Sun Fire™ T1000 Server Administration Guide

Sun Microsystems, Inc.
www.sun.com

Please
Recycle

Adobe PostScript™

Contents

Figures

Tables

Preface

The *Sun Fire™ T1000 Server Administration Guide* is written for experienced system administrators. It includes general descriptive information about the server and detailed instructions for configuring and administering the server. To use the information in this manual, you must have working knowledge of computer network concepts and terms, and advanced familiarity with the Solaris™ Operating System (Solaris OS).

How This Book Is Organized

The *Sun Fire T1000 Server Administration Guide* is divided into the following chapters:

■ Chapter 1 describes the system console and how to access it.

■ Chapter 2 describes the tools used to configure system firmware, including Sun™ Advanced Lights Out Manager (ALOM) CMT system controller environmental monitoring, automatic system recovery (ASR), and multipathing software. In addition, it describes how to unconfigure and reconfigure a device manually.

■ Chapter 3 describes redundant array of independent disks (RAID) concepts, and how to configure and manage RAID disk volumes using your server's on-board serial attached SCSI (SAS) disk controller.

This manual also includes the following appendix:

■ Appendix A provides a list of all OpenBoot™ configuration variables, and a short description of each.

Using UNIX Commands

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

- Software documentation that you received with your system
- Solaris OS documentation, which is at:

  http://docs.sun.com

Shell Prompts

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

Typographic Conventions

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

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

Related Documentation

Title	Description	Part Number
Sun Fire T1000 Server Site Planning Guide	Site planning information for the server	819-3749
Sun Fire T1000 Server Product Notes	Late-breaking information about the server. The latest notes are posted at: http://www.sun.com/documentation	819-3246
Sun Fire T1000 Server Getting Started Guide	Information about where to find documentation to get your system installed and running quickly	819-3244
Sun Fire T1000 Server Installation Guide	Detailed rackmounting, cabling, power-on, and configuration information	819-3247

Title	Description	Part Number
Sun Fire T1000 Server Service Manual	How to run diagnostics to troubleshoot your server and how to remove and replace parts in the server	819-3248
Advanced Lights Out Management (ALOM) CMT v1.3 Guide	How to use the Advanced Lights Out Manager (ALOM) software on the server	819-3250 (version 1.1) 819-6672 (version 1.2) 819-7981 (version 1.3)
Sun Fire T1000 Server Safety and Compliance Guide	Provides safety and compliance information that is specific to this server.	819-6674

Documentation, Support, and Training

Sun Function	URL
Documentation	http://www.sun.com/documentation/
Support	http://www.sun.com/support/
Training	http://www.sun.com/training/

Third-Party Web Sites

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

Sun Welcomes Your Comments

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

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

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

Sun Fire T1000 Server Administration Guide, part number 819-3249-12

Configuring the System Console

This chapter explains what the system console is, describes the different ways of configuring it on your server, and helps you understand its relation to the system controller.

Topics covered in this chapter include:

- "Communicating With the System" on page 1
- "Accessing the System Controller" on page 6
- "Switching Between the System Controller Prompt and the System Console" on page 15
- "ALOM CMT and The `sc>` Prompt" on page 16
- "OpenBoot `ok` Prompt" on page 18
- "System Console OpenBoot Configuration Variable Settings" on page 23

Communicating With the System

To install your system software or to diagnose problems, you need some way to interact at a low level with the system. The *system console* is Sun's facility for doing this. You use the system console to view messages and issue commands. There can be only one system console per computer.

The system console must be accessed through the system controller during initial system installation. After installation, you can configure the system console to accept input from and send output to different devices. TABLE 1-1 lists these devices and where they are discussed in the document.

TABLE 1-1   Ways of Communicating With the System

Devices Available	During Installation	After Installation	Further Information
A terminal server attached to the serial management port (SER MGT).			
	X	X	"Accessing the System Controller" on page 6
	X	X	"Accessing the System Console Through a Terminal Server" on page 9
	X	X	"System Console OpenBoot Configuration Variable Settings" on page 23
An alphanumeric terminal or similar device attached to the serial management port (SER MGT).			
	X	X	"Accessing the System Controller" on page 6
	X	X	"Accessing the System Console Through an Alphanumeric Terminal" on page 13
	X	X	"System Console OpenBoot Configuration Variable Settings" on page 23
A TIP line attached to the serial management port (SER MGT).			
	X	X	"Accessing the System Controller" on page 6
	X	X	"Accessing the System Console Through a TIP Connection" on page 10
		X	"Modifying the /etc/remote File" on page 12

TABLE 1-1 Ways of Communicating With the System *(Continued)*

Devices Available	During Installation	After Installation	Further Information
	X	X	"System Console OpenBoot Configuration Variable Settings" on page 23
An Ethernet line connected to the network management port (NET MGT).			
		X	"Activating the Network Management Port" on page 7

What the System Console Does

The system console displays status and error messages generated by firmware-based tests during system startup. After running those tests, you can enter special commands that affect the firmware and alter system behavior. For information about the tests that run during the boot process, refer to the service manual for your server.

Once the operating system is booted, the system console displays UNIX system messages and accepts UNIX commands. You can access the system console using the ALOM CMT `console` command.

What the System Controller Console Does

The system controller console displays the results of the ALOM CMT boot diagnostics and initialization.

If it receives no user input within 60 seconds, ALOM CMT automatically connects to the system console. To return to the system controller, type the console escape sequence `#.` (Hash-Period).

Using the System Console

To use the system console, you need to attach an I/O device to the system. Initially, you might have to configure that hardware, and load and configure appropriate software as well.

You must also ensure that the system console is directed to the appropriate port on your server's rear panel, generally, the one to which your hardware console device is attached (see FIGURE 1-1). You do this by setting the `input-device` and `output-device` OpenBoot configuration variables.
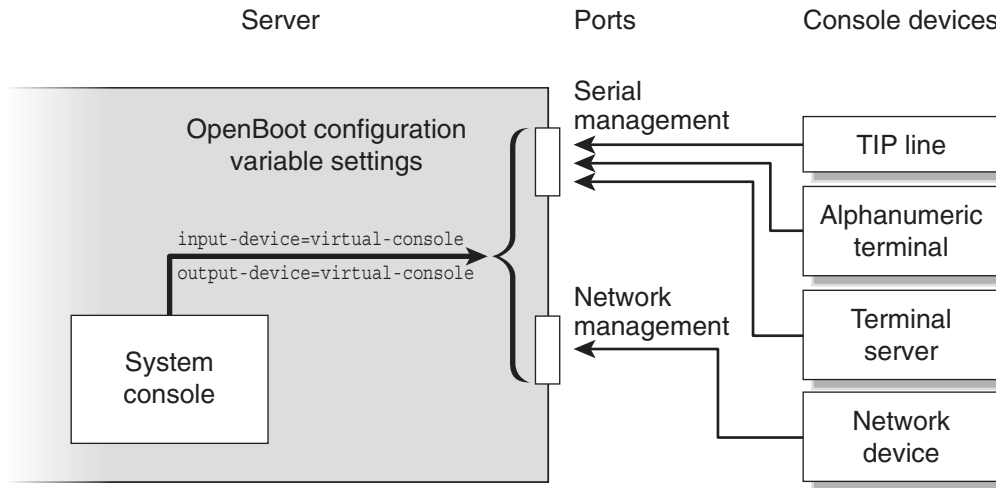
FIGURE 1-1 Directing the System Console

Default System Console Connection Through the Serial Management and Network Management Ports

On your server, the system console comes preconfigured to allow input and output only by means of ALOM CMT. ALOM CMT must be accessed either through the system controller's serial management port (SER MGT) or the network management port (NET MGT). By default, the network management port is configured to retrieve network configuration using DHCP and to allow connections using SSH. You can modify the network management port configuration after connecting to ALOM CMT through either the system controller's serial or network management ports.

Typically, you connect one of the following hardware devices to the serial management port:

- Terminal server
- Alphanumeric terminal or similar device
- TIP line connected to another computer

These constraints provide for secure access at the installation site.

TTYA
Serial Port

Net3

Net2

Net1

Net0

PCI-E Slot

SC NET MGT Port

SC SER MGT Port

FIGURE 1-2   Rear I/O Panel of the Chassis

Using a TIP line enables you to use windowing and operating system features on the system making the connection to your server.

The serial management port is not a general-purpose serial port. If you want to use a general-purpose serial port with your server, to connect a serial printer, for instance, use the regular 9-pin serial port on the rear panel of your server. The Solaris OS sees this port as TTYA.

- For instructions on accessing the system console through a terminal server, see "Accessing the System Console Through a Terminal Server" on page 9.

- For instructions on accessing the system console through an alphanumeric terminal, see "Accessing the System Console Through an Alphanumeric Terminal" on page 13.

- For instructions on accessing the system console through a TIP line, see "Accessing the System Console Through a TIP Connection" on page 10.

Once the network management port (NET MGT) has been assigned an IP address by a DHCP server, you can connect to ALOM CMT using Secure Shell (SSH). As an alternative to the (default) DHCP configuration, you can configure the network management port with a static IP address, and change the communication protocol

from SSH to Telnet. Up to eight simultaneous connections to the system controller sc> prompt are available through the network management port. For more information, see "Activating the Network Management Port" on page 7.

Accessing the System Controller

The following sections describe ways of accessing the system controller.

Using the Serial Management Port

When you are accessing ALOM CMT using a device connected to the system controller's serial management port, you will see the output of the ALOM CMT diagnostics when AC power is first applied or when the system controller has been reset. After the diagnostics have completed the serial management port is available for login.

For more information about the system controller card, refer to the ALOM CMT guide for your server.

▼ To Use the Serial Management Port

1. **Ensure that the serial port on your connecting device is set to the following parameters:**

   - 9600 baud
   - 8 bits
   - No parity
   - 1 stop bit
   - No handshaking

2. **Establish an system controller session.**

   For instructions on how to use the system controller, see the the ALOM CMT guide for your server.

Activating the Network Management Port

The network management port is configured by default to retrieve network settings using DHCP and allow connections using SSH. You may need to modify these settings for your network. If you are unable to use DHCP and SSH on your network, you must connect to ALOM CMT through the system controller using the serial management port to reconfigure the network management port. See "Using the Serial Management Port" on page 6

Note – There is no default password when connecting to the ALOM System Controller for the first time using the serial management port. When connecting to the ALOM System Controller using the network management port for the first time, the default password is the last 8 digits of the Chassis Serial Number. The Chassis Serial Number can be found printed on the back of the server or in the printed system information sheet which shipped with your server. You must assign a password during initial system configuration. For more information, refer to your server's installation guide and the ALOM CMT guide for your server.

You can assign the network management port a static IP address or you can configure the port to obtain an IP address using the Dynamic Host Configuration Protocol (DHCP) from another server. The network management port can be configured to accept connections from Telnet clients or SSH clients, but not both.

Data centers frequently devote a separate subnet to system management. If your data center has such a configuration, connect the network management port to this subnet.

Note – The network management port is a 10/100BASE-T port. The IP address assigned to the network management port is a unique IP address, separate from the main server IP address, and is dedicated for use only with ALOM CMT on the system controller.

▼ To Activate the Network Management Port

1. **Connect an Ethernet cable to the network management port.**

2. **Log in to the system controller through the serial management port.**

   For more information about connecting to the serial management port, see "Accessing the System Controller" on page 6.

3. **Type one of the following commands:**

■ **If your network uses static IP addresses, type:**

```
sc> setsc netsc_dhcp false
sc> setsc netsc_ipaddr ip-address
sc> setsc netsc_ipnetmask ip-address
sc> setsc netsc_ipgateway ip-address
```

■ **If your network uses Dynamic Host Configuration Protocol (DHCP), type:**

```
sc> setsc netsc_dhcp true
```

4. **Type one of the following commands:**

   ■ If you intend to use Secure Shell (SSH) to connect to ALOM CMT:

```
sc> setsc if_connection ssh
```

   ■ - If you intend to use Telnet to connect to ALOM CMT:

```
sc> setsc if_connection telnet
```

5. **Reset the system controller so that the new settings take affect:**

```
sc> resetsc
```

6. **After the system controller resets, log in to the system controller and issue the** shownetwork **command to verify network settings:**

```
sc> shownetwork
```

To connect through the network management port, use the telnet or ssh (based on the value you provided in Step 4) commands to the IP address you specified in Step 3 of the preceding procedure.

Accessing the System Console Through a Terminal Server

The following procedure assumes that you are accessing ALOM CMT on the system controller by connecting a terminal server to the serial management port (SER MGT) of your server.

▼ To Access The System Console Through a Terminal Server

1. **Complete the physical connection from the serial management port to your terminal server.**

   The serial management port on the server is a data terminal equipment (DTE) port. The pinouts for the serial management port correspond with the pinouts for the RJ-45 ports on the serial interface breakout cable supplied by Cisco for use with the Cisco AS2511-RJ terminal server. If you use a terminal server made by another manufacturer, check that the serial port pinouts of the server match those of the terminal server you plan to use.

   If the pinouts for the server serial ports correspond with the pinouts for the RJ-45 ports on the terminal server, you have two connection options:

   - Connect a serial interface breakout cable directly to the Sun Fire T1000 server. See "Accessing the System Controller" on page 6.
   - Connect a serial interface breakout cable to a patch panel and use the straight-through patch cable (supplied by Sun) to connect the patch panel to the server.

FIGURE 1-3   Patch Panel Connection Between a Terminal Server and Your Server

If the pinouts for the serial management port *do not* correspond with the pinouts for the RJ-45 ports on the terminal server, you need to make a crossover cable that takes each pin on the server serial management port to the corresponding pin in the terminal server's serial port.

TABLE 1-2 shows the crossovers that the cable must perform.

TABLE 1-2   Pin Crossovers for Connecting the Server to a Terminal Server

Sun Fire T1000 Serial Port (RJ-45 Connector) Pin	Terminal Server Serial Port Pin
Pin 1 (RTS)	Pin 1 (CTS)
Pin 2 (DTR)	Pin 2 (DSR)
Pin 3 (TXD)	Pin 3 (RXD)
Pin 4 (Signal Ground)	Pin 4 (Signal Ground)
Pin 5 (Signal Ground)	Pin 5 (Signal Ground)
Pin 6 (RXD)	Pin 6 (TXD)
Pin 7 (DSR /DCD)	Pin 7 (DTR)
Pin 8 (CTS)	Pin 8 (RTS)

2. **Open a terminal session on the connecting device, and type:**

```
% telnet IP-address-of-terminal-server  port-number
```

For example, for a server connected to port 10000 on a terminal server whose IP address is 192.20.30.10, you would type:

```
% telnet 192.20.30.10 10000
```

Accessing the System Console Through a TIP Connection

Use this procedure to access the server system console by connecting the serial management port (SER MGT) to the serial port of another Sun system (FIGURE 1-4).

FIGURE 1-4 TIP Connection Between a Server and Another Sun System

▼ To Access the System Console Through a TIP Connection

1. **Connect the RJ-45 serial cable and, if required, the DB-9 or DB-25 adapter provided.**

   The cable and adapter connect between another Sun system's serial port (typically TTYB) and the serial management port on the rear panel of the server. Pinouts, part numbers, and other details about the serial cable and adapter are provided in the service manual for your server.

2. **Ensure that the** `/etc/remote` **file on the Sun system contains an entry for** `hardwire`**.**

   Most releases of Solaris OS software shipped since 1992 contain an `/etc/remote` file with the appropriate `hardwire` entry. However, if the system is running an older version of Solaris OS software, or if the `/etc/remote` file has been modified, you might need to edit the file. See "Modifying the `/etc/remote` File" on page 12 for details.

3. **In a shell tool window on the remote system, type:**

   ```
   % tip hardwire
   ```

   The system responds by displaying:

   ```
   connected
   ```

   The shell tool is now a TIP window directed to your server through the remote system's serial port. This connection is established and maintained even when the server is completely powered off or just starting up.

Note – Use a shell tool or a CDE terminal (such as dtterm), not a command tool. Some TIP commands might not work properly in a command tool window.

Modifying the /etc/remote File

This procedure might be necessary if you are accessing the server using a TIP connection from a remote system running an older version of the Solaris OS software. You might also need to perform this procedure if the /etc/remote file on the remote system has been altered and no longer contains an appropriate hardwire entry.

Log in as superuser to the system console of a system that you intend to use to establish a TIP connection to your server.

▼ To Modify the /etc/remote File

1. **Determine the release level of Solaris OS software installed on the remote system. Type:**

```
# uname -r
```

The system responds with a release number.

2. **Do one of the following, depending on the number displayed.**

   ■ **If the number displayed by the** uname -r **command is 5.0 or higher:**

   The Solaris OS software shipped with an appropriate entry for hardwire in the /etc/remote file. If you suspect that this file was altered, and the hardwire entry modified or deleted, check the entry against the following example, and edit it as needed.

```
hardwire:\
        :dv=/dev/term/b:br#9600:el=^C^S^Q^U^D:ie=%$:oe=^D:
```

Note – If you intend to use the remote system's serial port A rather than serial port B, edit this entry by replacing /dev/term/b with /dev/term/a.

- **If the number displayed by the `uname -r` command is less than 5.0:**

  Check the `/etc/remote` file and add the following entry, if it does not already exist.

```
hardwire:\
        :dv=/dev/ttyb:br#9600:el=^C^S^Q^U^D:ie=%$:oe=^D:
```

Note – If you intend to use the remote system's serial port A rather than serial port B, edit this entry by replacing `/dev/ttyb` with `/dev/ttya`.

The `/etc/remote` file is now properly configured. Continue establishing a TIP connection to the server system console. See "TIP Connection Between a Server and Another Sun System" on page 11.

If you have redirected the system console to TTYB and want to change the system console settings back to use the serial management and network management ports, see "System Console OpenBoot Configuration Variable Settings" on page 23.

Accessing the System Console Through an Alphanumeric Terminal

Use this procedure when you access the server system console by connecting the serial port of an alphanumeric terminal to the serial management port (SER MGT) of the server.

▼ To Access the System Console Through an Alphanumeric Terminal

1. **Attach one end of the serial cable to the alphanumeric terminal's serial port.**

   Use a null modem serial cable or an RJ-45 serial cable and null modem adapter. Connect this cable to the terminal's serial port connector.

2. **Attach the opposite end of the serial cable to the serial management port on your server.**

3. **Connect the alphanumeric terminal's power cord to an AC outlet.**

4. **Set the alphanumeric terminal to receive:**
   - 9600 baud
   - 8 bits
   - No parity

- 1 stop bit
- No handshake protocol

Refer to the documentation accompanying your terminal for information about how to configure the terminal.

You can issue system commands and view system messages using the alphanumeric terminal. Continue with your installation or diagnostic procedure, as needed. When you are finished, type the alphanumeric terminal's escape sequence.

For more information about connecting to and using the system controller, refer to the ALOM CMT guide for your server.

Switching Between the System Controller Prompt and the System Console

The server features two management ports, labeled SER MGT and NET MGT, located on the server's rear panel. If the system console is directed to the virtual-console device (its default configuration), these ports provide access to both the system console and the ALOM CMT command-line interface (also called the system controller prompt, see FIGURE 1-5).

FIGURE 1-5 Switching between the System Console and the System Controller Prompt

If the system console is configured to use the virtual-console device, when you connect through one of these ports you can access either the system controller prompt) or the system console. You can switch between the system controller prompt and the system console at any time, but you cannot access both at the same time from a single terminal or shell tool.

The prompt displayed on the terminal or shell tool indicates which channel you are accessing:

- The # or % prompt indicates that you are at the system console and that the Solaris OS is running.
- The ok prompt indicates that you are at the system console and that the server is running under OpenBoot firmware control.
- The sc> prompt indicates that you are at the ALOM CMT command-line interface.

Note – If no text or prompt appears, it might be that no console messages were recently generated by the system. If this happens, pressing the terminal's Enter or Return key should produce a prompt. If the ALOM CMT session has timed out, pressing the terminal's Enter or Return key might not be effective. In that case, it might be necessary to issue the escape sequence is #. (Hash-Period) to return to ALOM CMT.

To reach the system console from the system controller prompt,

- Type the console command at the sc> prompt.

To reach ALOM CMT from the system console,

- Type the system controller escape sequence,

  By default, the escape sequence is #. (Hash-Period).

For more information about communicating with the system controller and system console, see the following:

Advanced Lights Out Management (ALOM) CMT v1.3 Guide

ALOM CMT and The sc> Prompt

The system controller runs independently of your server and regardless of system power state. When you connect your server to AC power, the system controller immediately starts up and begins monitoring the system.

Note – To view system controller boot messages, you must connect an alphanumeric terminal to the serial management port *before* connecting the AC power cords to the server.

You can log in to the system controller at any time, regardless of system power state, as long as AC power is connected to the system and you have a way of interacting with the system. The sc> prompt indicates that you are interacting with the system controller directly. The sc> prompt is the first prompt you see when you log in to the system through the serial management port or network management port.

Note – When you access the system controller for the first time and you issue an administrative command, you must create a password for the default username, admin, for subsequent access. After this initial configuration, you will be prompted to enter a user name and password every time you access the system controller.

For more information about navigating between the system console and ALOM CMT (the system controller prompt), see the following:

- "Getting to the ok Prompt" on page 21
- "Follow the appropriate instructions in TABLE 1-3." on page 22

Access Through Multiple Controller Sessions

Up to nine ALOM CMT sessions can be active concurrently, one session through the serial management port and up to eight sessions through the network management port. Users of each of these sessions can issue commands at the sc> prompt. For more information, see:

- "Accessing the System Controller" on page 6
- "Activating the Network Management Port" on page 7

Note – Only one user has active control of the system console at any time. Any additional ALOM CMT sessions afford passive views of system console activity, until the active user of the system console logs out. However, the console -f command enables users to seize access to the system console from one another. For more information, see the ALOM CMT guide for your server.

Reaching the sc> Prompt

There are several ways to get to the sc> prompt:

- You can log in directly to the system controller from a device connected to the serial management port. See "Accessing the System Controller" on page 6.

- You can log in directly to ALOM CMT on the system controller using a connection through the network management port. See "Activating the Network Management Port" on page 7.

- If you have logged in to ALOM CMT directly through the system controller and then directed the system console to the serial management and network management ports, you can return to the prior ALOM CMT session by typing the system controller escape sequence (#.).

OpenBoot ok Prompt

The server with the Solaris OS installed is capable of operating at different *run levels*. A synopsis of run levels follows. For a full description of run levels, refer to the Solaris system administration documentation.

Most of the time, you operate the server at run level 2 or run level 3, which are multiuser states with access to full system and network resources. Occasionally, you might operate the system at run level 1, which is a single-user administrative state. However, the lowest operational state is run level 0. At this state, it is safe to turn off power to the system.

When your server is at run level 0, the ok prompt appears. This prompt indicates that the OpenBoot firmware is in control of the system.

There are a number of scenarios under which OpenBoot firmware control can occur.

- By default, before the operating system is installed the system comes up under OpenBoot firmware control.

- When the auto-boot? OpenBoot configuration variable is set to false the system boots to the ok prompt.

- When the operating system is halted the system transitions to run level 0 in an orderly way.

- When the operating system crashes the system reverts to OpenBoot firmware control.

- During the boot process, when there is a serious hardware problem that prevents the operating system from running, the system reverts to OpenBoot firmware control.

- When a serious hardware problem develops while the system is running, the operating system transitions smoothly to run level 0.

- When you deliberately place the system under firmware control in order to execute firmware-based commands.

It is the last of these scenarios which most often concerns you as an administrator, since there will be times when you need to reach the ok prompt. The section "Methods To Reach the ok Prompt" on page 19 lists several ways. For detailed instructions, see "Getting to the ok Prompt" on page 21.

Methods To Reach the ok Prompt

There are several ways to reach the ok prompt, depending on the state of the system and the means by which you are accessing the system console.

Note – These methods of reaching the ok prompt work only if the system console has been redirected to the appropriate port. For details, see "System Console OpenBoot Configuration Variable Settings" on page 23.

The methods are:

- Graceful shutdown
- system controller break and console command pair
- L1-A (Stop-A) keys or Break key
- Manual system reset

A discussion of each method follows. For step-by-step instructions, see "Getting to the ok Prompt" on page 21.

Note – As a rule, before suspending the operating system, you should back up files, warn users of the impending shutdown, and halt the system in an orderly manner. However, it is not always possible to take such precautions, especially if the system is malfunctioning.

Graceful Shutdown

The preferred method of reaching the ok prompt is to shut down the operating system by issuing an appropriate command (for example, the shutdown, init, or uadmin command) as described in Solaris system administration documentation. You can also use the system Power button to initiate a graceful system shutdown.

Gracefully shutting down the system prevents data loss, enables you to warn users beforehand, and causes minimal disruption. You can usually perform a graceful shutdown, provided the Solaris OS is running and the hardware has not experienced serious failure.

ALOM CMT `break` or `console` Command

Typing `break` from the `sc>` prompt forces a running server to drop into OpenBoot firmware control. If the operating system is already halted, you can use the `console` command instead of `break` to reach the `ok` prompt.

> ⚠ **Caution –** After forcing the system into OpenBoot firmware control, be aware that issuing certain OpenBoot commands (like `probe-scsi`, `probe-scsi-all`, or `probe-ide`) might hang the system.

L1-A (Stop-A) Keys or Break Key

When it is impossible or impractical to shut down the system gracefully, you can get to the `ok` prompt by typing the L1-A (Stop-A) key sequence from a keyboard connected to the server (that is, if OpenBoot `input-device=keyboard`). If you have an alphanumeric terminal attached to the server, press the Break key.

> ⚠ **Caution –** After forcing the system into OpenBoot firmware control, be aware that issuing certain OpenBoot commands (like `probe-scsi`, `probe-scsi-all`, or `probe-ide`) might hang the system.

Manual System Reset

> ⚠ **Caution –** Forcing a manual system reset results in loss of system state data, and should be attempted only as a last resort. After a manual system reset, all state information is lost, which inhibits troubleshooting the cause of the problem until the problem reoccurs.

Use the system controller `reset` command, or `poweron` and `poweroff` commands, to reset the server. Reaching the `ok` prompt by performing a manual system reset or by power-cycling the system should be the method of last resort. Using these commands result in the loss of all system coherence and state information. A manual system reset could corrupt the server's file systems, although the `fsck` command usually restores them. Use this method only when nothing else works.

> ⚠ **Caution –** Accessing the `ok` prompt suspends the Solaris OS.

When you access the ok prompt from a functioning server, you are suspending the Solaris OS and placing the system under firmware control. Any processes that were running under the operating system are also suspended, and *the state of such processes might not be recoverable*.

After a manual system reset the system can be configured to boot automatically if the OpenBoot auto-boot? configuration variable is set to true. See "System Console OpenBoot Configuration Variable Settings" on page 23. If the server begins to boot automatically after a reset, you must abort the boot with the ALOM CMT break command or perform a graceful shutdown of the Solaris operating system once the boot has completed.

The commands you run from the ok prompt have the potential to affect the state of the system. This means that it is not always possible to resume execution of the operating system from the point at which it was suspended. Although the go command will resume execution in most circumstances, in general, each time you drop the system down to the ok prompt, you should expect to have to reboot the system to get back to the operating system.

For More Information about OpenBoot Firmware

For more information about the OpenBoot firmware, refer to the *OpenBoot 4.x Command Reference Manual.* An online version of the manual is included with the OpenBoot Collection AnswerBook that ships with Solaris software.

Getting to the ok Prompt

This procedure provides several ways of reaching the ok prompt. For details about when to use each method, see "OpenBoot ok Prompt" on page 18.

Caution – Forcing the server to the ok prompt suspends all application and operating system software. After you issue firmware commands and run firmware-based tests from the ok prompt, the system might not be able to resume where it left off.

If at all possible, back up system data before starting this procedure. Also exit or stop all applications and warn users of the impending loss of service. For information about the appropriate backup and shutdown procedures, see Solaris system administration documentation.

▼ To Get to the `ok` Prompt

1. Decide which method you need to use to reach the `ok` prompt.

See "OpenBoot `ok` Prompt" on page 18 for details.

2. Follow the appropriate instructions in TABLE 1-3**.**

TABLE 1-3   Ways of Accessing the `ok` Prompt

Access Method	What to Do
Graceful shutdown of the Solaris OS	• From a shell or command tool window, issue an appropriate command (for example, the `shutdown` or `init` command) as described in Solaris system administration documentation.
L1-A (Stop-A) keys or Break key	• From a Sun keyboard connected directly to the Sun Fire T1000 server, press the Stop and A keys simultaneously.* • From an alphanumeric terminal configured to access the system console, press the Break key.
system controller `break` and `console` commands	1. From the `sc>` prompt, type the `break` command. The `break` command should put the system in a state in which the operating environment software is not running and the server is under OpenBoot firmware control. 2. Then issue the `console` command.
Manual system reset	• From the `sc>` prompt, type the `reset` command.

* Requires the OpenBoot configuration variable `input-device=keyboard`. For more information, see "System Console OpenBoot Configuration Variable Settings" on page 23.

System Console OpenBoot Configuration Variable Settings

The server's system console is directed to the serial management and network management ports (SER MGT and NET MGT) by default.

Certain OpenBoot configuration variables control from where system console input is taken and to where its output is directed. The table below shows how to set these variables in order to use the serial management and network management ports.

TABLE 1-4 OpenBoot Configuration Variables That Affect the System Console

OpenBoot Configuration Variable Name	Serial and Network Management Ports
output-device	virtual-console
input-device	virtual-console

The serial management port does not function as a standard serial connection. (If you want to connect a conventional serial device (such as a printer) to the system, you must connect it to TTYA not the serial management port.)

It is important to note that the sc> prompt and POST messages are only available through the serial management port and network management port.

In addition to the OpenBoot configuration variables described in TABLE 1-4, there are other variables that affect and determine system behavior. These variables are discussed in more detail in Appendix A.

Managing RAS Features and System Firmware

This chapter describes how to manage reliability, availability, and serviceability (RAS) features and system firmware, including Sun Advanced Lights Out Manager (ALOM) system controller, and automatic system recovery (ASR). In addition, this chapter describes how to unconfigure and reconfigure a device manually, and introduces multipathing software.

This chapter contains the following sections:

Note – This chapter does not cover detailed troubleshooting and diagnostic procedures. For information about fault isolation and diagnostic procedures, refer to the diagnostics and troubleshooting guide for your server.

ALOM CMT and The System Controller

The system controller supports a total of nine concurrent ALOM CMT sessions per server, one connection through the serial management port and eight connections available through the network management port.

After you log in to your ALOM account, the system controller command prompt (sc>) appears, and you can enter system controller commands. If the command you want to use has multiple options, you can either enter the options individually or grouped together, as shown in the following example. The commands shown in the following example are identical.

```
sc> poweroff -f -y
sc> poweroff -fy
```

Logging In To The System Controller

All environmental monitoring and control is handled by the system controller. The system controller command prompt (sc>) provides you with a way of interacting with the system controller. For more information about the sc> prompt, see "ALOM CMT and The sc> Prompt" on page 16.

For instructions on connecting to the system controller, see:

- "Accessing the System Controller" on page 6
- "Activating the Network Management Port" on page 7

Note – This procedure assumes that the system console is directed to use the serial management and network management ports (the default configuration).

▼ To Log In To ALOM CMT

1. **If you are logged in to the system console, type** `#.` **(Hash-Period) to get to the** `sc>` **prompt.**

   Press the Hash key, followed by the Period key. Then press the Return key.

2. **At the ALOM CMT login prompt, enter the login name and press Return.**

   The default login name is `admin`.

   ```
   Sun(tm) Advanced Lights Out Manager CMT v1.3
   Please login: admin
   ```

3. **At the password prompt, enter the password and press Return to get to the** `sc>` **prompt.**

   ```
   Please Enter password:

   sc>
   ```

   Note – There is no default password when connecting to ALOM CMT for the first time using the serial management port. When connecting to the system controller using the network management port for the first time, the default ALOM CMT password is the last 8 digits of the Chassis Serial Number. The Chassis Serial Number can be found printed on the back of the server or in the printed system information sheet which shipped with your server. You must assign a password during initial system configuration. For more information, refer to the installation guide for your server and the ALOM CMT guide for your server.

   Caution – To provide optimum system security, change the default system login name and password during initial setup.

   Using the system controller, you can monitor the system, turn the Locator LED on and off, or perform maintenance tasks on the system controller card itself. For more information, refer to the ALOM CMT guide for your server.

▼ To View Environmental Information

1. **Log in to the system controller.**

2. **Use the** `showenvironment` **command to display a snapshot of the server's environmental status.**

The information this command can display includes temperature, power supply status, front panel LED status, and so on.

Note – Some environmental information might not be available when the server is in Standby mode.

Note – You do not need system controller user permissions to use this command.

Interpreting System LEDs

The behavior of LEDs on your server conform the the American National Standards Institute (ANSI) Status Indicator Standard (SIS). These standard LED behaviors are described in TABLE 2-1.

TABLE 2-1   LED Behavior and Meaning

LED Behavior	Meaning
Off	The condition represented by the color is not true.
Steady On	The condition represented by the color is true.
Standby Blink	The system is functioning at a minimal level and ready to resume full function.
Slow Blink	Transitory activity or new activity represented by the color is taking place.
Fast Blink	Attention is required.
Feedback Flash	Activity is taking place commensurate with the flash rate (such as disk drive activity).

The LEDs have assigned meanings, described in TABLE 2-2.

TABLE 2-2   LED Behaviors With Assigned Meanings

Color	Behavior	Definition	Description
White	Off	Steady State	
	Fast Blink	4 Hz repeating sequence, equal intervals On and Off.	This indicator helps you to locate a particular enclosure, board, or subsystem (for example, the Locator LED).
Blue	Off	Steady State	
	Steady On	Steady State	If blue is on, a service action can be performed on the applicable component with no adverse consequences (for example, the OK-to-Remove LED).
Yellow/Amber	Off	Steady State	
	Steady On	Steady State	This indicator signals the existance of a fault condition. Service is required (for example, the Service Required LED).
Green	Off	Steady State	
	Standby Blink	Repeating sequence consisting of a brief (0.1 sec.) on flash followed by a long off period (2.9 sec.)	The system is running at a minimum level and is ready to be quickly revived to full function (for example, the System Activity LED).
	Steady On	Steady State	Status normal; system or component functioning with no service actions required
	Slow Blink		A transitory (temporary) event is taking place for which direct proportional feedback is not needed or not feasible.

Controlling the Locator LED

You control the Locator LED from the `sc>` prompt or by the Locator button on the front of the chassis.

Locator Button/LED

FIGURE 2-1   Locator Button on the Front of the Server Chassis

- **To turn on the Locator LED from the system controller command prompt, type:**

```
sc> setlocator on
```

- **To turn off the Locator LED from the system controller command prompt, type:**

```
sc> setlocator off
```

- **To display the state of the Locator LED from the system controller command prompt, type:**

```
sc> showlocator
Locator LED is on.
```

Note – You do not need user permissions to use the `setlocator` and `showlocator` commands

Automatic System Recovery

The system provides for automatic system recovery (ASR) from failures in memory modules or PCI cards.

Automatic system recovery functionality enables the system to resume operation after experiencing certain nonfatal hardware faults or failures. When ASR is enabled, the system's firmware diagnostics automatically detect failed hardware components. An autoconfiguring capability designed into the system firmware enables the system to unconfigure failed components and to restore system operation. As long as the system is capable of operating without the failed component, the ASR features enable the system to reboot automatically, without operator intervention.

Note – ASR is not activated until you enable it. See "Enabling and Disabling Automatic System Recovery" on page 33.

For more information about ASR, refer to the service manual for your server.

AutoBoot Options

The system firmware stores a configuration variable called `auto-boot?`, which controls whether the firmware will automatically boot the operating system after each reset. The default setting for Sun platforms is `true`.

Normally, if a system fails power-on diagnostics, `auto-boot?` is ignored and the system does not boot unless an operator boots the system manually. An automatic boot is generally not acceptable for booting a system in a degraded state. Therefore, the server's OpenBoot firmware provides a second setting, `auto-boot-on-error?`. This setting controls whether the system will attempt a degraded boot when a subsystem failure is detected. Both the `auto-boot?` and `auto-boot-on-error?` switches must be set to `true` to enable an automatic degraded boot.

▼ To enable an Automatic Degraded Boot

● **Set the switches by typing:**

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

Note – The default setting for `auto-boot-on-error?` is `false`. The system will not attempt a degraded boot unless you change this setting to `true`. In addition, the system will not attempt a degraded boot in response to any fatal nonrecoverable error, even if degraded booting is enabled. For examples of fatal nonrecoverable errors, see "Error Handling Summary" on page 32.

Error Handling Summary

Error handling during the power-on sequence falls into one of the following three cases:

- If no errors are detected by POST or OpenBoot Diagnostics, the system attempts to boot if `auto-boot?` is `true`.

- If only nonfatal errors are detected by POST or OpenBoot Diagnostics, the system attempts to boot if `auto-boot?` is `true` and `auto-boot-on-error?` is `true`. Nonfatal errors include the following:

  - Ethernet interface failure.

  - Serial interface failure.

  - PCI-Express card failure.

  - Memory failure. When a DIMM fails, the firmware unconfigures the entire logical bank associated with the failed module. Another nonfailing logical bank must be present in the system for the system to attempt a degraded boot. Note that certain DIMM failures might not be diagnosable to a single DIMM. These failures are fatal, and result in both logical banks being unconfigured.

Note – If POST or OpenBoot Diagnostics detect a nonfatal error associated with the normal boot device, the OpenBoot firmware automatically unconfigures the failed device and tries the next-in-line boot device, as specified by the `boot-device` configuration variable.

- If a fatal error is detected by POST or OpenBoot Diagnostics, the system does not boot regardless of the settings of `auto-boot?` or `auto-boot-on-error?`. Fatal nonrecoverable errors include the following:

  - Any CPU failed
  - All logical memory banks failed
  - Flash RAM cyclical redundancy check (CRC) failure
  - Critical field-replaceable unit (FRU) PROM configuration data failure
  - Critical system configuration SEEPROM read failure
  - Critical application-specific integrated circuit (ASIC) failure

For more information about troubleshooting fatal errors, refer to the service manual for your server.

Reset Scenarios

Three ALOM CMT configuration variables, `diag_mode`, `diag_level`, and `diag_trigger`, control whether the system runs firmware diagnostics in response to system reset events.

The standard system reset protocol bypasses POST completely unless the virtual keyswitch or ALOM CMT variables are set as follows:

TABLE 2-3   Virtual Keyswitch Setting for Reset Scenario

Keyswitch	Value
virtual keyswitch	`diag`

TABLE 2-4   ALOM CMT Variable Settings for Reset Scenario

Variable	Value	Default
`diag-mode`	`normal` or `service`	`normal`
`diag-level`	`min` or `max`	`max`
`diag-trigger`	`power-on-reset` `error-reset`	`power-on-reset`

Therefore, ASR is enabled by default. For instructions, see "Enabling and Disabling Automatic System Recovery" on page 33.

Automatic System Recovery User Commands

The ALOM CMT commands are available for obtaining ASR status information and for manually unconfiguring or reconfiguring system devices. For more information, see:

- "Unconfiguring and Reconfiguring Devices" on page 35
- "To Reconfigure a Device Manually" on page 36
- "Obtaining Automatic System Recovery Information" on page 35

Enabling and Disabling Automatic System Recovery

The automatic system recovery (ASR) feature is not activated until you enable it. Enabling ASR requires changing configuration variables in ALOM CMT as well as OpenBoot firmware.

▼ To Enable Automatic System Recovery

1. At the `sc>` prompt, type:

```
sc> setsc diag-mode normal
sc> setsc diag-level max
sc> setsc diag-trigger power-on-reset
```

2. At the `ok` prompt, type:

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

Note – For more information about OpenBoot configuration variables, refer to the *Sun Fire T1000 Server Service Manual*.

3. To cause the parameter changes to take effect, type:

```
ok reset-all
```

The system permanently stores the parameter changes and boots automatically when the OpenBoot configuration variable `auto-boot?` is set to `true` (its default value).

Note – To store parameter changes, you can also power cycle the system using the front panel Power button.

▼ To Disable Automatic System Recovery

1. At the `ok` prompt, type:

```
ok setenv auto-boot-on-error? false
```

2. To cause the parameter changes to take effect, type:

```
ok reset-all
```

The system permanently stores the parameter change.

Note – To store parameter changes, you can also power cycle the system using the front panel Power button.

After you disable the automatic system recovery (ASR) feature, it is not activated again until you re-enable it.

Obtaining Automatic System Recovery Information

Use the following procedure to retrieve information about the status of system components affected by automatic system recovery (ASR).

● **At the** `sc>` **prompt, type:**

```
sc> showcomponent
```

In the `showcomponent` command output, any devices marked disabled have been manually unconfigured using the system firmware. The `showcomponent` command also lists devices that have failed firmware diagnostics and have been automatically unconfigured by the system firmware.

For more information, see:

Unconfiguring and Reconfiguring Devices

To support a degraded boot capability, the ALOM CMT firmware provides the `disablecomponent` command, which enables you to unconfigure system devices manually. This command flags the specified device as *disabled* by creating an entry in the ASR database.

▼ To Unconfigure a Device Manually

● **At the** `sc>` **prompt, type:**

```
sc> disablecomponent asr-key
```

The *asr-key* is one of the device identifiers from

Note – The device identifiers are not case-sensitive. You can type them as uppercase or lowercase characters.

TABLE 2-5   Device Identifiers and Devices

Device Identifiers	Devices
`MB/CMP`*cpu-number*`/P`*strand_number*	CPU strand (Number: 0-31)
`PCIE`*slot-number*	PCI-E slot (Number: 0)
`MB/PCIEa`	PCI-E leaf A (`/pci@780`)
`MB/PCIEb`	PCI-E leaf B (`/pci@7c0`)
`MB/CMP0/CH`*channel-number*`/R`*rank-number*`/D`*dimm-number*	DIMMS

▼ To Reconfigure a Device Manually

1. **At the** `sc>` **prompt, type:**

```
sc> enablecomponent asr-key
```

where *asr-key* is any device identifier from

Note – The device identifiers are not case-sensitive. You can type them as uppercase or lowercase characters.

You can use the ALOM CMT `enablecomponent` command to reconfigure any device that you previously unconfigured with the `disablecomponent` command.

Displaying System Fault Information

ALOM CMT software enables you to display current valid system faults. The `showfaults` command displays the fault ID, the faulted FRU device, and the fault message to standard output. The `showfaults` command also displays POST results.

▼ To Display System Fault Information

- **Type** `showfaults`

  For example:

  ```
  sc> showfaults
     ID FRU        Fault
      0 FT0.F2     SYS_FAN at FT0.F2 has FAILED.
  ```

  Adding the −v option displays additional information,

  ```
  sc> showfaults -v
     ID Time             FRU        Fault
      0  MAY 20 10:47:32 FT0.F2     SYS_FAN at FT0.F2 has FAILED.
  ```

  For more information about the `showfaults` command, refer to the *Advanced Lights Out Management (ALOM) CMT v1.3 Guide*.

Multipathing Software

Multipathing software enables you to define and control redundant physical paths to I/O devices, such as storage devices and network interfaces. If the active path to a device becomes unavailable, the software can automatically switch to an alternate path to maintain availability. This capability is known as *automatic failover*. To take advantage of multipathing capabilities, you must configure the server with redundant hardware, such as redundant network interfaces or two host bus adapters connected to the same dual-ported storage array.

For your server, three different types of multipathing software are available:

- Solaris IP Network Multipathing software provides multipathing and load-balancing capabilities for IP network interfaces.
- VERITAS Volume Manager (VVM) software includes a feature called Dynamic Multipathing (DMP), which provides disk multipathing as well as disk load balancing to optimize I/O throughput.
- Sun StorEdge™ Traffic Manager is an architecture fully integrated within the Solaris OS (beginning with the Solaris 8 release) that enables I/O devices to be accessed through multiple host controller interfaces from a single instance of the I/O device.

For More Information on Multipathing Software

For instructions on how to configure and administer Solaris IP Network Multipathing, consult the *IP Network Multipathing Administration Guide* provided with your specific Solaris release.

For information about VVM and its DMP feature, refer to the documentation provided with the VERITAS Volume Manager software.

For information about Sun StorEdge Traffic Manager, refer to your Solaris OS documentation.

Storing FRU Information

The `setfru` command enables you to store information on FRU PROMs. For example, you might store information identifying the server in which the FRUs have been installed.

▼ To Store Information in Available FRU PROMs

- **At the** `sc>` **prompt type:**

```
setfru -c data
```

Managing Disk Volumes

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

This chapter contains the following sections:

- "RAID Requirements" on page 41
- "Disk Volumes" on page 41
- "RAID Technology" on page 42
- "Hardware Raid Operations" on page 44

RAID Requirements

To configure and use RAID disk volumes on the server, you must install the appropriate patches. For the latest information on patches for the server, see your server's firmware product notes for this release. Patches are available for download from `http://www.sun.com/sunsolve`. Installation procedures for patches are included in text README files that accompany the patches.

Disk Volumes

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

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

The on-board disk controller of the server provides for the creation of one hardware RAID volume. The controller supports either a two-disk RAID 1 (integrated mirror, or IM) volume or a two-disk RAID 0 (integrated stripe, or IS) volume.

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

Volume migration (relocating all RAID volume disk members from one server chassis to another) is not supported. If this operation must be performed, please contact your service provider.

RAID Technology

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

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

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

Integrated Stripe Volumes (RAID 0)

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

Integrated stripe volumes provide for a logical unit (LUN) that is equal in capacity to the sum of all its member disks. For example, a two-disk IS volume configured on 72 GByte drives will have a 144 GByte capacity.

FIGURE 3-1 Graphical Representation of Disk Striping

⚠

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

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

Integrated Mirror Volumes (RAID 1)

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


FIGURE 3-2 Graphical Representation of Disk Mirroring

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

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

Hardware Raid Operations

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

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

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

If your system encounters a disk error, often you can find messages about failing or failed disks in the system console. This information is also logged in the `/var/adm/messages` files.

These error messages typically refer to a failed hard drive by its physical device name (such as `/devices/pci@7c0/pci@0/pci@8/scsi@2/sd@1,0`) or by its logical device name (such as `c0t0d0`). In addition, some applications might report a disk slot number (0 or 1).

You can use TABLE 3-1 to associate internal disk slot numbers with the logical and physical device names for each hard drive.

TABLE 3-1   Disk Slot Numbers, Logical Device Names, and Physical Device Names

Disk Slot Number	Logical Device Name*	Physical Device Name
Slot 0	c0t0d0	/devices/pci@7c0/pci@0/pci@8/scsi@2/sd@0,0
Slot 1	c0t1d0	/devices/pci@7c0/pci@0/pci@8/scsi@2/sd@1,0

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

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

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

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

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

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

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

   ok devalias disk
   disk                  /pci@7c0/pci@0/pci@8/scsi@2/disk@0,0
   ```

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

   ```
   ok boot net -s
   ```

3. Verify that the member disks are available and that there is not a volume already created. using the `raidctl` **command:**

The server's on-board SAS controller can configure one RAID volume. Prior to volume creation, ensure that the member disks are available and that there is not a volume already created.

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

See "Physical Disk Slot Numbers, Physical Device Names, and Logical Device Names for Non-RAID Disks" on page 44.

The preceding example indicates that no RAID volume exists. In another example, a single IM volume has been enabled. It is fully synchronized and is online:

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

4. **Create the RAID 1 volume:**

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

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

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

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

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

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

5. **Check the status of a RAID mirror.**

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

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

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

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

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

```

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

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

The preceding example indicates that the RAID mirror is still resynchronizing with the backup drive.

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

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

Under RAID 1 (disk mirroring), all data is duplicated on both drives. If a disk fails, see the service manual for your server for instructions.

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

6. **Relabel the disk using the** `format` **utility.**

```
# format
Searching for disks...done

AVAILABLE DISK SELECTIONS:
       0. c0t0d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424>
          /pci@7c0/pci@0/pci@8/scsi@2/sd@0,0
Specify disk (enter its number): 0
selecting c0t0d0
[disk formatted]

FORMAT MENU:
...
format> type

AVAILABLE DRIVE TYPES:
       0. Auto configure
       ...
       19. SUN72G
       20. other
Specify disk type (enter its number)[19]: 0
c0t0d0: configured with capacity of 68.00GB
<LSILOGIC-LogicalVolume-3000 cyl 65533 alt 2 hd 16 sec 136>
selecting c0t0d0
[disk formatted]
format> label
Ready to label disk, continue? yes

format> disk

AVAILABLE DISK SELECTIONS:
       0. c0t0d0 <LSILOGIC-LogicalVolume-3000 cyl 65533 alt 2 hd
16 sec 136>
          /pci@7c0/pci@0/pci@8/scsi@2/sd@0,0
Specify disk (enter its number)[0]: 0
selecting c0t0d0
[disk formatted]
format> quit
#
```

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

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

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

▼ To Create a Hardware Striped Volume

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

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

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

   ok devalias disk
   disk                 /pci@7c0/pci@0/pci@8/scsi@2/disk@0,0
   ```

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

   ```
   ok boot net –s
   ```

3. **Verify that the member disks are available and that there is not a volume already created.**

   The server's on-board SAS controller can configure one RAID volume. Prior to volume creation, ensure that the member disks are available and that there is not a volume already created.

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

   See .

   The preceding example indicates that no RAID volume exists.

4. **Create the RAID 0 volume.**

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

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

```
# raidctl -c -r 0 c0t0d0 c0t1d0
Creating RAID volume c0t1d0 will destroy all data on member disks,
proceed
(yes/no)? yes
Volume 'c0t0d0' created
#
```

When you create a RAID striped volume, the other member drives (in this case, c0t1d0) disappear from the Solaris device tree.

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

```
# raidctl -f -c -r 0 c0t0d0 c0t1d0
Volume 'c0t0d0' created
#
```

5. **Check the status of a RAID striped volume.**

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

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

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

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

6. **Relabel the disks using the** format **utility.**

```
# format
Searching for disks...done

AVAILABLE DISK SELECTIONS:
       0. c0t0d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424>
          /pci@7c0/pci@0/pci@8/scsi@2/sd@0,0
Specify disk (enter its number): 0
selecting c0t0d0
[disk formatted]

FORMAT MENU:
...
format> type

AVAILABLE DRIVE TYPES:
       0. Auto configure
       ...
       19. SUN72G
       20. other
Specify disk type (enter its number)[19]: 0
c0t0d0: configured with capacity of 68.00GB
<LSILOGIC-LogicalVolume-3000 cyl 65533 alt 2 hd 16 sec 136>
selecting c0t0d0
[disk formatted]
format> label
Ready to label disk, continue? yes

format> disk

AVAILABLE DISK SELECTIONS:
       0. c0t0d0 <LSILOGIC-LogicalVolume-3000 cyl 65533 alt 2 hd
16 sec 136>
          /pci@7c0/pci@0/pci@8/scsi@2/sd@0,0
Specify disk (enter its number)[0]: 0
selecting c0t0d0
[disk formatted]
format> quit
#
```

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

▼ To Delete a Hardware RAID Volume

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

   See "Disk Slot Numbers, Logical Device Names, and Physical Device Names" on page 45.

2. **Determine the name of the RAID volume.**

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

In this example, the RAID volume is c0t1d0.

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

---

3. **Delete the volume.**

```
# raidctl -d volume
```

For example:

```
# raidctl -d c0t0d0
```

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

```
# raidctl -d c0t0d0
Are you sure you want to delete RAID-1 Volume c0t0d0(yes/no)? yes
/pci@7c0/pci@0/pci@8/scsi@2 (mpt0):
        Volume 0 deleted.
/pci@7c0/pci@0/pci@8/scsi@2 (mpt0):
        Physical disk 0 deleted.
/pci@7c0/pci@0/pci@8/scsi@2 (mpt0):
        Physical disk 1 deleted.
Volume 'c0t0d0' deleted.
#
```

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

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

4. **Confirm that you have deleted the RAID array.**

```
# raidctl
```

For example:

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

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

5. **To re-label all of the member disks of the volume using the** `format` **command, select the disk name that represents the RAID volume that you have configured.**

In this example, `c0t0d0` is the logical name of the volume.

```
# format
Searching for disks...done

AVAILABLE DISK SELECTIONS:
       0. c0t0d0 <LSILOGIC-LogicalVolume-3000 cyl 65533 alt 2 hd
16 sec 136>
          /pci@7c0/pci@0/pci@8/scsi@2/sd@0,0
       1. c0t1d0 <LSILOGIC-LogicalVolume-3000 cyl 65533 alt 2 hd
16 sec 136>
          /pci@7c0/pci@0/pci@8/scsi@2/sd@1,0
Specify disk (enter its number): 0
selecting c0t0d0
[disk formatted]

FORMAT MENU:
        disk       - select a disk
        type       - select (define) a disk type
        partition  - select (define) a partition table
        current    - describe the current disk
        format     - format and analyze the disk
        repair     - repair a defective sector
        label      - write label to the disk
        analyze    - surface analysis
        defect     - defect list management
        backup     - search for backup labels
        verify     - read and display labels
        save       - save new disk/partition definitions
        inquiry    - show vendor, product and revision
        volname    - set 8-character volume name
        !<cmd>     - execute <cmd>, then return
        quit
```

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

   For example:

   ```
   format> type

   AVAILABLE DRIVE TYPES:
           0. Auto configure
           1. Quantum ProDrive 80S
           2. Quantum ProDrive 105S
           3. CDC Wren IV 94171-344
           4. SUN0104
           5. SUN0207
           6. SUN0327
           7. SUN0340
           8. SUN0424
           9. SUN0535
           10. SUN0669
           11. SUN1.0G
           12. SUN1.05
           13. SUN1.3G
           14. SUN2.1G
           15. SUN2.9G
           16. Zip 100
           17. Zip 250
           18. Peerless 10GB
           19. LSILOGIC-LogicalVolume-3000
           20. other
    Specify disk type (enter its number)[19]: 0
    c0t0d0: configured with capacity of 68.35GB
    <SUN72G  cyl 14087 alt 2 hd 24 sec 424>
    selecting c0t0d0
    [disk formatted]
   ```

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

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

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

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

9. **Verify that the new labels have been written by printing the disk list using the** `disk` **command.**

```
format> disk

AVAILABLE DISK SELECTIONS:
       0. c0t0d0 <SUN72G  cyl 14087 alt 2 hd 24 sec 424>
          /pci@7c0/pci@0/pci@8/scsi@2/sd@0,0
       1. c0t1d0 <LSILOGIC-LogicalVolume-3000 cyl 65533 alt 2 hd
16 sec 136>
          /pci@7c0/pci@0/pci@8/scsi@2/sd@1,0
Specify disk (enter its number)[0]: 1
selecting c0t1d0
[disk formatted]
```

Note that `c0t1d0` now has a type indicating it is an `LSILOGIC-LogicalVolume`.

10. **Repeat the labeling process for the second disk.**

11. **Exit the** `format` **utility.**

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

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

OpenBoot Configuration Variables

TABLE A-1 describes the OpenBoot firmware configuration variables stored in non-volatile memory on the system. The OpenBoot configuration variables are printed here in the order in which they appear when you issue the showenv command.

TABLE A-1   OpenBoot Configuration Variables Stored on the System Configuration Card

Variable	Possible Values	Default Value	Description
local-mac-address?	true, false	true	If true, network drivers use their own MAC address, not the server MAC address.
fcode-debug?	true, false	false	If true, include name fields for plug-in device FCodes.
scsi-initiator-id	0-15	7	SCSI ID of the Serial Attached SCSI controller.
oem-logo?	true, false	false	If true, use custom OEM logo; otherwise, use Sun logo.
oem-banner?	true, false	false	If true, use custom OEM banner.
ansi-terminal?	true, false	true	If true, enables ANSI terminal emulation.
screen-#columns	0-n	80	Sets number of columns on screen.
screen-#rows	0-n	34	Sets number of rows on screen.
ttya-rts-dtr-off	true, false	false	If true, operating system does not assert rts (request-to-send) and dtr (data-transfer-ready) on serial management port.
ttya-ignore-cd	true, false	true	If true, operating system ignores carrier-detect on serial management port.

TABLE A-1 OpenBoot Configuration Variables Stored on the System Configuration Card *(Continued)*

Variable	Possible Values	Default Value	Description
ttya-mode	9600,8,n,1,-	9600,8,n,1,-	Serial management port (baud rate, bits, parity, stop, handshake). The serial management port only works at the default values.
output-device	virtual-console, screen	virtual-console	Power-on output device.
input-device	virtual-console, keyboard	virtual-console	Power-on input device.
auto-boot-on-error?	true, false	false	If true, boot automatically after system error.
load-base	0-n	16384	Address.
auto-boot?	true, false	true	If true, boot automatically after power on or reset.
boot-command	*variable-name*	boot	Action following a boot command.
boot-file	*variable-name*	none	File from which to boot if diag-switch? is false.
boot-device	*variable-name*	disk net	Devices from which to boot if diag-switch? is false.
use-nvramrc?	true, false	false	If true, execute commands in NVRAMRC during server startup.
nvramrc	*variable-name*	none	Command script to execute if use-nvramrc? is true.
security-mode	none, command, full	none	Firmware security level.
security-password	*variable-name*	none	Firmware security password if security-mode is not none (never displayed). *Do not set this directly.*
security-#badlogins	*variable-name*	none	Number of incorrect security password attempts.

Variable	Possible Values	Default Value	Description
diag-switch?	true, false	false	If true: • OpenBoot verbosity is set to maximum If false: • OpenBoot verbosity is set to minimum
error-reset-recovery	boot, sync, none	boot	Command to execute following a system reset generated by an error.
network-boot-arguments	[*protocol*,] [*key=value,*]	none	Arguments to be used by the PROM for network booting. Defaults to an empty string. Use network-boot-arguments to specify the boot protocol (RARP/DHCP) and a range of system knowledge to be used in the process. For further information, see the eeprom (1M) man page or your Solaris reference manual.

Index