense information and status for all subsystem devices including disks. Device data consists of inquiry, capacity, and configuration information.

Supported on:

- Sun StorEdge A5x00
- Sun Fire 880 internal storage array
- Individual FC-AL devices
- SPARCstorage Array

```
luxadm display enclosure[,dev]... | pathname...
luxadm display -p pathname...
luxadm display –r enclosure[,dev]... | pathname...
luxadm display -v enclosure[,dev]... | pathname...
```

TABLE 2-7   `display` Options and Arguments

Option/Argument	Description
enclosure	The *box_name* of a Sun StorEdge A5x00 IB or Sun Fire 880 SES that you assign with the `luxadm enclosure_name` subcommand. See "About Addressing a Disk or Disk Array" on page 3 for more information.
dev	The slot number of a specific disk in an enclosure. For a Sun StorEdge A5x00 device this is [,f] or [,r] and the slot number; for a Sun Fire 880 device this is [,s] and the slot number. See "About Addressing a Disk or Disk Array" on page 3 for more information.
pathname	The physical or logical path name of a Sun StorEdge A5x00 array or individual disk, SPARCStorage Array, RSM controller (`cN name`), Sun Fire 880 individual disk or enclosure; or the WWN of a Sun Fire 880 disk or SES or a Sun StorEdge A5x00 IB or disk.
`-p`	Display performance information for the specified device or subsystem.
`-r`	Display error information for the specified device or subsystem.
`-v`	Display in verbose mode, including mode sense data.

Example: Display of Sun Fire 880 internal storage array enclosure information

```
# luxadm display DAKTEST

SUNWGS INT FCBPL
DISK STATUS
SLOT   DISKS              (Node WWN)
0      On (O.K.)          20000020371b1edd
1      On (O.K.)          20000020371b153f
2      On (O.K.)          200000203708c412
3      Not Installed
4      Not Installed
5      On (O.K.)          20000020371b1f06
6      On (O.K.)          20000020371b1fa8
7      On (O.K.)          200000203708c525
8      On (O.K.)          20000020371b1fe3
9      Not Installed
10     Not Installed
11     On (O.K.)          200000203708c468
SUBSYSTEM STATUS
FW Revision:9203   Box ID:0
Node WWN:50000800208618f7    Enclosure Name:DAKTEST
SSC100's - 0=Base Bkpln, 1=Base LoopB, 2=Exp Bkpln, 3=Exp LoopB
      SSC100 #0:    O.K.(11.A)
      SSC100 #1:    Not Installed
      SSC100 #2:    O.K.(11.A)
      SSC100 #3:    Not Installed
            Temperature Sensors - 0 Base, 1 Expansion
            0:27°C 1:24°C  (All temperatures are NORMAL.)

 Default Language is USA English, ASCII
#
```

Example: Display of Sun StorEdge A5x00 enclosure information

```
# luxadm display /dev/es/ses5


                              DISK STATUS
 SLOT    FRONT DISKS       (Node WWN)          REAR DISKS         (Node WWN)
 0       On (O.K.)      20000020370bd55a    On (O.K.)        20000020370bc45b
 1       On (O.K.)      20000020370bd634    On (O.K.)        20000020370bd23a
 2       On (O.K.)      20000020370b8b7b    On (O.K.)        20000020370bcaba
 3       On (O.K.)      20000020370bd633    On (O.K.)        20000020370bbcb0
 4       On (O.K.)      20000020370bafbe    On (O.K.)        20000020370bab00
 5       On (O.K.)      20000020370bd574    On (O.K.)        20000020370bd55e
 6       On (O.K.)      20000020370bb713    On (O.K.)        20000020370bc835
                            SUBSYSTEM STATUS
FW Revision:1.09  Box ID:0  Node WWN: 508002000000eda0  Enclosure Name:DRIVEBAY
 Power Supplies (0,2 in front, 1 in rear)
         0 O.K.(rev.-02) 1 O.K.(rev.-02) 2 O.K.(rev.-02)
 Fans (0 in front, 1 in rear)
         0 O.K.(rev.-05) 1 O.K.(rev.-00)
 ESI Interface board(IB) (A top, B bottom)
         A: O.K.(rev.-04)
                 GBIC module (1 on left, 0 on right in IB)
                 0 O.K.(mod.-01)
                 1 Not Installed
         B: O.K.(rev.-04)
                 GBIC module (1 on left, 0 on right in IB)
                 0 Not Installed
                 1 Not Installed
 Disk backplane (0 in front, 1 in rear)
         Front Backplane: O.K.(rev.-04)
           Temperature sensors (on front backplane)
           0:34ºC 1:36ºC 2:37ºC 3:36ºC 4:36ºC 5:36ºC
           6:36ºC  (All temperatures are NORMAL.)
         Rear Backplane:  O.K.(rev.-04)
           Temperature sensors (on rear backplane)
           0:37ºC 1:36ºC 2:34ºC 3:37ºC 4:36ºC 5:37ºC
           6:36ºC  (All temperatures are NORMAL.)
 Interconnect assembly
         O.K.(rev.-02)
 Loop  configuration
         Loop A is configured as a single loop.
         Loop B is configured as a single loop.
 Language       USA English

#
```

Displaying Disk Information (`inquiry`)

Use the `inquiry` subcommand to display information for a specific disk.

 Supported on:

- Sun StorEdge A5x00
- Sun Fire 880 internal storage array
- Individual FC-AL devices
- SPARCstorage Array
- SPARCstorage RSM

```
luxadm [ -v ] inquiry { enclosure[,dev]... | pathname...}
```

TABLE 2-8   `inquiry` Options and Arguments

Option/Argument	Description
enclosure	The *box_name* of a Sun StorEdge A5x00 IB or Sun Fire 880 SES that you assign with the `luxadm enclosure_name` subcommand. See "About Addressing a Disk or Disk Array" on page 3 for more information.
dev	The slot number of a specific disk in an enclosure. For a Sun StorEdge A5x00 device this is [,f] or [,r] and the slot number; for a Sun Fire 880 device this is [,s] and the slot number. See "About Addressing a Disk or Disk Array" on page 3 for more information.
pathname	The physical or logical path name of a Sun StorEdge A5x00 array or individual disk, SPARCStorage Array, RSM controller (cN name), Sun Fire 880 individual disk or enclosure; or the WWN of a Sun Fire 880 disk or SES, or a Sun StorEdge A5x00 IB or disk.

Example:

```
# luxadm inquiry macs1
INQUIRY:
  Physical path:
  /devices/pci@8,700000/pci@2/SUNW,qlc@4/fp@0,0/ses@w5080020000061899,0:0
Vendor:                         SUN
Product:                        SENA
Revision:                       1.05
Device type:                    0xd (SES device)
Removable media:                no
Medium Changer Element:         no
ISO version:                    0
ECMA version:                   0
ANSI version:                   3 (Device complies to SCSI-3)
Terminate task:                 no
Response data format:           2
Additional length:              0x7b
Command queueing:               no
                VENDOR-SPECIFIC PARAMETERS
Byte#                 Hex Value                                ASCII
51    00 00 00 00                                              ....
95    6d 61 63 73 31 00 00 00 00 00 00 00 00 00 00 00     macs1...........
      00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00     ................
#
```

Downloading Firmware and fcode

Downloading Firmware (`download`)

Use the `download` subcommand to download a PROM image to the FEPROMs on a Sun StorEdge A5x00 disk array IB or a SPARCstorage Array controller board.

Use the `download` subcommand to restore the original flash image into the flash PROMs on both backplanes on a Sun Fire 880 system.

In a Sun StorEdge A5x00 disk array or Sun Fire 880 internal storage array, when the download is complete, the disk array is reset and the downloaded code is executed.

In a SPARCstorage Array, when the download is complete, you must reset the SPARCstorage Array to execute the downloaded code.

Supported on:

- Sun StorEdge A5x00
- Sun Fire 880 internal storage array
- SPARCstorage Array

```
luxadm [ -v ] download [-s][-w WWN] [-f filename-path] enclosure...| pathname
```

⚠ **Caution –** When using the -s option, the download modifies the FEPROM in the Sun StorEdge A5x00 disk array.

Note – The -s option does not apply to the SPARCstorage Array controller, as it always writes the downloaded firmware into the FEPROM. The -s option does not apply to the Sun Fire 880 internal storage array, as it always writes the downloaded firmware into the flash memory.

TABLE 2-9   download Options and Arguments

Option/Argument	Description
-s	Saves the downloaded firmware in the FEPROM in a Sun StorEdge A5x00 disk array. If -s is not specified, the downloaded firmware will not be saved across power cycles.
-f filename	Downloads the PROM image in *filename-path*. If you do not specify the -f option, a default PROM image may be used.
	The default PROM image in a Sun StorEdge A5x00 is in the directory /usr/lib/locale/C/LC_MESSAGES and is named ibfirmware
	The default PROM image in a SPARCstorage Array is in the directory /usr/lib/firmware/ssa and is named ssafirmware.
	For the Sun Fire 880 internal storage array, the -f option is required. There is no default PROM image for the Sun Fire 880 system. You must load the image from the Supplement CD provided with the Solaris release. Once loaded, the image resides in the directory /usr/platform/SUNW,Sun-Fire-880/lib/images/int_fcbpl_fw and you can use the -f option to download the image from there.
-w WWN	This option is for the SPARCstorage Array only. See "Changing a SPARCstorage Array WWN (download)" on page 25.

Changing a SPARCstorage Array WWN (`download`)

Use the `download` subcommand to change the WWN of a SPARCstorage Array controller board.

Supported on:

■ SPARCstorage Array

```
luxadm [ -v ] download [-w WWN] pathname
```

TABLE 2-10   `download` Options and Arguments

Option/Argument	Description
pathname	A SPARCstorage Array controller.
`-w` *WWN*	Changes the World Wide Name for the SPARCstorage Array. `WWN` is a 12-digit hex number; leading zeros are required. The new SPARCstorage Array controller's image will have the least-significant 6 bytes of the 8-byte WWN modified to `WWN`.

Downloading fcode to FC25/S Host Adapters (`fc_s_download`)

Use the `fc_s_download` subcommand to download new fcode into all the FC25/S SBus cards. (This is the 25-MHz host adapter card that connects to SPARCstorage Arrays.)

The `fc_s_download` subcommand is interactive; it waits for user confirmation before downloading the fcode.

The version of the FC25/S SBus cards fcode that was released with this version of the operating system is located in the directory `usr/lib/firmware/fc_s` and is named `fc_s_fcode`.

Supported on:

■ SPARCstorage Array

/!\ **Caution –** Ensure that you download the `/usr/lib/firmware/fc_s/fc_s_fcode` file.

Caution – Only use the `fc_s_download` subcommand in single-user mode; otherwise, the FC25/S card could be reset.

```
luxadm [ -v ] fc_s_download [-F] [-f fcode-file]
```

TABLE 2-11 `fc_s_download` Options and Arguments

Option	Description
`-F`	Forcibly downloads the fcode. The subcommand expects user confirmation before the download.
`-f` *fcode-file*	The name of the file that has the new fcode. When the `fc_s_download` subcommand is invoked without the [`-f` *fcode-file*]option, the current version of the fcode in each FC25/S SBus card is printed.

Downloading fcode to FC100 Host Adapters (`fcal_s_download`)

Use the `fcal_s_download` subcommand to download new fcode into all the FC100/S SBus or FC100/P PCI host adapters or to display the current version of the fcode in each host adapter. (This is the 100-MHz host adapter card that connects to the Sun StorEdge A5x00 disk array.)

The `fcal_s_download` subcommand is interactive and waits for user confirmation before downloading the fcode.

Supported on:

■ Sun StorEdge A5x00

Caution – Ensure that you download the `/usr/lib/firmware/fc_s/fcal_s_fcode` file.

Caution – Do not attempt to download fcode to a FC100/S SBus or FC100/P PCI card that is in your boot path. Boot from another device, such as a CD-ROM, and then download the fcode.

```
luxadm [ -v ] fcal_s_download [ -f fcode-file ]
```

TABLE 2-12 `fcal_s_download` Options and Arguments

Option	Description
`-f` *fcode-file*	This is the name of the file that has the new fcode. If you invoke the `fcal_s_download` subcommand without the [`-f` *fcode-file*] option, the current version of the fcode in each FC100/S SBus card is displayed. The version of the FC100/S SBus cards fcode released with the operating system is located in the directory `usr/lib/firmware/fc_s` and is named `fcal_s_fcode`.

Downloading fcode to FC/S, FC100/S, FC100/P, and FC100/2P Host Adapters (`fcode_download`)

Use the `fcode_download` subcommand to locate the FC/S, FC100/S, FC100/P, and FC100/2P host adapter cards, download the fcode contained in the directory *dir-name* to the appropriate cards, and to display the current version of the fcode in each host adapter.

The `fcode_download` subcommand is interactive and waits for user confirmation before downloading the fcode.

Supported on:

■ Sun StorEdge A5x00

Caution – Only use the `fcode_download` subcommand in single-user mode; otherwise, the host adapter card could be reset.

```
luxadm [ -v ] fcode_download [-p] [-d dir-name]
```

TABLE 2-13 `fcode_download` Options

Option	Description
`-p`	Displays the current version of the fcode in each host adapter card. When the `-p` option is used, no download is performed.
`-d` *dir-name*	The name of the directory that contains the new fcode. When the `fcode_download` subcommand is invoked without the [`-d` dir-name] option, the default directory `usr/lib/firmware/fc_s` is used.

Downloading fcode to FC100/P and FC100/2P Host Adapters (`qlgc_s_download`)

Use the `qlgc_s_download` subcommand to download the fcode contained in the file *fcode-file* into all the FC100/P and FC100/2P PCI host adapter cards or to display the current version of the fcode in each host adapter.

The `qlgc_s_download` subcommand is interactive and waits for user confirmation before downloading the fcode.

Supported on:

- Sun StorEdge A5x00
- Sun Fire 880 internal storage array

Caution – Only use the `qlgc_s_download` subcommand in single-user mode; otherwise, the host adapter card could be reset.

```
luxadm [ -v ] qlgc_s_download [-f fcode-file]
```

TABLE 2-14   `qlgc_s_download` Options

Option	Description
-f fcode-file	The name of the file that contains the new fcode. When the `qlgc_s_download` subcommand is invoked without the [-f *fcode-file*] option, the current version of the fcode in each FC100/2P host adapter card is printed.

Accessing the Enclosure Services Card

The `env_display` and various `alarm` subcommands apply only to an Enclosure Services Card (SES) in a RSM tray in a SPARCstorage Array. The RSM tray is addressed by using the logical or physical path of the SES device or by specifying the controller followed by the tray number. The controller is addressed by `cN` or the physical path to the controller in the SPARCstorage Array.

Displaying Environmental Information (`env_display`)

Use the `env_display` subcommand to display the environmental information for a SPARCstorage Array or SPARCstorage RSM.

 Supported on:

- SPARCstorage Array
- SPARCstorage RSM

```
luxadm [ -v ] env_display { pathname | controller tray-number }
```

TABLE 2-15 `env_display` Options and Arguments

Option/Argument	Description
pathname	The path to an SES device.
controller	The path to a SPARCstorage Array controller.
tray-number	An RSM tray number. *tray-number* is valid only for an RSM tray in a SPARCstorage Array.

Disabling the Alarm (`alarm_off`)

Use the `alarm_off` subcommand to disable the audible alarm for this enclosure. When invoked without an option, the current state of audible alarm is printed.

 Supported on:

- SPARCstorage Array
- SPARCstorage RSM

```
luxadm [ -v ] alarm_off { pathname | controller tray-number }
```

TABLE 2-16 `alarm_off` Options and Arguments

Option/Argument	Description
pathname	The path to an SES device.
controller	The path to a SPARCstorage Array controller.
tray-number	An RSM tray number. *tray-number* is valid only for an RSM tray in a SPARCstorage Array.

Enabling the Alarm (`alarm_on`)

Use the `alarm_on` subcommand to enable the audible alarm for this enclosure. When invoked without an option, the current state of audible alarm is printed.

Supported on:

- SPARCstorage Array
- SPARCstorage RSM

```
luxadm [ -v ] alarm_on { pathname | controller tray-number }
```

TABLE 2-17 `alarm_on` Options and Arguments

Option/Argument	Description
pathname	The path to an SES device.
controller	The path to a SPARCstorage Array controller.
tray-number	An RSM tray number. *tray-number* is valid only for an RSM tray in a SPARCstorage Array.

Setting the Alarm (`alarm_set`)

Use the `alarm_set` subcommand to set the duration of the audible alarm to a specified number of seconds.

Supported on:

- SPARCstorage Array
- SPARCstorage RSM

```
luxadm [ -v ] alarm_set { pathname | controller tray-number } [seconds]
```

TABLE 2-18 `alarm_set` Options and Arguments

Option/Argument	Description
pathname	The path to an SES device.
controller	The path to a SPARCstorage Array controller.
tray-number	An RSM tray number. *tray-number* is valid only for an RSM tray in a SPARCstorage Array.
seconds	The number of seconds of the audible alarm.

Performing Enclosure and Disk Operations

Renaming a Sun StorEdge A5x00 Disk Array or Sun Fire 880 Internal Storage Array (`enclosure_name`)

Use the `enclosure_name` subcommand to change the enclosure name of a Sun StorEdge A5x00 array or a Sun Fire 880 enclosure.

Supported on:

- Sun StorEdge A5x00
- Sun Fire 880 internal storage array

```
luxadm [ -v ] enclosure_name new-name enclosure | pathname
```

TABLE 2-19   `enclosure_name` Options and Arguments

Option/Argument	Description
new-name	The name you assign to the enclosure. The new name must be 16 or fewer alphabetic or numeric characters. *New-name* specifies the *box_name* of the enclosure or interface board.
enclosure	The enclosure name of a Sun StorEdge A5x00 disk array or a Sun Fire 880 internal storage array. Use the `probe` subcommand to list the enclosure name.
pathname	The physical or logical path name of a Sun StorEdge A5x00 disk array or Sun Fire 880 internal storage array. Use a path name instead of `enclosure` if you do not know the enclosure name. Use the `probe` or `probe-p` command to list the path names and World Wide Name.

Collecting Performance Statistics (`perf_statistics`)

Use the `perf_statistics` subcommand to enable or disable the accumulation of performance statistics for a specific SPARCstorage Array controller.

Supported on:

- SPARCstorage Array

Note – You must enable the accumulation of performance statistics before you can use the `display -p` subcommand.

```
luxadm [ -v ] perf_statistics [ -e ] pathname
```

TABLE 2-20 `perf_statistics` Options and Arguments

Option/Argument	Description
pathname	A SPARCstorage Array controller
-e	Enables the accumulation of performance statistics

Powering Off an Enclosure or Disk Drive (`power_off`)

Use the `power_off` subcommand to set an enclosure to the power-save mode.

Note – Sun StorEdge A5x00 disk drives are not available when in the power-save mode.

When an Enclosure Services Card in a SPARCstorage Array is addressed, the RSM tray is powered off.

When a disk drive in a Sun StorEdge A5x00 is addressed, the drive is set to the drive off/unmated state. When it is set to the drive off/unmated state, the drive is spun down (stopped) and put in the bypass mode.

Supported on:

- Sun StorEdge A5x00
- SPARCstorage Array

```
luxadm [ -v ] power_off { enclosure[,dev]... | pathname...}
```

TABLE 2-21 `power_off` Options and Arguments

Option/Argument	Description
enclosure	The enclosure name of a Sun StorEdge A5x00
dev	The name of a specific disk in an enclosure
pathname	The physical or logical path name of a Sun StorEdge A5x00, SPARCstorage Array, or a specific disk in an array

Powering On an Enclosure or Disk Drive (`power_on`)

Use the `power_on` subcommand to set a drive to its normal power-on state. If you specify a Sun StorEdge A5x00 disk drive, the `power_on` subcommand sets the specified disks to the normal start-up state.

Supported on:

- Sun StorEdge A5x00

```
luxadm [ -v ] power_on { enclosure[,dev]... | pathname...}
```

TABLE 2-22 `power_on` Options and Arguments

Option	Description
enclosure	The enclosure name of a Sun StorEdge A5x00 disk array
dev	The name of a specific disk in a disk array
pathname	The physical or logical path name of a Sun StorEdge A5x00 disk array or a specific disk in a disk array

Releasing Disks (`release`)

Use the `release` subcommand to release one or more disk drives from reservation.

Supported on:

- Sun StorEdge A5x00
- SPARCstorage Array
- SPARCstorage RSM

```
luxadm [ -v ] release { pathname...}
```

TABLE 2-23 `release` Options and Arguments

Option	Description
pathname	The physical or logical path name of a Sun StorEdge A5x00 disk array, SPARCstorage Array, or a specific disk in a disk array

Reserving Disks (`reserve`)

Use the `reserve` subcommand to reserve the specified disk(s) for exclusive use by the host from which the subcommand was issued.

Supported on:

- Sun StorEdge A5x00
- SPARCstorage Array
- SPARCstorage RSM

```
luxadm [ -v ] reserve { pathname...}
```

TABLE 2-24 `reserve` Options and Arguments

Option	Description
pathname	The physical or logical path name of a Sun StorEdge A5x00 disk array, SPARCstorage Array, or a specific disk in a disk array

Setting the Boot Device Variable (`set_boot_dev`)

Use the `set_boot_dev` subcommand to set the boot device variable in the system PROM to a physical device name. The `set_boot_dev` subcommand normally runs interactively; it requests confirmation for setting the default boot device in the PROM.

Supported on:

- StorEdge A5x00
- SPARCstorage Array
- SPARCstorage RSM

```
luxadm [ -v ] set_boot_dev [ -y ] pathname
```

TABLE 2-25   `set_boot_dev` Options and Arguments

Option	Description
-y	Runs non-interactively; no confirmation is requested or required
pathname	A block special device or a mount point

Starting Disks (`start`)

Use the `start` subcommand to spin up the specified disk(s). If *pathname* specifies the SPARCstorage Array controller, this action applies to all disks in the SPARCstorage Array.

Supported on:

- SPARCstorage Array

```
luxadm [ -v ] start [ -t tray-number ] pathname
```

TABLE 2-26   `start` Options and Arguments

Option	Description
-t	Spin up all disks in the tray specified by *tray-number*
pathname	The physical or logical path name of a SPARCstorage Array controller

Stopping Disks (`stop`)

Use the `stop` subcommand to spin down one or more disks.

Supported on:

■ SPARCstorage Array

```
luxadm [ -v ] stop [ -t tray-number ] pathname
```

TABLE 2-27 `stop` Options and Arguments

Option	Description
-t	Spin down all disks in the tray specified by *tray-number*
pathname	The physical or logical path name of a SPARCstorage Array controller

Displaying and Setting NVRAM

Enabling and Disabling Fast Write (`fast_write`)

Use the `fast_write` subcommand to enable or disable the use of the NVRAM to enhance the performance of writes in the SPARCstorage Array.

Supported on:

■ SPARCstorage Array

```
luxadm [ -v ] fast_write [-s] -c -d -e pathname
```

TABLE 2-28 `fast_write` Options and Arguments

Option	Description
pathname	A SPARCstorage Array controller or an individual disk
-s	Causes the SPARCstorage Array to save the change so it will persist across power cycles
-c	Enables fast writes for synchronous writes only
-e	Enables fast writes
-d	Disables fast writes

Displaying Fast Write Data (`nvram_data`)

Use the `nvram_data` subcommand to display the amount of fast write data in the NVRAM for a specific disk. This command can only be used for an individual disk.

 Supported on:

- SPARCstorage Array

```
luxadm [ -v ] nvram_data pathname
```

TABLE 2-29 `nvram_data` Options and Arguments

Option	Description
pathname	A SPARCstorage Array controller or an individual disk

Purging Fast Write Data From NVRAM (`purge`)

Use the `purge` subcommand to purge any fast write data from NVRAM for one or more disks.

 Supported on:

- SPARCstorage Array

⚠ **Caution –** Use the `purge` subcommand with caution, usually only when a drive has failed.

```
luxadm [ -v ] purge pathname
```

TABLE 2-30   `purge` Options and Arguments

Option	Description
pathname	A SPARCstorage Array controller or an individual disk. If you specify a SPARCstorage Array controller, fast write data for all disks associated with that controller is purged.

Flushing NVRAM (`sync_cache`)

Use the `sync_cache` subcommand to flush all outstanding writes for one or more disks from NVRAM to the media. If *pathname* specifies the controller, this action applies to all disks in the SPARCstorage Array subsystem.

 Supported on:

- SPARCstorage Array

```
luxadm [ -v ] sync_cache pathname
```

TABLE 2-31   `sync_cache` Options and Arguments

Option	Description
pathname	A SPARCstorage Array controller or an individual disk. If you specify a SPARCstorage Array controller, outstanding writes for all disks associated with that controller are flushed.

Removing, Inserting, and Replacing Enclosures and Disks

This section discusses how to remove, insert, and replace disk drives, enclosures, or a chain of enclosures. For more detailed instructions on performing these operations with a Sun StorEdge A5x00 disk array or Sun Fire 880 internal storage array, see Chapter 3 and refer to the service manual for your enclosure.

Removing Devices (`remove_device`)

Use the `remove_device` subcommand to hot-plug a disk drive, enclosure, or a chain of enclosures. This subcommand interactively guides you through hot-plugging of one or more devices.

Supported on:

- Sun StorEdge A5x00
- Sun Fire 880 internal storage array
- SPARCstorage RSM

In the Sun StorEdge A5x00 disk array and the Sun Fire 880 internal storage array, the `remove_device` subcommand:

- Checks whether the device is busy and if so warns you.
- Takes the device offline (this fails if the disk drive is open).
- Informs you when you can safely remove the the device(s).
- Informs you which device to remove by blinking the yellow LED on the Sun StorEdge A5x00 disk array and turning on an OK-to-Remove LED on the Sun Fire 880 internal storage array.
- Displays a list of devices to be removed and requests confirmation.
- Removes the logical device names for the device that was removed.

In the SPARCstorage RSM, the remove_device subcommand:

- Takes the device offline.
- Quiesces the bus for buses that support quiescing.
- Informs you that you can safely replace the device.
- Requests confirmation that the device has been replaced.
- Unquiesces the bus for buses that support quiescing.
- Brings the device back online.
- Removes the logical device name for the device that was removed.

```
luxadm [ -v ] remove_device [ -F ] { enclosure,dev... | pathname... }
```

TABLE 2-32 remove_device Options and Arguments

Option/Argument	Description
-F	Forces the hot-plug operation on one or more devices even if those devices are being used by the host (and are, therefore, busy). Caution— Removing devices that are in use will cause unpredictable results. Try to hot-plug normally (without -F) first, resorting to this option only when you are sure of the consequences of overriding normal hot-plug checks.
enclosure	The *box_name* of a Sun StorEdge A5x00 IB or Sun Fire 880 SES that you assign with the luxadm enclosure_name subcommand. See "About Addressing a Disk or Disk Array" on page 3 for more information.
dev	The slot number of a specific disk in an enclosure. For a Sun StorEdge A5x00 device this is [,f] or [,r] and the slot number; for a Sun Fire 880 device this is [,s] and the slot number. See "About Addressing a Disk or Disk Array" on page 3 for more information.
pathname	The physical or logical path name of an individual disk in an array, or of a Sun StorEdge A5x00 array, SPARCStorage Array, or RSM controller (cN name). You can also specifiy the WWN of an individual Sun Fire 880 or Sun StorEdge A5x00 disk, or a Sun StorEdge A5x00 IB.

Example: Sun StorEdge A5x00 remove_device display

```
# luxadm remove_device macs1,f1


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



The list of devices which will be removed is:

  1: Box Name "macs1" frontslot 1

Please enter 'q' to Quit OR <Return> to Continue:

stopping: Drive in "macs1" front slot 1....Done
offlining: Drive in "macs1" front slot 1....Done

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

You must physically remove the device at this time. After you press a key, the
following data is displayed:

```
  Drive in Box Name "macs1" front slot 1
     Removing Logical Nodes:
     Removing c2t1d0s0
     Removing c2t1d0s1
     Removing c2t1d0s2
     Removing c2t1d0s3
     Removing c2t1d0s4
     Removing c2t1d0s5
     Removing c2t1d0s6
     Removing c2t1d0s7
  #
```

Inserting Devices (`insert_device`)

Use the `insert_device` subcommand to hot-plug a new disk drive, enclosure, or chain of enclosures. If you specify more than one enclosure, you can perform concurrent hot-plug operations on multiple buses.

 Supported on:

- Sun StorEdge A5x00
- Sun Fire 880 internal storage array
- SPARCstorage RSM

The `insert_device` subcommand interactively guides you through the hot-plug procedure of one or more devices. In the Sun StorEdge A5x00 disk array and the Sun Fire 880 internal storage array, the `insert_device` subcommand:

- Informs you when you can safely insert the device(s).
- Requests confirmation that the list(s) is as expected.
- Informs you which slot to insert the new drive into by turning on an OK-to Remove LED (Sun Fire 880 systems only).
- Creates the logical device names for the new devices.
- Displays the logical path name for the devices.

In the SPARCstorage RSM, the `insert_device` subcommand:

- Quiesces the bus for buses that support quiescing.
- Informs you that you can safely insert the device.
- Requests confirmation that the device has been inserted.
- Unquiesces the bus for buses that support quiescing.
- Creates the logical device name for the new device.

```
luxadm [ -v ] insert_device enclosure,dev...]
```

TABLE 2-33   `insert_device` Options and Arguments

Option/Argument	Description
enclosure	The *box_name* of a Sun StorEdge A5x00 IB or Sun Fire 880 SES that you assign with the `luxadm enclosure_name` subcommand. See "About Addressing a Disk or Disk Array" on page 3 for more information.
dev	The slot number of a specific disk in an enclosure. For a Sun StorEdge A5x00 device this is [,f] or [,r] and the slot number; for a Sun Fire 880 device this is [,s] and the slot number. See "About Addressing a Disk or Disk Array" on page 3 for more information.

Example: Sun StorEdge A5x00 insert_device display

```
# luxadm insert_device, macs1,f1

The list of devices which will be inserted is:
  1: Box Name "macs1" front slot 1

Please enter 'q' to Quit or <Return> to Continue:


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

You must physically install the disk drive at this time. After pressing any key, the following data is displayed:

```
 Drive in Box Name "macs1" front slot 1
  Logical Nodes under /dev/dsk and /dev/rdsk :
    c2t1d0s0
    c2t1d0s1
    c2t1d0s2
    c2t1d0s3
    c2t1d0s4
    c2t1d0s5
    c2t1d0s6
    c2t1d0s7
#
```

Replacing Devices (`replace_device`)

Use the `replace_device` subcommand to hot-plug a device.

The `replace_device` subcommand interactively guides you through the hot- plug procedure of one or more devices. The `replace_device` subcommand:

- Takes the device offline.
- Quiesces the bus for buses that support quiescing.
- Informs you that you can safely replace the device.
- Requests confirmation that the device has been replaced.
- Unquiesces the bus for buses that support quiescing.
- Brings the device back online.

Supported on:

■ SPARCstorage RSM

```
luxadm [ -v ] replace_device [ -F ] pathname
```

TABLE 2-34 replace_device Options and Arguments

Option/Argument	Description
-F	Forces the hot-plug operation on one or more devices even if those devices are being used by the host (and are, therefore, busy). Caution— Replacing devices which are in use will cause unpredictable results. Try to hot-plug normally (without -F) first, resorting to this option only when you are sure of the consequences of overriding normal hot-plug checks.
pathname	A SPARCstorage Array controller or an individual disk.

Hot-Plug Procedures for FC-AL Disks and Disk Arrays

Hot-plugging is the process of installing or removing an individual FC-AL disk drive or an entire Sun StorEdge A5x00 enclosure while the power is on and the operating system is running. This chapter describes how to hot-plug individual FC-AL disk drives installed in a Sun StorEdge A5x00 disk array or in a Sun Fire 880 internal storage subsystem.

This chapter covers hot-plug procedures for systems running UNIX File System (UFS) operations, VERITAS Volume Manager, or Solstice DiskSuite software.

Caution – As with other products that have high reliability, availability, and serviceability (RAS), you should not randomly remove disk drives. If the drive is active, you must stop all activity before removing it. You can do this without bringing down the operating system or powering down the unit; however, there are software considerations that you must take into account. Follow the procedures in this chapter when removing, replacing, or adding disk drives.

This chapter covers the following topics and procedures:

About Hot-Plugging FC-AL Disks and Disk Arrays

Three specific cases exist where the hot-plug feature is useful:

- Adding a disk drive to a system to increase storage capacity. See "How to Add an FC-AL Disk Drive" on page 50.
- Replacing a faulty disk drive while the system is running. See "How to Replace an FC-AL Disk Drive" on page 66.
- Removing a disk drive from a system that no longer needs it. See "How to Remove an FC-AL Disk Drive" on page 62.

The way in which you hot-plug a disk drive depends on the application you are using. Each application is different, but each requires that you:

- Prepare the disk drive for removal. This means stopping activity to the disk drive and unconfiguring it from the operating environment. See "How to Prepare an FC-AL Drive for Removal" on page 56.
- Remove the disk drive and either install a replacement drive or leave the slot empty. See "How to Replace an FC-AL Disk Drive" on page 66, "How to Add an FC-AL Disk Drive" on page 50, or "How to Remove an FC-AL Disk Drive" on page 62.
- Reconfigure the operating environment to use the new drive. See "How to Reconfigure an FC-AL Disk Drive" on page 68.

Identifying a Faulty Drive

Different applications provide various levels of error logging. In general, you can find messages about failing or failed disks in your system console window. The information is also logged in the `/usr/adm/messages` file. See the documentation that came with your application for more information.

Preparing Spare Drives

If possible, prepare replacement disk drives in advance. Format, label, and partition each replacement disk drive in the same way as the disk it will replace. See the documentation for your application for instructions on how to format and partition the disk and add that disk to your application.

Adding, Removing, and Replacing Drives

The FC-AL disk hot-plug procedures use the `luxadm insert_device` and `remove_device` subcommands to add, remove, and replace disk drives. For detailed information about the syntax of these commands, see "Removing, Inserting, and Replacing Enclosures and Disks" on page 39 in this manual.

Refer to the disk enclosure's installation or service manual for details on physically adding or removing disk drives.

If you are replacing a faulty drive, install the new drive in the same slot from which you removed the faulty drive.

Finding the Logical Device Name

When you unconfigure or configure a disk drive for an application, you may need to specify the drive by using its logical device name.

The naming convention for disks attached to a host port or host adapter is **c$w$tx$d$ys$z$**. **c$w$tx$d$ys$z$** is the logical device name, where:

w corresponds to the FC-AL controller

x corresponds to the disk slot

y is the logical unit for the disk drive (always 0)

z is the slice or partition on the disk

To obtain the logical device name for a mounted disk drive, use the `df` command. Refer to the `df (1M)` man page for more information. To obtain the logical device name for an unmounted drive, use the `luxadm display` command. You can also use the `format` command. Refer to the `format(1M)` man page for more information.

Assigning a Box Name to an Enclosure

You can specify a device to the luxadm subcommands by using a path name, WWN, or enclosure name and slot number.

If you use an enclosure name, you need to assign a box name.

The enclosure name for a Sun Fire 880 SES is specified as:

```
box_name,[s]slot_number
```

The enclosure name for a Sun StorEdge A5x00 IB is specified as:

```
box_name,[f|r]slot_number
```

A *box_name* is the name you assign to the enclosure with the luxadm enclosure_name subcommand or if you are using a Sun StorEdge A5x00, the front panel module. When used without the optional *slot_number*, the *box_name* identifies the Sun StorEdge A5x00 subsystem IB or a Sun Fire 880 internal storage array.

To assign the *box_name* and the *slot_number*, follow these steps:

1. **Use the probe subcommand to determine the enclosure name, type:**

```
#luxadm probe
```

A list of all attached subsystems and disks is displayed including the logical path name, the WWNs, and the enclosure names.

2. **Use the enclosure_name subcommand to assign a** *box_name* **to the enclosure name, type:**

```
#luxadm enclosure_name new-name enclosure | pathname
```

TABLE 3-1 enclosure_name Options and Arguments

Options	Description
new-name	The name you assign to the enclosure name. The new name must be 16 or fewer alphabetic or numeric characters. *New-name* specifies the *box_name* of the enclosure or interface board.
enclosure	The enclosure name of a Sun StorEdge A5x00 disk array or a Sun Fire 880 internal storage array. Use the probe command to display the enclosure name.
pathname	The physical or logical path name of a Sun StorEdge A5x00 disk array or Sun Fire 880 internal storage array. Use a path name instead of enclosure if you do not know the enclosure name. Use the probe (or probe -p) command to display the path names and World Wide Name.

3. **Use the** display **subcommand to determine the slot number for an individual disk.**

The display command returns a list of slot numbers and WWN numbers for each disk. Use the *box_name* from Step 2 and the *slot_number* from Step 3 to specify an individual disk to a luxadm subcommand.

Example:

The following command assigns the box name dak to a Sun Fire 880 enclosure using the enclosure_name subcommand with a logical path name.

```
# luxadm enclosure_name dak /dev/es/ses1
```

How to Add an FC-AL Disk Drive

This procedure describes how to add a disk drive while the power is on and the operating system is running. Use this procedure to add a new FC-AL disk drive to a Sun Fire 880 system or to a Sun StorEdge A5x00 array.

After you install a new drive, you need to configure the file system so that the Solaris Operating Environment recognizes the new drive. If you are running Volume Manager or Solstice DiskSuite software, you need to configure your application to recognize the new drive.

Caution – You must be a qualified system administrator to perform this procedure.

Before You Begin

- Know how to specify the disk to the `luxadm insert_device` subcommand. Specifying a disk depends on whether the enclosure is a Sun Fire 880 system or a Sun StorEdge A5x00 disk array. See "About Addressing a Disk or Disk Array" on page 3 for more information.

What to Do

1. **Become superuser.**

2. **Select any available slot for the new disk drive.**

   For reference when you configure the software environment, make a note of which slot (and enclosure) you choose.

3. Determine the address for the new device.

You need to specify the new device to the `luxadm insert_device` command. To specify an individual Sun Fire 880 disk, use *box_name*[, s]*slot_number*. To specify an individual Sun StorEdge A5x00 disk, use *box_name*[, f | , r]*slot_number*. Use a box name without a slot number to specify an enclosure. To determine the box name and slot number, use the `probe`, `enclosure_name`, and `display` subcommands:

- Use the `probe` subcommand to display the enclosure name. If you are using a Sun StorEdge A5x00, you may also obtain the box name from the front panel module, if necessary.

- Use the `enclosure_name` subcommand to assign a box name to the enclosure.

- Use the `display` subcommand to determine the slot number.

For more information, see "Assigning a Box Name to an Enclosure" on page 48. For more detailed information about all of the addressing options, see "About Addressing a Disk or Disk Array" on page 3.

4. Use the `luxadm insert_device` command to insert the new device.

This command is interactive. You are informed when you can insert the new device and guided through the procedure for creating a new device entry or chain of devices.

a. Type the `luxadm insert_device` command:

```
# luxadm insert_device [enclosure,dev...]
```

where *enclosure,dev* is the box name and slot number determined in Step 3.

After you press Return, `luxadm` displays the list of device(s) to be inserted and asks you to verify that the list is correct.

The following example inserts a new drive into slot 5 of a Sun Fire 880 enclosure named dak.

```
# luxadm insert_device dak,s5
```

The following example inserts a new drive into the first slot in the front of a Sun StorEdge A5x00 array named macs1.

```
# luxadm insert_device macs1,f1
```

b. Type c **at the prompt or press Return if the list of devices to be added is correct.**

A message similar to the following is displayed.

```
# Searching directory /dev/es for links to enclosures.

Hit <Return> after inserting the devices(s)
```

c. Physically insert the new drive, then press Return.

Refer to the disk enclosure's installation or service manual for information about installing a disk drive.

The luxadm insert_device subcommand configures the drive for the Solaris Operating Environment by creating a new device entry for the drive in the /dev/dsk and /dev/rdsk directories. The new drive is assigned a WWN.

After you insert the drive and press Return, the luxadm command informs you that the disk has been inserted and displays the logical device names for the device, for example:

```
 Device dak5 inserted

Drive in Box Name "dak" slot 5
  Logical Nodes under /dev/dsk and /dev/rdsk :
        c2t5d0s0
         c2t5d0s1
         c2t5d0s2
         c2t5d0s3
         c2t5d0s4
         c2t5d0s5
         c2t5d0s6
         c2t5d0s7
```

Note – For reference when you configure the application, make a note of the logical device name (cwtxdysz) for the disk you just added. You need to enter this device name when you configure the disk drive for your application.

5. Configure the new disk drive for your application.

Continue the procedure for adding a drive by configuring the disk drive for your application. The procedure you use depends on whether your system is running UFS, VERITAS Volume Manager, or Solstice DiskSuite software. See "How to Configure a New FC-AL Disk Drive" on page 53.

How to Configure a New FC-AL Disk Drive

⚠ **Caution –** You must be a qualified system administrator to perform this procedure. Performing a hot-plug operation on an active disk drive can result in data loss or data corruption.

After you install a new disk drive into a Sun Fire 880 enclosure or a Sun StorEdge A5x00 array, you need to configure your application to accept the new drive. Each application is different. This section provides procedures for UFS, VERITAS Volume Manager, and Solstice DiskSuite software. Select the appropriate procedure for your application and follow the steps.

Note – To configure a disk drive, you need the logical device name (**c**w**t**x**d**y**s**z**)** of the new disk. The logical device name is displayed after you use the luxadm insert_device subcommand to physically install the disk.

▼ Configuring a New FC-AL Drive for UFS

1. **Become superuser.**

2. **Verify that the device label meets your requirements.**

   Use the prtvtoc command to inspect the label for your disk. To modify the label, use the format command. Refer to the prtvtoc(1M) and format(1M) man pages for more information.

3. **Select a disk slice for your UFS file system and check if it has a clean file system, type:**

```
# fsck /dev/rdsk/cwtxdysz
```

where **c***w***t***x***d***y***s***z* is the logical device name for the new disk.

For example:

```
# fsck /dev/rdsk/c1t2d0s2
```

If you get an error message, you need to use the newfs command to create a new file system on the slice, type:

```
# newfs /dev/rdsk/cwtxdysz
```

Refer to the newfs(1M) man page for more information.

4. **If necessary, create a mount point for the new file system, type:**

```
# mkdir mount_point
```

where *mount_point* is a fully qualified path name. Refer to the mount(1M) man page for more information.

5. **Mount the new file system, type:**

```
# mount mount_point
```

where: *mount_point* is the directory you created in Step 4.

6. **After you have created the file system and mount point, modify the** /etc/vfstab **file to reflect the new file system.**

See the vfstab(4) man page for more details.

The new disk is ready to be used.

▼ Configuring a New FC-AL Disk Drive for Volume Manager

1. Become superuser.

2. Configure the Volume Manager to recognize the disk drive, type:

```
# vxdctl enable
```

3. Add the new disk to a new or existing Volume Manager disk group, type:

```
# vxdiskadd cwtxdysz
```

where **c***w***t***x***d***y***s***z* is logical device name of the new disk. This command is interactive. You will be guided through the procedure for adding a new disk to Volume Manager.

Refer to the vxdiskadd (1M) man page for further details.

The disk is now ready for use with Volume Manager as part of a new volume, added to an existing volume as a plex, or to increase an existing volume. Refer to your *Sun StorEdge Volume Manager User's Guide* for more information.

4. Quit the vxdiskadd **utility.**

▼ Configuring a New FC-AL Disk Drive for Solstice DiskSuite

Refer to the Solstice DiskSuite documentation for information about configuring the new disk drive.

How to Prepare an FC-AL Drive for Removal

Before you remove a device from a Sun StorEdge A5x00 array or a Sun Fire 880 enclosure, you need to stop activity to the drive and remove the drive from the application. The way you prepare a disk drive for removal depends on whether you are using UFS, VERITAS Volume Manager, or Solstice DiskSuite software. Each application is different.

This section provides procedures for UFS, VERITAS Volume Manager, and Solstice DiskSuite software. Select the appropriate procedure for your application and follow the steps.

> **Caution –** You must be a qualified system administrator to perform this procedure. Performing a hot-plug operation on an active disk drive can result in data loss or data corruption.

▼ Preparing a Disk Drive for Removal From UFS

Use this procedure to unconfigure a disk that is being used by one or more UFS file systems.

1. **Become superuser.**

2. **Identify activities or applications attached to the device you plan to remove.**

   Commands to use are mount, showmount -a, df, and ps -ef. See the mount(1M), showmount(1M), and ps(1) man pages for more details.

   For example, where the device to be removed is c0t11d0:

```
# mount | grep c0t11d0
/export/home1 on /dev/dsk/c0t11d0s2 setuid/read/write on
# showmount -a | grep /export/home1
cinnamon:/export/home1/archive
austin:/export/home1
swlab1:/export/home1/doc
# ps -f | grep c0t11d0
root  1225   450   4 13:09:58  pts/2   0:00 grep c0t11
```

In this example, the file system /export/home1 on the faulty disk is being remotely mounted by three different systems—cinnamon, austin, and swlab1. The only process running is grep, which has finished.

3. **Stop any activity or application processes on the file systems to be unconfigured.**

4. **Back up your system.**

5. **Determine and save the partition table for the disk.**

   If you are replacing the disk and the replacement disk is the same type as the faulty disk, you can use the format command to save the partition table of the disk. Use the format save command to save a copy of the partition table to the /etc/format.dat file. This enables you to configure the replacement disk so that its layout matches the current disk.

   Refer to the format(1M) man page for more information.

6. **Unmount any file systems on the disk.**

Note – If the file systems are on a disk that is failing or has failed, the umount operation may not unmount the file systems. A large number of error messages may be displayed in the system console and in the /var directory during the umount operation. If this happens and the umount command does not complete its operation, you may have to restart the system.

For each file system, type:

```
# umount filesystem
```

filesystem is the first field for each file system returned.

For example:

```
# umount /export/home1
```

7. **Verify that the file system has been unmounted, type:**

```
# df
```

The disk is now ready to be removed or replaced. See "How to Remove an FC-AL Disk Drive" on page 62.

▼ Preparing a Disk Drive for Removal From Volume Manager

You will need the logical device name of the disk to complete this procedure.

1. Become superuser.

2. Identify the faulty disk drive.

Different applications provide various levels of error logging. In general, you can find messages about failing or failed disks in your system console window. The information is also logged in the /usr/adm/messages file. See the documentation that came with your application for more information.

3. Back up your system.

Refer to the documentation that came with your system for backup details.

4. Identify the disk media name for the disk you intend to replace, type:

```
# vxdisk list | grep cwtxdysz
```

For example, if the disk to be removed is c2t1d0, type:

```
# vxdisk list | grep c2t1d0
c2t1d0s2    sliced    disk01       rootdg        online
```

The disk media name is the third field in the output above: disk01.

You can use the vxdiskadm utility to prepare the disk for replacement.

5. Type vxdiskadm **in a shell window.**

```
# vxdiskadm
```

This operation is interactive and requires your confirmation of the operation.

6. **If you are planning to replace the disk, select the "Remove a disk for replacement" option. Otherwise select the "Remove a disk" option.**

When prompted for a disk name to replace or remove, type the disk media name. The vxdiskadm utility marks the disk for replacement and saves the subdisk information to be rebuilt on the replacement disk.

Redundant data is automatically recovered after the replacement disk has been reattached to Volume Manager. Nonredundant data is identified as unusable and must be re-created from backups.

Refer to the vxdiskadm(1M) man page for further details.

7. **Quit the** vxdiskadm **utility.**

The disk is now ready to be removed or replaced. See "How to Remove an FC-AL Disk Drive" on page 62.

▼ Preparing a Disk Drive for Removal From Solstice DiskSuite

1. **Become superuser.**

2. **Identify the disk to be replaced by examining the** /var/adm/messages **file and** metastat **output.**

3. **Use the** metadb **command to locate any local metadevice state database replicas that may have been placed on the problem disk.**

Errors may be reported for the replicas located on the failed disk. In this example, c0t1d0 is the problem device.

```
# metadb
        flags          first blk          block count
    a m    u       16            1034            /dev/dsk/c0t0d0s4
    a      u       1050          1034            /dev/dsk/c0t0d0s4
    a      u       2084          1034            /dev/dsk/c0t0d0s4
    W   pc luo     16            1034            /dev/dsk/c0t1d0s4
    W   pc luo     1050          1034            /dev/dsk/c0t1d0s4
    W   pc luo     2084          1034            /dev/dsk/c0t1d0s4
```

The output above shows three state database replicas on slice 4 of each of the local disks, c0t0d0 and c0t1d0. The *W* in the flags field of the c0t1d0s4 slice indicates that the device has write errors. Three replicas on the c0t0d0s4 slice are still good.

Caution – If, after deleting the bad state database replicas, you are left with three or fewer replicas, add more state database replicas before continuing. This will ensure that your system reboots correctly.

4. **Record the slice name where the replicas reside and the number of replicas, then delete the state database replicas.**

   The system obtains the number of replicas by counting the number of appearances of a slice in `metadb` output in Step 2. In this example, the three state database replicas that exist on `c0t1d0s4` are deleted.

   ```
   # metadb -d c0t1d0s4
   ```

5. **Locate any submirrors using slices on the problem disk and detach them.**

   a. **Use the `metastat` command to show the affected mirrors.**

   ```
    # metastat
   metastat

     d5: Mirror
        Submirror 0: d4
          State: Okay
        Submirror 1: d3
          State: Okay
        Pass: 1
        Read option: roundrobin (default)
        Write option: parallel (default)
        Size: 1213380 blocks

   d4: Submirror of d5
        State: Okay
        Size: 1213380 blocks
        Stripe 0:
           Device              Start Block  Dbase State       Hot Spare
             c1t117d0s3                  0     No    Okay
        Stripe 1:
           Device              Start Block  Dbase State       Hot Spare
             c3t112d0s3                  0     No    Okay

   ```

b. Use the `metadetach` **command to detach the submirrors identified in the previous step.**

```
# metadetach d5 d3
   d5: submirror d3 is detached
```

6. **Delete hot spares on the problem disk.**

```
  # metahs -d hsp000 c0t1d0s6
hsp000: Hotspare is deleted
```

7. **Preserve the disk label if the disk is using multiple partitions.**

   Perform this step if you are using a slice other than `s2`.

```
# prtvtoc /dev/rdsk/c2t17d0s2 > /var/tmp/c2t17d0.vtoc
```

   See the `prtvtoc` (1M) man page for more information.

8. **Use the** `metareplace` **command to replace the disk slices that are not hot spares.**

```
# metareplace d1 c2t17d0s2 c2t16d0s2
d1: device c2t17d0s2 is replaced with c2t16d0s2
```

   The disk is now ready to be removed or replaced. See "How to Remove an FC-AL Disk Drive" on page 62.

How to Remove an FC-AL Disk Drive

This procedure describes how to remove a disk drive or an entire array while the power is on and the operating system is running. Use this procedure to remove an FC-AL disk drive from a Sun Fire 880 server or a Sun StorEdge A5x00 array.

Caution – You must be a qualified system administrator to perform this procedure. Performing a hot-plug operation on an active disk drive can result in data loss or data corruption.

Before You Begin

- Identify the disk to be removed. Different applications provide various levels of error logging. In general, you can find messages about failing or failed disks in your system console window. The information is also logged in the /usr/adm/messages file(s). See the documentation that came with your application for more information.

- Know how to specify the disk to luxadm. How you specify a device depends on whether the enclosure is a Sun Fire 880 system or a Sun StorEdge disk array. See Chapter 1 for more information.

What to Do

1. **Determine an address for the disk to be removed.**

   You need to specify the device to the luxadm remove_device command by using a path name, a WWN, or a *box_name* and *slot_number*. Use the probe, enclosure_name, and display subcommands to determine an address.

   - Use the probe subcommand to display the enclosure name, logical path name and WWN of the enclosure.

   - Use the probe -p subcommand to display a physical path name of the enclosure.

   - Use the enclosure_name subcommand to assign a *box_name* to the enclosure, if needed.

   - Use the display subcommand to display the WWNs and slot numbers of individual disks.

   For more information about using a box name and slot number, see "Assigning a Box Name to an Enclosure" on page 48. For information about all of the addressing options, see "About Addressing a Disk or Disk Array" on page 3.

2. Stop any activity to the drive and unconfigure the drive from your application.

See "How to Prepare an FC-AL Drive for Removal" on page 56 and follow the steps for your application.

3. Use the luxadm remove_device **command to remove the device.**

This command is interactive. You are guided through the procedure for removing a new device entry or chain of devices. This command checks if the device is busy, makes the device go offline, and informs you that the device can be removed.

a. Type the luxadm remove_device **command:**

```
# luxadm remove_device [-F] enclosure[,dev]...| pathname...
```

where *enclosure*[,*dev*]...| *pathname*... is the address determined in Step 1.

Note – If you are running VERITAS Volume Manager or Solstice DiskSuite software, use the luxadm remove_device -F command to remove the disk drive. The -F option is required to take disks offline.

Caution – Removing devices that are in use will cause unpredictable results. Try to hot-plug the devices normally (without -F) first, resorting to this option only when you are sure of the consequences of overriding normal hot-plug checks.

After you press Return, luxadm displays a list of the devices to be removed and asks you to verify that the list is correct.

The following example shows the command to remove a drive from slot 10 in a Sun Fire 880 enclosure named newdak.

```
# luxadm remove_device newdak,s10
```

The following example shows the command to remove a disk in slot 1 in the front of a Sun StorEdge A5x00 array named macs.

```
# luxadm remove_device macs,f1
```

b. Type c **at the prompt or press Return if the list of devices to be removed is correct.**

luxadm prepares the disk(s) or enclosure(s) for removal and displays a message similar to the following:

```
Searching directory /dev/es for links to enclosures
stopping: Drive in "DAK1" slot 1....Done
offlining:Drive in "DAK1" slot 1....Done
Hit <Return> after removing the device(s).
```

Note – If a message is displayed indicating that the list of devices is being used by the host, you will need to take the devices offline. See "How to Prepare an FC-AL Drive for Removal" on page 56 and follow the steps for your application.

c. Physically remove the drive, then press Return.

The luxadm command indicates which device you can remove by the status of the LEDs.

On a Sun StorEdge 5000 array, the yellow LED on the designated disk drive(s) will be flashing. On a Sun Fire 880 enclosure, the disk's OK-to-Remove LED will light.

For a Sun Fire 880 system, you may remove the disk drive when the OK-to-Remove LED is lit. The green power LED may also be lit or blinking.

For a Sun StorEdge 5000 array, you may remove the disk drive when the OK-to-Remove LED is blinking.

⚠ **Caution –** When the OK-to Remove LED is lit on a Sun Fire 880 system or blinking on a Sun StorEdge A5x00 system, the disk is logically ready to be removed. However, the spindle will continue to rotate for 30 seconds or more. It is safe to remove the disk before it completely stops spinning if you are careful. Do not use sudden movements and do not drop the drive.

See your service manual for more information about removing a disk drive.

After you remove the disk drive and press Return, luxadm informs you that the disk has been removed and displays the logical device names for the removed device. For example, after you remove a disk from slot 10 of Sun Fire 880 enclosure, dak, and press Return, a message similar to the following is displayed:

```
Device DISK10 removed

Drive in Box Name "dak" slot 10
  Logical Nodes being removed under /dev/dsk/ and /dev/rdsk:
  Logical Nodes being removed under /dev/dsk/ and /dev/rdsk:
        c1t12d0s0
        c1t12d0s1
        c1t12d0s2
        c1t12d0s3
        c1t12d0s4
        c1t12d0s5
        c1t12d0s6
        c1t12d0s7
 #
```

This drive is now removed from the enclosure and your application.

What Next

If you are replacing the drive, go to "How to Replace an FC-AL Disk Drive" on page 66 and continue the procedure at Step 3. Otherwise, if you are running UFS, edit the /etc/vfstab file to delete any references to the removed devices. See the vfstab(4) man page for additional details.

How to Replace an FC-AL Disk Drive

This procedure describes how to replace an FC-AL disk drive while the power is on and the operating system is running. Before you remove a disk drive, you need to stop activity to the drive and remove the drive from your application. After you replace the drive you need to reconfigure the drive for your application.

Note – If you are familiar with the `luxadm` command and the procedures for hot-plugging a disk, see the quick reference checklists in Appendix B for a summary of the tasks required for disk replacement.

Caution – You must be a qualified system administrator to perform this procedure. Performing a hot-plug operation on an active disk drive can result in data loss or data corruption.

Before You Begin

- Identify the faulty disk drive. Different applications provide various levels of error logging. In general, you can find messages about failing or failed disks in your system console window. The information is also logged in the `/usr/adm/messages` file(s). See the documentation that came with your application for more information.

- Prepare a replacement disk ahead of time with the same format, label, and partition as the disk it will replace.

- Know how to specify the disk to `luxadm`. How you specify a device depends on whether the enclosure is a Sun Fire 880 system or a Sun StorEdge A5x00 disk array. See "About Addressing a Disk or Disk Array" on page 3 for more information.

What to Do

1. **Determine an address for the disk to be removed.**

   You need to specify the disk to luxadm. You can specify the disk with a path name, a WWN, or a *box_name* and *slot_number*. To determine an address, you need to use the probe, enclosure_name, and display subcommands:

   - Use the probe subcommand to display the logical path name and WWN of the enclosure.
   - Use the probe-p subcommand to display a physical path name.
   - Use the enclosure_name subcommand to assign a box name to the enclosure, if needed.
   - Use the display subcommand to display the WWN and slot numbers of individual disks.

   To specify a disk or an array by *box_name* and *slot_number*, see "Assigning a Box Name to an Enclosure" on page 48. For more detailed information about all of the addressing options, see "About Addressing a Disk or Disk Array" on page 3.

2. **Stop all activity to the drive and unconfigure the drive from your application, if you have not already done so.**

   Your system may be running UNIX File system, VERITAS Volume Manager, or Solstice DiskSuite software. You must stop activity to the disk and notify the application that you are removing the disk drive.

   See "How to Prepare an FC-AL Drive for Removal" on page 56 and follow the steps for your application.

3. **Use the luxadm remove_device command to remove the device.**

   See "How to Remove an FC-AL Disk Drive" on page 62 and follow the steps.

4. **Use the luxadm insert_device command to add the new device.**

   See "How to Add an FC-AL Disk Drive" on page 50 and follow the steps. Insert the new drive into the same slot as the one you removed.

5. **Reconfigure the disk drive for your application.**

   Continue the disk replacement procedure by reconfiguring the disk drive within your application. The procedure you use depends on whether your system is running UFS or Volume Manager or Solstice DiskSuite software. See "How to Reconfigure an FC-AL Disk Drive" on page 68.

How to Reconfigure an FC-AL Disk Drive

After you replace a faulty FC-AL disk drive, it is necessary to reconfigure the drive for the application running on your system.

This section provides procedures for UFS, VERITAS Volume Manager, and Solstice DiskSuite software. Use the reconfiguration procedure appropriate for the application running on your system.

> **Caution –** You must be a qualified system administrator to perform this procedure. Performing a hot-plug operation on an active disk drive can result in data loss and/or data corruption.

▼ Reconfiguring a Disk Drive for UFS

1. **Verify that the device's partition table satisfies the requirements of the file system(s) you intend to re-create.**

   You can use the `prtvtoc` command to inspect the label for your device. If you need to modify the label, use the `format` command. Refer to the `prtvtoc(1M)` and `format(1M)` man pages for more information.
   For example:

   ```
   # prtvtoc /dev/rdsk/cwtxdysz
   ```

   If you have saved a disk partition table using the `format` utility and the replacement disk type matches the old disk type, then you can use the `format` utility's `partition` section to configure the partition table of the replacement disk. See the `select` and `label` commands in the partition section.

   If the replacement disk is of a different type than the disk it replaced, you can use the partition size information from the previous disk to set the partition table for the replacement disk. Refer to the `prtvtoc(1M)` and `format(1M)` man pages for more information.

2. **Select a disk slice for your UFS file system and create a new file system on the slice:**

```
# newfs /dev/rdsk/cwtxdysz
```

Refer to the newfs(1M) man page for more information.

3. **Mount the new file system using the** mount **command, type:**

```
# mount mount_point
```

where *mount_point* is the directory on which the faulty disk was mounted.

The new disk is ready to be used. You can now restore data from your backups.

▼ Reconfiguring a Disk Drive for Volume Manager

To re-create the replaced disk on the new drive:

1. **Configure the Volume Manager to recognize the disk drive, type:**

```
# vxdctl enable
```

2. **Use the** vxdiskadm **utility.**

Select the "Replace a failed or removed disk" option.

vxdiskadm supplies a list of available disks to be used as replacements.

3. **Select the replacement drive.**

vxdiskadm automatically configures the replacement drive to match the failed drive.

Redundant data is recovered automatically. Space for nonredundant data is created and identified. Nonredundant data must be recovered from backing store.

▼ Reconfiguring a Disk Drive for Solstice DiskSuite

1. Restore the disk label, if necessary.

```
# cat /var/tmp/c2t17d0.vtoc | fmhard -s - /dev/rdsk/c2t17d0s2
```

2. If you deleted replicas, add the same number back to the appropriate slice. In this example, /dev/dsk/c-t1d0s4 **is used.**

```
# metadb -a c 3 c0t1d0s4
```

3. Depending on how the disk was used, you may have a variety of tasks to do.

Use the following table to decide what to do next.

TABLE 3-2   Disk Replacement Decision Table

Type of Device	Do the Following...
Slice	Use normal data recovery procedures.
Unmirrored Stripe or Concatenation	If the stripe/concat is used for a file system, run newfs(1M), mount the file system, then restore data from backup. If the stripe/concat is used as an application that uses the raw device, that application must have its own recovery procedures.
Mirror (Submirror)	Run metattach(1M) to reattach a detached submirror.
RAID5 Metadevice	Run metareplace(1M) to re-enable the slice. This causes the resyncs to start.
Trans Metadevice	Run fsck(1M) to repair the trans metadevice.

4. Replace hot spares that were deleted, and add them to the appropriate hot spare pool(s).

```
  # metahs -a hsp000 c0t0d0s6
hsp000: Hotspare is added
```

5. Validate the data.

Check the user and application data on all metadevices. You may have to run an application-level consistency checker or use some other method to check the data.

APPENDIX **A**

Expert Mode Subcommands

Caution – Use the expert mode commands only if you are a qualifed administrator who is knowledgeable about the systems you are managing.

Expert mode subcommands are listed in TABLE A-1.

TABLE A-2 lists the expert mode subcommands supported on the Sun StorEdge A5x00, SPARCstorage Array, SPARCstorage RSM trays, and Sun Fire 880 internal storage array.

The command line must contain the luxadm –e (expert mode) option and a subcommand.

```
luxadm –e subcommand pathname
```

TABLE A-1   Expert Mode Subcommands

Subcommand	Description
bus_getstate	Gets and displays the state of the specified bus or the bus controlling the specified device.
bus_quiesce	Quiesces the specified bus or the bus controlling the specified device.
bus_reset	Resets the specified bus or the bus controlling the specified device.
bus_resetall	Resets the specified bus or the bus controlling the specified device, and all devices on that bus.
bus_unquiesce	Unquiesces the specified bus or the bus controlling the specified device.
dev_getstate	Gets and displays the state of the specified device.
dev_reset	Resets the specified device.

TABLE A-1   Expert Mode Subcommands *(Continued)*

Subcommand	Description
forcelip	Forces the link to reinitialize using the Loop Initialize Primitive (LIP) sequence. This subcommand is supported on the Sun StorEdge A5x00 disk array and the Sun Fire 880 internal storage array only.
offline	Takes the specified device offline.
online	Puts the specified device online.
rdls	Reads the link error status block from a specified device. This subcommand also displays the link error status information for the host adapter associated with the specified device, if available. The rdls subcommand is supported on the Sun StorEdge A5x00 disk array and the Sun Fire 880 internal storage array only.

TABLE A-2   Expert Mode Subcommand Support Matrix

Subcommand	Sun StorEdge A5x00	SPARCstorage Array	SPARCstorage RSM	Sun Fire 880 Internal Storage Array
bus_getstate	yes	yes	yes	
bus_quiesce	yes	yes	yes	
bus_reset	yes	yes	yes	
bus_resetall	yes	yes	yes	
bus_unquiesce	yes	yes	yes	
dev_getstate	yes	yes	yes	
dev_reset	yes	yes	yes	
forcelip	yes			yes
offline	yes	yes	yes	
online	yes	yes	yes	
rdls	yes			yes

Checklists for Replacing an FC-AL Disk Drive

This section provides a summary of the tasks required to replace a disk drive on a Sun StorEdge A5x00 array or a Sun Fire 880 system running UFS, Volume Manager, or Solstice DiskSuite software. To complete the tasks listed below, you need to complete these procedures:

- "How to Prepare an FC-AL Drive for Removal" on page 56
- "How to Replace an FC-AL Disk Drive" on page 66
- "How to Reconfigure an FC-AL Disk Drive" on page 68

Caution – You must be a qualified system administrator to perform these procedures. Performing hot-plug operations on an active disk drive can result in data loss or data corruption.

Replacing a Disk Drive Using UFS

- Identify the faulty disk drive.
- Use the df, mount, showmount-a, and ps-ef commands to identify activities or applications attached to the drive you plan to replace.
- Stop any activity to the drive.
- Back up your system.
- Determine and save the partition table for the disk.
- Use the umount command to unmount any file systems on the disk.

- If your system is running Volume Manager, use the Volume Manager command `vxdisk offline` to take the disk offline. This is required even if the disk is not being managed by Volume Manager.

- Use the `luxadm probe`, `enclosure_name`, and `display` subcommands to determine an address for the drive. An address can be a path name, a WWN, or a box name and slot number.

- Use the `luxadm remove_device` subcommand to remove the disk from the Solaris Operating Environment (deleting the `/devices` and `/dev` entries). Use `remove_device -F` if you are running Solstice DiskSuite software.

- Remove the disk drive from the enclosure when prompted by the `luxadm remove_device` subcommand.

- Use the `luxadm insert_device` subcommand to add the disk to the Solaris Operating Environment (adding the disk's WWN to the `/devices` and `/dev` directories).

- Insert the disk drive into the enclosure when prompted by the `luxadm insert_device` subcommand.

- Use the `prtvtoc` and `format` commands to verify that the partition table is correct for the file system you need.

- Use the `newfs` command to create a new file system on the slice and the `mount` command to mount the new file system.

Replacing a Disk Drive Using Volume Manager

- Use `vxdisk` to identify the disk you need to replace.

- Use `vxdiskadm` to remove the disk from Volume Manager.

- Use the `luxadm probe`, `enclosure_name`, and `display` subcommands to determine an address for the drive. An address can be a path name, a WWN, or a box name and slot number.

- Use the `luxadm remove_device -F` subcommand to take the disk offline from Volume Manager and remove the disk from Solaris (deleting the `/devices` and `/dev` entries).

- Remove the disk drive from the enclosure when prompted by the `luxadm remove_device` subcommand.

- Use the `luxadm insert_device` subcommand to add the disk to the Solaris Operating Environment (adding the disks WWN to the `/devices` and `/dev` directories) and physically install the drive.

- Insert the disk drive into the enclosure when prompted by the `luxadm insert_device` subcommand.
- Use `vxdctl enable` to notify Volume Manager of the new disk.
- Use `vxdiskadm` to replace the failed disk in Volume Manager.

Replacing a Disk Drive Using Solstice DiskSuite

- Use `/var/adm/messages` and `metastat` to identify the faulty disk drive.
- Use `metadb` to locate and delete any metadevice database replicas on the faulty disk.
- Use `metastat` to locate any submirrors using slices on the faulty disk.
- Use `metadetach` to detach submirrors that use slices on the faulty disk.
- Use `metahs` to delete hot spares marked "Available" that are on the problem disk.
- Use `prtvtoc` to preserve the disk label if multiple partitions are being used on the disk.
- Use `metareplace` to replace the disk slices if they are not hot spared.
- Use the `luxadm probe`, `enclosure_name`, and `display` subcommands to determine an address for the drive. An address can be a path name, a WWN, or a box_name and slot_number.
- Use the `luxadm remove_device -F` subcommand to take the disk offline from Solstice DiskSuite software and remove the disk from the Solaris Operating Environment (deleting the `/devices` and `/dev` entries).
- Remove the disk drive from the enclosure when prompted by the `luxadm remove_device` subcommand.
- Use the `luxadm insert_device` command to add the disk to the Solaris Operating Environment (adding the disks WWN to the `/devices` and `/dev` directories) .
- Insert the disk drive into the enclosure when prompted by the `luxadm insert_device` subcommand.
- If necessary, restore the disk label.
- Use `metadb -a` to add the same number of metadevice state database replicas that were deleted.

- Use `metahs -a` to replace hot spares that were deleted to hot spare pools.
- Perform one of the tasks listed in TABLE B-1, depending on how the slice that failed was used.

TABLE B-1 Slice and command Information

Slice	Action
Simple Slice	Use normal recovery procedures.
Stripe or Concatenation	Run `newfs` for the entire metadevice; restore from backup.
Mirror	Reattach detached submirrors.
RAID5 metadevice	Resync (enable) affected slices.
Trans Metadevice	run `fsck(1M)`.

Index