Sun SPARC Enterprise T5140 and T5240 Servers

Administration Guide

Please
Recycle

Adobe PostScript

Contents

Preface

This administration guide is for experienced system administrators of Oracle's Sun SPARC Enterprise T5140 and T5240 servers. The guide includes general descriptive information about the servers, and detailed instructions for configuring and administering the servers. To use the information in this document, you must have working knowledge of computer network concepts and terms, and advanced familiarity with the Oracle Solaris Operating System (Solaris OS).

This preface contains the following topics:

- "UNIX Commands" on page vii
- "Shell Prompts" on page viii
- "Related Documentation" on page viii
- "Documentation, Support, and Training" on page ix
- "Documentation Feedback" on page ix

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. Refer to the following for this information:

- Software documentation that you received with your system
- Solaris™ Operating System 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	#

Related Documentation

The documents listed as online are available at:

(http://docs.sun.com/app/docs/prod/sparc.t5140)

(http://docs.sun.com/app/docs/prod/sparc.t5240)

Application	Title	Part Number	Format	Location
Product Notes	*Sun SPARC Enterprise T5140 and T5240 Servers Product Notes*	820-3312	PDF	Online
Getting Started	*Sun SPARC Enterprise T5140 Server Getting Started Guide*	820-3875	Printed	Ships with system
Getting Started	*Sun SPARC Enterprise T5140 Server Getting Started Guide (DC)*	820-5840	Printed	Ships with system
Getting Started	*Sun SPARC Enterprise T5240 Server Getting Started Guide*	820-3521	Printed	Ships with system
Getting Started	*Sun SPARC Enterprise T5240 Server Getting Started Guide (DC)*	820-5841	Printed	Ships with system
Overview	*Sun SPARC Enterprise T5140 and T5240 Servers Overview Guide*	820-3313	PDF HTML	Online

Application	Title	Part Number	Format	Location
Planning	*Sun SPARC Enterprise T5140 and T5240 Servers Site Planning Guide*	820-3314	PDF HTML	Online
Installation	*Sun SPARC Enterprise T5140 and T5240 Servers Installation Guide*	820-3315	PDF HTML	Online
Administration	*Sun SPARC Enterprise T5140 and T5240 Servers Administration Guide*	820-3316	PDF HTML	Online
Service	*Sun SPARC Enterprise T5140 and T5240 Servers Service Manual*	820-3318	PDF HTML	Online
Safety	*Sun SPARC Enterprise T5140 and T5240 Servers Safety and Compliance Guide*	820-3319	PDF	Online
Remote Management	*Oracle Integrated Lights Out Manager (ILOM) 3.0 Supplement for Sun SPARC Enterprise T5140 and T5240 Servers*	820-6684	PDF HTML	Online
Remote Management	*Oracle Intergrated Lights-Out Manager (ILOM) 3.0 Getting Started Guide*	820-5523	PDF HTML	Online

Documentation, Support, and Training

The following additional resources are available:

- Documentation (http://www.sun.com/documentation)
- Support (http://www.sun.com/support)
- Training (http://www.sun.com/training)

Documentation Feedback

Submit comments about this document by clicking the Feedback[+] link at (http://docs.sun.com).

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

Sun SPARC Enterprise T5140 and T5240 Servers Administration Guide, part number 820-3316-14.

Communicating With the System

This section includes information on low-level communication with the system, using the Oracle Integrated Lights Out Manager (ILOM) tool andthe system console.

ILOM Overview

The ILOM service processor runs independently of the server and regardless of the system power state as long as AC power is connected to the system. When you connect your server to AC power, the ILOM service processor immediately starts up, and begins monitoring the system. All environmental monitoring and control is handled by ILOM

The `->` prompt indicates that you are interacting with the ILOM service processor directly. This prompt is the first prompt you see when you log in to the system through the serial management port or network management port, regardless of the host's power state.

You can also access the ILOM service processor prompt (`->`) from the OpenBoot `ok` prompt, or from the Solaris # or % prompt, provided the system console is configured to be accessible through the serial management and network management ports.

The ILOM service processor supports a total of five concurrent sessions per server, four SSH connections available through the network management port and one connection available through the serial management port.

Related Information

-

- Integrated Lights Out Manager (ILOM) Documentation

    (http://docs.sun.com/app/docs/prod/int.lights.mgr)

- *Integrated Lights Out Manager (ILOM) Supplement for SPARC Enterprise T5140 and T5240 Servers*

▼ Log In to ILOM

This procedure assumes the default configuration of the service processor as described in your server's installation guide.

● **Open an SSH session and connect to the service processor by specifying its IP address.**

The ILOM default username is *root* and the default password is *changeme*.

```
% ssh root@xxx.xxx.xxx.xxx
...
Are you sure you want to continue connecting (yes/no) ? yes

...
Password: password (nothing displayed)
Waiting for daemons to initialize...

Daemons ready

Oracle Integrated Lights Out Manager

Version 3.x.x.x

Copyright 2010 Oracle and/or its affiliates. All rights reserved.
Use is subject to license terms.

->
```

You are now logged in to ILOM. Perform tasks as needed.

Note – To provide optimum system security, change the default system password.

Related Information

- "ILOM Overview" on page 1
- "Log In to the System Console" on page 3

▼ Log In to the System Console

1. "Log In to ILOM" on page 2.

2. **To access the system console from ILOM, type:**

```
-> start /SP/console
Are you sure you want to start /SP/console (y/n) ? y
Serial console started. To stop, type #.
.
.
.
```

You are logged in to the system console. Perform tasks as needed.

Note – If the Solaris OS is not running, the system displays the ok prompt.

Related Information

- "Display the ILOM ->Prompt" on page 4
- "Use a Local Graphics Monitor" on page 5

▼ Display the ok Prompt

This procedure assumes the default system console configuration.

● **Choose the appropriate shutdown method from the following table to reach the ok prompt.**

> ⚠️ **Caution –** When possible, reach the ok prompt by performing a graceful shutdown of the OS. Any other method used might result in the loss of system state data.

System State	What To Do
OS Running and Responsive	**Shut down the system using one of these methods:** • From a shell or command tool window, issue an appropriate command (for example, the shutdown or init 0 command) as described in Solaris system administration documentation. • From the ILOM -> prompt, type: `-> Stop /SYS` • Use the system Power button.

OS Unresponsive	**Shut down the system from ILOM**
	(provided the operating system software is not running and the server is already under OpenBoot firmware control).
	From the ILOM -> prompt, type:
	`-> set /HOST send_break_action=break`
	Press Enter.
	Then type:
	`-> start /SP/console`

OS Unresponsive and Need to Prevent auto-boot	**Shut down the system from ILOM and disable autoboot.**
	From the ILOM -> prompt, type:
	`-> set /HOST/bootmode script="setenv auto-boot? false"`
	Press Enter.
	Then type:
	`-> reset /SYS`
	`-> start /SP/console`

Related Information

- "Handling Faults" on page 43
- "OpenBoot Configuration Variables on the SCC" on page 53

▼ Display the ILOM ->Prompt

● **Use one of the following ways to get to the ILOM -> prompt:**

  ■ From the system console, type the ILOM escape sequence (#.).

- Log in to ILOM directly from a device connected to the serial management port or network management port.
- Log in to ILOM through an SSH connection. See "Log In to ILOM" on page 2.

Related Information

- "ILOM Overview" on page 1

▼ Use a Local Graphics Monitor

Though it is *not* recommended, the system console can be redirected to the graphics frame buffer. You *cannot* use a local graphics monitor to perform initial system installation, nor can you use a local graphics monitor to view power-on self-test (POST) messages.

To install a local graphics monitor, you must have the following items:

- Supported PCI-based graphics accelerator card and software driver
- Monitor with appropriate resolution to support the frame buffer
- Supported USB keyboard
- Supported USB mouse

1. **Install the graphics card into an appropriate PCI slot.**

   Installation must be performed by a qualified service provider. For further information, refer to the service manual for your server or contact your qualified service provider.

2. **Attach the monitor's video cable to the graphics card's video port.**

   Tighten the thumbscrews to secure the connection.

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

4. **Connect the USB keyboard cable to one USB port.**

5. **Connect the USB mouse cable to another USB port on the Sun SPARC Enterprise T5120 or T5220 server.**

6. "Display the ok Prompt" on page 3

7. **Set OpenBoot configuration variables appropriately.**

   From the existing system console, type:

```
ok setenv input-device keyboard
ok setenv output-device screen
```

Note – There are many other system configuration variables. Although these variables do not affect which hardware device is used to access the system console, some of the variables affect which diagnostic tests the system runs and which messages the system displays at its console. For details, refer to the service manual for your server.

8. To cause the changes to take effect, type:

```
ok reset-all
```

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

Note – To cause the parameter changes to take effect, you can also power cycle the system using the front panel Power button.

You can now type system commands and view system messages using your local graphics monitor. Continue with your installation or diagnostic procedure, as needed.

Performing Common Tasks

This section includes procedures for some common tasks performed on the servers.

- "Power On the System" on page 7
- "Power Off the System" on page 8
- "Reset the System" on page 9
- "Update the Firmware" on page 9

▼ Power On the System

1. "Log In to ILOM" on page 2

2. **At the ILOM -> prompt, type:**

```
-> start /SYS
Are you sure you want to start /SYS (y/n) ? y
Starting /SYS

->
```

Note – To force a power-on sequence, use the start -script /SYS command.

Related Information

- "Power Off the System" on page 8
- "Reset the System" on page 9

▼ Power Off the System

1. Shut down the Solaris OS.

At the Solaris prompt, type:

```
# shutdown -g0 -i0 -y
# svc.startd: The system is coming down. Please wait.
svc.startd: 91 system services are now being stopped.
Jun 12 19:46:57 wgs41-58 syslogd: going down on signal 15
svc.stard: The system is down.
syncing file systems...done
Program terminated
r)eboot o)k prompt, h)alt?
```

2. Switch from the system console prompt to the service processor console prompt. Type:

```
ok #.
->
```

3. From the ILOM -> prompt, type:

```
-> stop /SYS
Are you sure you want to stop /SYS (y/n)? y
Stopping /SYS

->
```

Note – To perform an immediate shutdown, use the `stop -force -script /SYS` command. Ensure that all data is saved before entering this command.

Related Information

-
-

▼ Reset the System

It is not necessary to power the system off and on to perform a reset.

● **To reset the system, from the Solaris prompt, type:**

```
# shutdown -g0 -i6 -y
```

Related Information

- ■ "Power Off the System" on page 8
- ■ "Power On the System" on page 7

▼ Update the Firmware

1. **Ensure that the ILOM service processor network management port is configured.**

   See the server's installation guide for instructions.

2. **Open an SSH session to connect to the service processor.**

```
% ssh root@xxx.xxx.xxx.xxx
...
Are you sure you want to continue connecting (yes/no) ? yes

...
Password: password (nothing displayed)
Waiting for daemons to initialize...

Daemons ready

Oracle Integrated Lights Out Manager

Version 3.x.x.x

Copyright 2010 Oracle and/or its affiliates. All rights reserved.
Use is subject to license terms.

->
```

3. **Power off the host. Type:**

```
-> stop /SYS
```

4. **Set the** keyswitch_state **parameter to** normal**. Type:**

```
-> set /SYS keyswitch_state=normal
```

5. **Type the** load **command with the path to the new flash image.**

   The load command updates the service processor flash image and the host firmware. The load command requires the following information:

   ■ IP address of a TFTP server on the network that can access the flash image

   ■ Full path name to the flash image that the IP address can access

   The command usage is as follows:

   load [-script] -source tftp://*xxx.xxx.xx.xx/pathname*

   where:

   ■ -script - Does not prompt for confirmation and acts as if yes was specified

   ■ -source - Specifies the IP address and full path name (URL) to the flash image

```
-> load -source tftp://129.168.10.101/pathname
NOTE: A firmware upgrade will cause the server and ILOM to be reset.
It is recommended that a clean shutdown of the server be done prior
to the upgrade procedure.
An upgrade takes about 6 minutes to complete. ILOM will enter a
special mode to load new firmware.
No other tasks can be performed in ILOM until the firmware upgrade
is complete and ILOM is reset.
Are you sure you want to load the specified file (y/n)?y
Do you want to preserve the configuration (y/n)? y
..............................................................
Firmware update is complete.
ILOM will now be restarted with the new firmware.
Update Complete. Reset device to use new image.
->
```

   After the flash image has been updated, the system automatically resets, runs diagnostics, and returns to the login prompt on the serial console.

```
U-Boot 1.1.1 (May 23 2008 - 21:30:12)
***
POST cpu PASSED
POST ethernet PASSED
Hit any key to stop autoboot: 0
```

```
## Booting image at fe080000    ***

IP Protocols: ICMP, UDP, TCP, IGMP

Checking all file systems...
fsck 1.37 (21-Mar-2005)
Setting kernel variable ...
... done.
Mounting local filesystems...
Cleaning /tmp /var/run /var/lock.

Identifying DOC Device Type(G3/G4/H3) ...
OK

Configuring network interfaces....Internet Systems Consortium DHCP
Client V3.0.1
Copyright 2007 Internet Systems Consortium
All rights reserved.
For info, please visit http://www.isc.org/products/DHCP

eth0: config: auto-negotiation on, 100FDX, 100HDX, 10FDX, 10HDX.
Listening on LPF/eth0/00:14:4f:3f:8c:af
Sending on   LPF/eth0/00:14:4f:3f:8c:af
Sending on   Socket/fallback
DHCPDISCOVER on eth0 to 255.255.255.255 port 67 interval 6
eth0: link up, 100Mbps Full Duplex, auto-negotiation complete.
DHCPDISCOVER on eth0 to 255.255.255.255 port 67 interval 15
Hostname: hostname
Starting portmap daemon: portmap.
Initializing random number generator...done.
INIT: Entering runlevel: 3
Starting system log daemon: syslogd and klogd.
Starting periodic command scheduler: cron.
Starting IPMI Stack........ Done.
Starting OpenBSD Secure Shell server: sshd.
Starting Servicetags listener: stlistener.
Starting FRU update program: frutool.

hostname login:
```

Managing Disks

This section describes how to configure and manage RAID disk volumes using the Sun SPARC Enterprise T5140 and T5240 server's on-board serial attached SCSI (SAS) disk controller, and how to hot-plug a disk.

Hardware RAID Support

RAID technology enables the construction of a logical volume, made up of several physical disks, in order to provide data redundancy, increased performance, or both. The Sun SPARC Enterprise T5140 and T5240 server's on-board disk controller supports both RAID 0 (striping) and RAID 1 (mirroring) volumes using the Solaris OS `raidctl` utility.

To configure and use RAID disk volumes on the Sun SPARC Enterprise T5140 and T5240 servers, you must install the appropriate patches. For the latest information on patches, see the product notes for your system.

Volume migration (relocating all RAID volume disk members from one Sun SPARC T5120 or T5220 chassis) is not supported. If you must perform this operation, contact your service provider.

The Sun SPARC Enterprise T5140 and T5240 servers can also be configured with a Sun StorageTek SAS RAID HBA. To manage RAID volumes on servers configured with these controllers, see the *Sun StorageTek RAID Manager's Software User's Guide*.

Related Information

- "Creating Hardware RAID Volumes" on page 14
- "Delete a Hardware RAID Volume" on page 25

Creating Hardware RAID Volumes

Caution – Creating RAID volumes using the on-board disk controller destroys all data on member disks.

- "Create a Hardware Mirrored Volume" on page 15
- "Create a Hardware Mirrored Volume of the Default Boot Device" on page 19
- "Create a Hardware Striped Volume" on page 20
- "Configure a Hardware RAID Volume for the Solaris OS" on page 22
- "Create a Hardware Mirrored Volume" on page 15
- "Create a Hardware Mirrored Volume of the Default Boot Device" on page 19
- "Create a Hardware Striped Volume" on page 20
- "Configure a Hardware RAID Volume for the Solaris OS" on page 22

▼ Create a Hardware Mirrored Volume

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

```
# raidctl
Controller: 1
Disk: 0.0.0
Disk: 0.1.0
Disk: 0.2.0
Disk: 0.3.0
Disk: 0.4.0
Disk: 0.5.0
Disk: 0.6.0
Disk: 0.7.0
```

See "Disk Slot Numbers" on page 35.

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

```
# raidctl
Controller: 1
Volume:c1t0d0
Disk: 0.0.0
Disk: 0.1.0
Disk: 0.2.0
Disk: 0.3.0
Disk: 0.4.0
Disk: 0.5.0
Disk: 0.6.0
Disk: 0.7.0
```

In this example, a single volume (c1t0d0) has been enabled.

The Sun SPARC Enterprise T5120 or T5220 server's on-board SAS controller can configure as many as two RAID volumes. Prior to volume creation, ensure that the member disks are available and that there are not two volumes already created.

The RAID status might be:

- OPTIMAL – Indicating that the RAID volume is online and fully synchronized.
- SYNC – Indicating that the data between the primary and secondary member disks in an IM are still synchronizing.
- DEGRADED – Indicating that a member disk is failed or otherwise offline.

- 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 GOOD, indicating that it is online and functioning properly, or it might be FAILED, 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 -l c1t0d0
Volume                   Size    Stripe  Status    Cache  RAID
       Sub                       Size                      Level
                 Disk
-----------------------------------------------------------------
c1t0d0                   136.6G  N/A     DEGRADED  OFF    RAID1
                 0.1.0   136.6G          GOOD
                 N/A     136.6G          FAILED
```

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

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

---

2. Type the following command:

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

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

```
# raidctl -c c1t0d0 c1t1d0
Creating RAID volume c1t0d0 will destroy all data on member disks,
proceed (yes/no)? yes
...
Volume c1t0d0 is created successfully!
#
```

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 c1t0d0 c1t1d0
Volume c1t0d0 is created successfully!
#
```

When you create a RAID mirror, the secondary drive (in this case, c1t1d0) disappears from the Solaris device tree.

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

```
# raidctl -l c1t0d0
Volume                    Size    Stripe  Status   Cache  RAID
        Sub                       Size                    Level
                  Disk
----------------------------------------------------------------
c1t0d0                    136.6G  N/A     SYNC     OFF    RAID1
                  0.0.0   136.6G          GOOD
                  0.1.0   136.6G          GOOD
```

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 -l c1t0d0
Volume                    Size    Stripe  Status   Cache  RAID
        Sub                       Size                    Level
                  Disk
----------------------------------------------------------------
c1t0d0                    136.6G  N/A     OPTIMAL  OFF    RAID1
                  0.0.0   136.6G          GOOD
                  0.1.0   136.6G          GOOD
```

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

Under RAID 1 (disk mirroring), all data is duplicated on both drives. If a disk fails, replace it with a working drive and restore the mirror. For instructions, see "Delete a Hardware RAID Volume" on page 25.

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

Related Information

-
-

▼ Create a Hardware Mirrored Volume of the Default Boot Device

Due to the volume initialization that occurs on the disk controller when a new volume is created, the volume must be configured and labeled using the format(1M) utility prior to use with the Solaris Operating System (see "Configure a Hardware RAID Volume for the Solaris OS" on page 22). 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@0/pci@0/pci@2/scsi@0/disk@0,0
```

2. Type the boot net –s **command.**

```
ok boot net –s
```

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

See "Create a Hardware Mirrored Volume" on page 15. For example:

```
# raidctl -c –r 1 c1t0d0 c1t1d0
Creating RAID volume c1t0d0 will destroy all data on member disks,
proceed (yes/no)? yes
...
Volume  c1t0d0  is created successfully!
#
```

4. Install the volume with the Solaris OS using any supported method.

The hardware RAID volume `c1t0d0` 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.

Related Information

- "Disk Slot Numbers" on page 35
- "Create a Hardware Mirrored Volume" on page 15
- "Configure a Hardware RAID Volume for the Solaris OS" on page 22

▼ Create a Hardware Striped Volume

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

See "Disk Slot Numbers" on page 35.

To verify the current RAID configuration, type:

```
# raidctl
Controller: 1
Disk: 0.0.0
Disk: 0.1.0
Disk: 0.2.0
Disk: 0.3.0
Disk: 0.4.0
Disk: 0.5.0
Disk: 0.6.0
Disk: 0.7.0
```

The preceding example indicates that no RAID volume exists.

2. Type the following command:

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

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

```
# raidctl -c -r 0 c1t1d0 c1t2d0 c1t3d0
Creating RAID volume will destroy all data on spare space of member
disks, proceed (yes/no)? yes
May 16 16:33:30 wgs57-06 scsi: /pci@0/pci@0/pci@2/scsi@0 (mpt0):
May 16 16:33:30 wgs57-06         Physical disk 0 created.
May 16 16:33:30 wgs57-06 scsi: /pci@0/pci@0/pci@2/scsi@0 (mpt0):
May 16 16:33:30 wgs57-06         Physical disk 1 created.
May 16 16:33:31 wgs57-06 scsi: /pci@0/pci@0/pci@2/scsi@0 (mpt0):
May 16 16:33:31 wgs57-06         Physical disk 2 created.
May 16 16:33:31 wgs57-06 scsi: /pci@0/pci@0/pci@2/scsi@0 (mpt0):
May 16 16:33:31 wgs57-06         Volume 3 is |enabled||optimal|
May 16 16:33:31 wgs57-06 scsi: /pci@0/pci@0/pci@2/scsi@0 (mpt0):
May 16 16:33:31 wgs57-06         Volume 3 is |enabled||optimal|
May 16 16:33:31 wgs57-06 scsi: /pci@0/pci@0/pci@2/scsi@0 (mpt0):
May 16 16:33:31 wgs57-06         Volume 3 created.
Volume c1t3d0 is created successfully!
#
```

When you create a RAID striped volume, the other member drives (in this case, c1t2d0 and c1t3d0) 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 c1t1d0 c1t2d0 c1t3d0
...
Volume c1t3d0 is created successfully!
#
```

3. To verify the existence of a RAID volume, type the following command:

```
# raidctl -l
Controller: 1
        Volume:c1t3d0
        Disk: 0.0.0
        Disk: 0.1.0
        Disk: 0.2.0
        Disk: 0.3.0
        Disk: 0.4.0
```

```
          Disk: 0.5.0
          Disk: 0.6.0
          Disk: 0.7.0
```

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

```
# raidctl -l c1t3d0
Volume                  Size    Stripe  Status   Cache  RAID
        Sub                     Size                    Level
                Disk
----------------------------------------------------------------
c1t3d0                  N/A     64K     OPTIMAL  OFF    RAID0
                0.3.0   N/A             GOOD
                0.4.0   N/A             GOOD
                0.5.0   N/A             GOOD
```

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

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

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

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

Related Information

-
-

▼ Configure a Hardware RAID Volume for the Solaris OS

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

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

```
# format
```

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

2. **Select the disk name that represents the RAID volume that you have configured.**

   In this example, `c1t2d0` is the logical name of the volume.

```
# format
Searching for disks...done
AVAILABLE DISK SELECTIONS:
       0. c1t0d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424>
          /pci@0/pci@0/pci@2/scsi@0/sd@0,0
       1. c1t1d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424>
          /pci@0/pci@0/pci@2/scsi@0/sd@1,0
       2. c1t2d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424>
          /pci@0/pci@0/pci@2/scsi@0/sd@2,0
       3. c1t3d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424>
          /pci@0/pci@0/pci@2/scsi@0/sd@3,0
       4. c1t4d0 <SUN73G cyl 14087 alt 2 hd 24 sec 424>
          /pci@0/pci@0/pci@2/scsi@0/sd@4,0
       5. c1t5d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424>
          /pci@0/pci@0/pci@2/scsi@0/sd@5,0
       6. c1t6d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424>
          /pci@0/pci@0/pci@2/scsi@0/sd@6,0
       7. c1t7d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424>
          /pci@0/pci@0/pci@2/scsi@0/sd@7,0
Specify disk (enter its number): 2
selecting c1t2d0
[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
```

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

For example:

```
format> type

VAILABLE 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. SUN72G
        21. SUN73G
        22. other
Specify disk type (enter its number)[19]: 0
c1t2d0: configured with capacity of 136.71GB
<SUN146G cyl 14087 alt 2 hd 24 sec 848>
selecting c1t2d0
[disk formatted]
```

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

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

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

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

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

```
format> disk

AVAILABLE DISK SELECTIONS:
        0. c1t0d0 <SUN72G cyl 14084 alt 2 hd 24 sec 424>
           /pci@0/pci@0/pci@2/scsi@0/sd@0,0
        1. c1t1d0 <SUN72G cyl 14084 alt 2 hd 24 sec 424>
           /pci@0/pci@0/pci@2/scsi@0/sd@1,0
        2. c1t2d0 <LSILOGIC-LogicalVolume-3000 cyl 65533 alt 2 hd
16 sec 273>
           /pci@0/pci@0/pci@2/scsi@0/sd@2,0
...
```

Note that `c1t2d0` now has a type indicating it is an `LSILOGIC-LogicalVolume`.

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

The volume can now be used in the Solaris OS.

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

Related Information

■ "Create a Hardware Mirrored Volume" on page 15
■ "Create a Hardware Mirrored Volume of the Default Boot Device" on page 19
■ "Create a Hardware Striped Volume" on page 20
■ "Delete a Hardware RAID Volume" on page 25

▼ Delete a Hardware RAID Volume

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

   See "Disk Slot Numbers" on page 35.

2. To determine the name of the RAID volume, type:

```
# raidctl
Controller: 1
Volume:c1t0d0
Disk: 0.0.0
Disk: 0.1.0
          ...
```

In this example, the RAID volume is c1t0d0.

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

---

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

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

For example:

```
# raidctl -d c1t0d0
Deleting RAID volume c1t0d0 will destroy all data it contains,
proceed (yes/no)? yes
/pci@0/pci@0/pci@2/scsi@0 (mpt0):
        Volume 0 deleted.
/pci@0/pci@0/pci@2/scsi@0 (mpt0):
        Physical disk 0 deleted.
/pci@0/pci@0/pci@2/scsi@0 (mpt0):
        Physical disk 1 deleted.
Volume c1t0d0 is deleted successfully!
```

If the RAID volume is an IS volume, the deletion of the RAID volume is interactive, for example:

```
# raidctl -d c1t0d0
Deleting volume c1t0d0 will destroy all data it contains, proceed
(yes/no)? yes
...
Volume c1t0d0 is deleted successfully!
#
```

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 c1t0d0
Volume c1t0d0 is deleted successfully!
#
```

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

```
# raidctl
```

For example:

```
# raidctl
Controller: 1
Disk: 0.0.0
Disk: 0.1.0
...
```

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

Related Information

- "Disk Slot Numbers" on page 35
- "Hot-Plug a Mirrored Disk" on page 28
- "Hot-Plug a Nonmirrored Disk" on page 30
- "Creating Hardware RAID Volumes" on page 14

▼ Hot-Plug a Mirrored Disk

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

See "Disk Slot Numbers" on page 35.

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

```
# raidctl
```

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

For example:

```
# raidctl -l c1t0d0
Volume                   Size    Stripe  Status    Cache  RAID
        Sub                      Size                     Level
                Disk
----------------------------------------------------------------
c1t0d0                   136.6G  N/A     DEGRADED  OFF    RAID1
                0.0.0    136.6G          GOOD
                0.1.0    136.6G          FAILED
```

This example indicates that the disk mirror has degraded due to a failure in disk `c1t2d0` (`0.1.0`).

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

---

3. **Remove the hard drive, as described in your server's service manual**.

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

4. **Install a new hard drive, as described in your server's service manual.**

The RAID utility automatically restores the data to the disk.

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

```
# raidctl
```

For example:

```
# raidctl -l c1t0d0
Volume                    Size    Stripe  Status   Cache  RAID
        Sub                       Size                    Level
                  Disk
----------------------------------------------------------------
c1t0d0                    136.6G  N/A     SYNC     OFF    RAID1
                  0.0.0   136.6G          GOOD
                  0.1.0   136.6G          GOOD
```

This example indicates that RAID volume c1t1d0 is resynchronizing.

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

```
# raidctl -l c1t0d0
Volume                    Size    Stripe  Status   Cache  RAID
        Sub                       Size                    Level
                  Disk
----------------------------------------------------------------
c1t0d0                    136.6G  N/A     OPTIMAL  OFF    RAID1
                  0.0.0   136.6G          GOOD
                  0.1.0   136.6G          GOOD
```

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

Related Information

■ "Disk Slot Numbers" on page 35
■ "Hot-Plug a Nonmirrored Disk" on page 30

▼ Hot-Plug a Nonmirrored Disk

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

See "Disk Slot Numbers" on page 35.

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

2. Type the following command:

```
# cfgadm –al
```

For example:

```
# cfgadm –al
Ap_Id            Type        Receptacle  Occupant      Condition
c1               scsi-bus    connected   configured    unknown
c1::dsk/c1t0d0   disk        connected   configured    unknown
c1::dsk/c1t1d0   disk        connected   configured    unknown
c1::dsk/c1t2d0   disk        connected   configured    unknown
c1::dsk/c1t3d0   disk        connected   configured    unknown
c1::dsk/c1t4d0   disk        connected   configured    unknown
c1::dsk/c1t5d0   disk        connected   configured    unknown
c1::dsk/c1t6d0   disk        connected   configured    unknown
c1::dsk/c1t7d0   disk        connected   configured    unknown
usb0/1           unknown     empty       unconfigured ok
usb0/2           unknown     empty       unconfigured ok
usb0/3           unknown     empty       unconfigured ok
usb1/1           unknown     empty       unconfigured ok
usb1/2           unknown     empty       unconfigured ok
usb2/1           unknown     empty       unconfigured ok
usb2/2           usb-storage connected   configured    ok
usb2/3           unknown     empty       unconfigured ok
usb2/4           usb-hub     connected   configured    ok
usb2/4.1         unknown     empty       unconfigured ok
usb2/4.2         unknown     empty       unconfigured ok
usb2/4.3         unknown     empty       unconfigured ok
usb2/4.4         unknown     empty       unconfigured ok
usb2/5           unknown     empty       unconfigured ok
#
```

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

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

```
# cfgadm -x remove_device c1::dsk/c1t3d0
Removing SCSI device: /devices/pci@0/pci@0/pci@2/scsi@0/sd@3,0
This operation will suspend activity on SCSI bus: c1
Continue (yes/no)? yes
cfgadm: Hardware specific failure: failed to suspend:
     Resource              Information
------------------  -----------------------
/dev/dsk/c1t1d0s0   mounted filesystem "/"
```

This warning is issued because these commands attempt to quiesce the (SAS) SCSI bus, but the server firmware prevents it. This warning message can be safely ignored in the Sun SPARC Enterprise T5120 or T5220 server, but the following step avoids this warning message altogether.

3. Remove the hard drive from the device tree.

Type the following command:

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

For example:

```
# cfgadm -c unconfigure c1::dsk/c1t3d0
```

This example removes c1t3d0 from the device tree. The blue OK-to-Remove LED lights.

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

Type the following command:

```
# cfgadm -al
Ap_Id              Type         Receptacle   Occupant     Condition
c1                 scsi-bus     connected    configured   unknown
c1::dsk/c1t0d0     disk         connected    configured   unknown
c1::dsk/c1t1d0     disk         connected    configured   unknown
c1::dsk/c1t2d0     disk         connected    configured   unknown
c1::dsk/c1t3d0     disk         connected    unconfigured unknown
c1::dsk/c1t4d0     disk         connected    configured   unknown
c1::dsk/c1t5d0     disk         connected    configured   unknown
c1::dsk/c1t6d0     disk         connected    configured   unknown
c1::dsk/c1t7d0     disk         connected    configured   unknown
usb0/1             unknown      empty        unconfigured ok
usb0/2             unknown      empty        unconfigured ok
usb0/3             unknown      empty        unconfigured ok
usb1/1             unknown      empty        unconfigured ok
usb1/2             unknown      empty        unconfigured ok
usb2/1             unknown      empty        unconfigured ok
usb2/2             usb-storage  connected    configured   ok
usb2/3             unknown      empty        unconfigured ok
usb2/4             usb-hub      connected    configured   ok
usb2/4.1           unknown      empty        unconfigured ok
usb2/4.2           unknown      empty        unconfigured ok
usb2/4.3           unknown      empty        unconfigured ok
usb2/4.4           unknown      empty        unconfigured ok
usb2/5             unknown      empty        unconfigured ok
#
```

Note that c1t3d0 is now unknown and unconfigured. The corresponding hard drive OK-to-Remove LED is lit.

5. Remove the hard drive, as described in your server's service manual.

The blue OK-to-Remove LED is extinguished when you remove the hard drive.

6. Install a new hard drive, as described in your server's service manual.

7. Configure the new hard drive.

Type the following command:

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

For example:

```
# cfgadm -c configure c1::dsk/c1t3d0
```

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

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

Type the following command:

```
# cfgadm -al
Ap_Id             Type         Receptacle    Occupant      Condition
c1                scsi-bus     connected     configured    unknown
c1::dsk/c1t0d0    disk         connected     configured    unknown
c1::dsk/c1t1d0    disk         connected     configured    unknown
c1::dsk/c1t2d0    disk         connected     configured    unknown
c1::dsk/c1t3d0    disk         connected     configured    unknown
c1::dsk/c1t4d0    disk         connected     configured    unknown
c1::dsk/c1t5d0    disk         connected     configured    unknown
c1::dsk/c1t6d0    disk         connected     configured    unknown
c1::dsk/c1t7d0    disk         connected     configured    unknown
usb0/1            unknown      empty         unconfigured ok
usb0/2            unknown      empty         unconfigured ok
usb0/3            unknown      empty         unconfigured ok
usb1/1            unknown      empty         unconfigured ok
usb1/2            unknown      empty         unconfigured ok
usb2/1            unknown      empty         unconfigured ok
usb2/2            usb-storage  connected     configured    ok
usb2/3            unknown      empty         unconfigured ok
usb2/4            usb-hub      connected     configured    ok
usb2/4.1          unknown      empty         unconfigured ok
usb2/4.2          unknown      empty         unconfigured ok
usb2/4.3          unknown      empty         unconfigured ok
usb2/4.4          unknown      empty         unconfigured ok
usb2/5            unknown      empty         unconfigured ok
#
```

Note that c1t3d0 is now listed as configured.

Related Information

- "Disk Slot Numbers" on page 35
- "Hot-Plug a Mirrored Disk" on page 28

Disk Slot Numbers

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

These error messages typically refer to a failed hard drive by its physical device name (such as `/devices/pci@1f,700000/scsi@2/sd@1,0`) or by its logical device name (such as `c1t1d0`). In addition, some applications might report a disk slot number (0 through 3).

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

Disk Slot Number	Logical Device Name[*]	Physical Device Name
Slot 0	c1t0d0	/devices/pci@0/pci@0/pci@2/scsi@0/sd@0,0
Slot 1	c1t1d0	/devices/pci@0/pci@0/pci@2/scsi@0/sd@1,0
Slot 2	c1t2d0	/devices/pci@0/pci@0/pci@2/scsi@0/sd@2,0
Slot 3	c1t3d0	/devices/pci@0/pci@0/pci@2/scsi@0/sd@3,0

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

Related Information

- "Managing Disks" on page 13

Managing Devices

This section contains information about managing devices in the servers, and the multipathing software supported.

▼ Unconfigure a Device Manually

The ILOM firmware provides a set *Device-Identifier* component_state= disabled command, which enables you to unconfigure system devices manually. This command marks the specified device as *disabled*. Any device marked disabled, whether manually or by the system firmware, is removed from the system's machine description prior to transfer of control to other layers of system firmware, such as OpenBoot PROM.

1. "Log In to ILOM" on page 2.

2. **At the ILOM -> prompt, type:**

```
-> set Device-Identifier component_state=disabled
```

Related Information
- "Reconfigure a Device Manually" on page 38
- "Devices and Device Identifiers" on page 38

▼ Reconfigure a Device Manually

The ILOM firmware provides a set *Device-Identifier* component_state=enabled command, which enables you to reconfigure system devices manually. Use this command to mark the specified device as *enabled*.

1. "Log In to ILOM" on page 2.

2. **At the ILOM -> prompt, type:**

```
-> set Device-Identifier component_state=enabled
```

Related Information

■ "Devices and Device Identifiers" on page 38

■ "Unconfigure a Device Manually" on page 37

Devices and Device Identifiers

The device identifiers are case-sensitive.

Device Identifiers	Devices
/SYS/MB/CMP*cpu_number*/P*strand_number*	CMP (0-1) CPU strand (0-63)
/SYS/MB/CMP*n*/RISER*n*/PCIE*slot_number*	CMP (0-1) Riser (0-1) PCIe slot (0-5)
/SYS/MB/CMP*n*/RISER*n*/XAUI*card_number*	CMP (0-1) Riser (0-1) XAUI card (0-1)
/SYS/MB/NET*network_number*	Network interfaces (0-3)
/SYS/MB/PCIE	PCIe root complex
/SYS/MB/USB*number*	USB ports (0-1, located on rear of chassis)
/SYS/MB/CMP*n*/L2_BANK*number*	CMP (0-1) Bank (0-7)
/SYS/DVD	DVD

Device Identifiers *(Continued)*	Devices *(Continued)*
/SYS/USBBD/USB*number*	USB ports (2-3, located on front of chassis)
/SYS/TTYA	DB9 Serial Port
/SYS/MB/CMP*n*/MR*n*/BR/*branch_number*/CH*channel_number*/D*dimm_number*	CMP (0-1) Riser (0-1) Branch (0-1) Channel (0-1) DIMM (0-3)

Related Information

Sun SPARC Enterprise T5x40 Device Tree

The following table shows the correspondence of the Sun SPARC Enterprise T5140 and T5240 server devices to the Solaris Operating System device tree.

Device (as Indicated on Chassis Label)	Solaris OS Device Tree
DVD Drive	/pci@400/pci@0/pci@1/pci@0/usb@0,2/storage@2/disk
HDD disk[0-*n*]*	/pci@400/pci@0/pci@8/scsi@0/disk@[0-*n*]
NET 0	/pci@500/pci@0/pci@8/network@0
NET 1	/pci@500/pci@0/pci@8/network@0,1
NET 2	/pci@500/pci@0/pci@8/network@0,2
NET 3	/pci@500/pci@0/pci@8/network@0,3
PCIe 0	/pci@500/pci@0/pci@9
PCIe 1	/pci@400/pci@0/pci@c
PCIe 2	/pci@400/pci@0/pci@9
PCIe 3 (T5240 only)	/pci@400/pci@0/pci@d
PCIe 4 (T5240 only)	/pci@500/pci@0/pci@d
PCIe 5 (T5240 only)	/pci@500/pci@0/pci@c

Device (as Indicated on Chassis Label)	Solaris OS Device Tree
USB 0 (rear) USB 1.x	`/pci@400/pci@0/pci@1/pci@0/usb@0/storage@1`[†]
USB 2.0	`/pci@400/pci@0/pci@1/pci@0/usb@0,2/storage@1`
USB 1 (rear) USB 1.x	`/pci@400/pci@0/pci@1/pci@0/usb@0,1/storage@2`
USB 2.0	`/pci@400/pci@0/pci@1/pci@0/usb@0,1/storage@3`
USB 2 (front)	`/pci@400/pci@0/pci@1/pci@0/usb@0,2/hub@4/storage@1`
USB 3 (front)	`/pci@400/pci@0/pci@1/pci@0/usb@0,2/hub@4/storage@2`
XAUI 0 (PCIe 0 slot)	`/pci@500/pci@0/pci@8/pci@0`
XAUI 1 (PCIe 1 slot)	`/pci@500/pci@0/pci@8/pci@0,1`

* The number of hard drives varies, based on the model of the server.

† The USB node string (`storage`) changes based on the kind of device that is connected to the USB port. For example, if you connect a keyboard, the `storage` string changes to `keyboard`.

Related Information

-
-
-

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 the Sun SPARC Enterprise T5120 or T5220 servers, 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 StorageTek™ 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.

Related Information

- 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 StorageTek Traffic Manager, refer to your Solaris OS documentation.

Handling Faults

The Sun SPARC Enterprise T5140 and T5240 servers provide many ways to find faults, including LEDs, ILOM and POST. For specific information about LEDs, and additional troubleshooting information, refer to the service manual for your server.

- "Discovering Faults" on page 43
- "Bypassing Minor Faults" on page 46
- "Clear a Fault" on page 49

Related Information

- "Performing Common Tasks" on page 7

Discovering Faults

This section contains information about finding system faults using pre-OS tools, including ILOM and POST.

- "Discover Faults Using ILOM" on page 44
- "Discover Faults Using POST" on page 44
- "Locate the System" on page 45

Related Information

- "Bypassing Minor Faults" on page 46

▼ Discover Faults Using ILOM

● **Type:**

```
-> show /SP/faultmgmt
```

This command displays the fault ID, the faulted FRU device, and the fault message to standard output. The show /SP/faultmgmt command also displays POST results.

For example:

```
-> show /SP/faultmgmt
 /SP/faultmgmt
Targets:
0 (/SYS/PS1)
Properties:
Commands:
cd
show
->
```

For more information about the show /SP/faultmgmt command, refer to the ILOM guide and the ILOM supplement for your server.

Related Information

▼ Discover Faults Using POST

The virtual keyswitch can be used to run full POST diagnostics without having to modify the diagnostic property settings. Note that POST diagnostics can take a significant amount of time to run at system reset.

1. **.**

2. **At the ILOM -> prompt, type:**

```
-> set /SYS keyswitch_state=diag
```

The system is set to run full POST diagnostics on system reset.

3. **To return to your normal diagnostic settings *after* running POST, type:**

```
-> set /SYS keyswitch_state=normal
```

Related Information

- "Discover Faults Using ILOM" on page 44
- "Locate the System" on page 45
- "Clear a Fault" on page 49
- "Bypassing Minor Faults" on page 46

▼ Locate the System

1. **To turn on the Locator LED, from the ILOM service processor command prompt, type:**

```
-> set /SYS/LOCATE value=Fast_Blink
```

2. **To turn off the Locator LED, from the ILOM service processor command prompt, type:**

```
-> set /SYS/LOCATE value=off
```

3. **To display the state of the Locator LED, from the ILOM service processor command prompt, type:**

```
-> show /SYS/LOCATE
```

Note – You do not need administrator permissions to use the set /SYS/LOCATE and show /SYS/LOCATE commands.

Related Information

- "Discover Faults Using ILOM" on page 44
- "Discover Faults Using POST" on page 44

Bypassing Minor Faults

This section includes information about configuring your server to automatically recover from minor faults.

- "Automatic System Recovery" on page 46
- "Enable ASR" on page 47
- "Disable ASR" on page 47
- "View Information on Components Affected by ASR" on page 48

Related Information

- "Discovering Faults" on page 43

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 "Enable ASR" on page 47.

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

Related Information

- "Enable ASR" on page 47
- "Disable ASR" on page 47
- "View Information on Components Affected by ASR" on page 48

▼ Enable ASR

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

```
-> set /HOST/diag mode=normal
-> set /HOST/diag level=max
-> set /HOST/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 service manual for your server.

---

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).

Related Information

▼ Disable ASR

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.

After you disable the ASR feature, it is not activated again until you re-enable it.

Related Information

- "View Information on Components Affected by ASR" on page 48
- "Automatic System Recovery" on page 46
- "OpenBoot Configuration Variables on the SCC" on page 53

▼ View Information on Components Affected by ASR

● **At the -> prompt, type:**

```
-> show /SYS/component component_state
```

In the show /SYS/*component* component_state command output, any devices marked disabled have been manually unconfigured using the system firmware. The command output also shows devices that have failed firmware diagnostics and have been automatically unconfigured by the system firmware.

Related Information

- "Automatic System Recovery" on page 46
- "Enable ASR" on page 47
- "Disable ASR" on page 47
- "Unconfigure a Device Manually" on page 37
- "Reconfigure a Device Manually" on page 38

▼ Clear a Fault

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

```
-> set /SYS/component clear_fault_action=true
```

Setting `clear_fault_action` to `true` clears the fault at the component and all levels below it in the `/SYS` tree.

Related Information

Managing Logical Domains Software

Oracle's Sun SPARC Enterprise servers support the Logical Domains (LDoms) software that is used to create and manage logical domains. The software comprises LDoms-enabling code in the Solaris OS, LDoms-enabling code in System Firmware, and the Logical Domains Manager, which is the command-line interface. See your LDoms documentation for the latest information.

- "Logical Domains Software" on page 51
- "Logical Domain Configurations" on page 52

Logical Domains Software

LDoms software enables you to create and manage as many as 32 logical domains, depending on the hardware configuration of the server on which the Logical Domains Manager has been installed. You can virtualize resources and define network, storage, and other I/O devices as services that can be shared between domains.

A *logical domain* is a discrete logical grouping with its own operating systems, resources, and identity within a single computer system. Applications software can run in logical domains. Each logical domain can be created, destroyed, reconfigured, and rebooted independently. There are several roles that logical domains can perform as shown in the following table.

TABLE: Logical Domain Roles

Domain Role	Description
Control domain	Domain in which the Logical Domains Manager runs, enabling you to create and manage other logical domains and allocate virtual resources to other domains. There can be only one control domain per server. The initial domain created when installing Logical Domains software is a control domain and is named primary.
Service domain	Domain that provides virtual device services to other domains, such as a virtual switch, a virtual console concentrator, and a virtual disk server.
I/O domain	Domain that has direct ownership of and direct access to physical I/O devices, such as a network card in a PCI Express controller. Shares the devices to other domains in the form of virtual devices. You can have a maximum of two I/O domains, one of which also must be the control domain.
Guest domain	Domain that is managed by the control domain and uses services from the I/O and service domains.

Related Information

■ "Logical Domain Configurations" on page 52

Logical Domain Configurations

The Logical Domain configurations are stored on the service processor (SP). Using Logical Domains Manager CLI commands, you can add a configuration, specify a configuration to be used, and list the configurations on the service processor. You can also use the ILOM `set /HOST/bootmode config=`*configfile* command to specify an LDoms boot configuration. For further information about `/HOST/bootmode`, see your server's ILOM supplement.

Related Information

■ "Logical Domains Software" on page 51

OpenBoot Configuration Variables

This section supplies information about variables storing configuration on the SCC.

-

OpenBoot Configuration Variables on the SCC

TABLE: OpenBoot Configuration Variables Stored on the System Configuration Card on page 53 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 type the following command:

```
ok printenv
```

TABLE: 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 the server manufacturer's logo.
oem-banner?	true, false	false	If true, use custom OEM banner.
ansi-terminal?	true, false	true	If true, enable ANSI terminal emulation.

TABLE: OpenBoot Configuration Variables Stored on the System Configuration Card *(Continued)*

Variable	Possible Values	Default Value	Description
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.
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.
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. network-boot-arguments can be used to specify the boot protocol (RARP/DHCP) to be used and a range of system knowledge to be used in the process. For further information, see the eeprom (1M) man page or your Oracle Solaris Reference Manual.

Index