SPARC T3 Series Servers

Administration Guide

Adobe PostScript™

Contents

Using This Documentation

This administration guide is for experienced system administrators of the Oracle SPARC T3 server series. The guide includes general descriptive information about the server, and detailed instructions for configuring and administering the server. 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 (Oracle Solaris OS).

Note – The *SPARC T3 Series Servers Administration Guide* applies to several server and server module products. Some examples used in this document may be based on particular server models. Your output may vary from the examples based on your product.

- "Related Documentation" on page x
- "Feedback" on page x
- "Support and Accessibility" on page x

Related Documentation

Documentation	Links
All Oracle products	http://www.oracle.com/documentation
Oracle Solaris OS and other systems software	http://www.oracle.com/technetwork/indexes/documentation/index.html#sys_sw
Oracle Integrated Lights Out Manager (ILOM) 3.0	http://www.oracle.com/pls/topic/lookup?ctx=ilom30
Oracle VTS 7.0	http://www.oracle.com/pls/topic/lookup?ctx=OracleVTS7.0

Feedback

Provide feedback on this documentation at:

http://www.oracle.com/goto/docfeedback

Support and Accessibility

Description	Links
Access electronic support through My Oracle Support	http://support.oracle.com
	For hearing impaired:
	http://www.oracle.com/accessibility/support.html
Learn about Oracle's commitment to accessibility	http://www.oracle.com/us/corporate/accessibility/index.html

Understanding System Administration Resources

These topics provide a summary of common tools used to administer the server.

Oracle ILOM Overview

Oracle Integrated Lights Out Manager (Oracle ILOM) is system management firmware that is preinstalled on some SPARC servers. Oracle ILOM enables you to actively manage and monitor components installed in your server. Oracle ILOM provides a browser-based interface and a command-line interface, as well as SNMP and IPMI interfaces.

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

The -> prompt indicates that you are interacting with the Oracle ILOM SP directly. This prompt is the first prompt you see when you log in to the server through the SER MGT port or NET MGT port, regardless of the host's power state. On a modular system, that prompt is also present when you log in to a server module either directly or through Oracle ILOM on the CMM of the modular system.

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

The Oracle ILOM service processor supports a total of ten concurrent sessions per server: nine SSH or web connections are available through the NET MGT port and one connection is available through the SER MGT port.

For more information about how to work with Oracle ILOM features that are common to all platforms managed by Oracle ILOM, see the Oracle ILOM documentation at:

http://www.oracle.com/pls/topic/lookup?ctx=ilom30

Related Information

- "Platform-Specific Oracle ILOM Features" on page 2
- "Log In to Oracle ILOM" on page 9

Platform-Specific Oracle ILOM Features

ILOM operates on many platforms, supporting features that are common to all platforms. Some ILOM features belong to only a subset of platforms. This topic describes the difference between ILOM features supported on the server and the common set of features described in the ILOM 3.0 base documentation.

Note – To perform some procedures documented in Oracle's ILOM 3.0 base documentation, you must create a serial connection to the server and activate the Physical Presence switch on the server. For information about creating a serial connection, see the installation guide for your server.

Among the Oracle ILOM features supported on other platforms, Oracle ILOM does *not* support the following features on this server:

- Chassis monitoring module (CMM) features, such as single signon.

Note – The T3 blade servers in a modular system do support the CMM features.

- POST diagnostics user-reset trigger is not available.

Oracle ILOM supports the following feature on this server, which may not be available on other platforms:

- POST diagnostics `hw-change` trigger. This new trigger (`hw-change error-reset`) is the *default setting* for the server, and causes POST to run each time the server is AC power-cycled, or the top cover is removed (if applicable). For more information on POST, see the service manual for your server.

Related Information

- "Oracle ILOM Overview" on page 1

Oracle Solaris OS Overview

The Oracle Solaris OS includes commands and other software resources to use for server administration. For an introduction to management tools in your Oracle Solaris release, see *System Administration Guide: Basic Administration* in the Oracle Solaris documentation collection.

Oracle Solaris software includes SunVTS software. SunVTS tests and validates Oracle hardware by verifying the connectivity and functionality of hardware devices, controllers and peripherals.

In addition to the SunVTS information in the Oracle Solaris documentation, SunVTS documentation collections are available at:

http://www.oracle.com/pls/topic/lookup?ctx=E19719-01&id=homepage

Related Information

- "OpenBoot Overview" on page 3

OpenBoot Overview

The OpenBoot firmware starts the OS, validates installed hardware, and can be used for other server administration tasks below the OS level. For more information about OpenBoot commands, see the *OpenBoot 4.x Command Reference Manual* in the Oracle Solaris documentation collection.

Related Information

- "Oracle Solaris OS Overview" on page 3

Oracle VM Server for SPARC Overview

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.

Oracle VM Server for SPARC software enables you to create and manage as many as 32 logical domains, depending on the hardware configuration of the server on which the Oracle VM Server for SPARC 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.

The Oracle VM Server for SPARC configurations are stored on the SP. Using Oracle VM Server for SPARC 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 Oracle VM Server boot configuration.

Related Information

- "Configuring Boot Mode" on page 53
- Oracle VM Server for SPARC documentation

  `http://www.oracle.com/technetwork/documentation/vm-sparc-19428`
  `7.html`

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 server, three different types of multipathing software are available:

- Oracle Solaris IP Network Multipathing software provides multipathing and load-balancing capabilities for IP network interfaces. For instructions on how to configure and administer Oracle Solaris IP Network Multipathing, consult the *IP Network Multipathing Administration Guide* provided with your specific Oracle Solaris release.

- VVM software includes a feature called DMP, which provides disk multipathing as well as disk load balancing to optimize I/O throughput. For information about VVM and its DMP feature, refer to the documentation provided with the VERITAS Volume Manager software.

- StorageTek Traffic Manager is an architecture fully integrated within the Oracle Solaris OS (beginning with the Oracle 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 information about StorageTek Traffic Manager, refer to your Oracle Solaris OS documentation.

Related Information

- "Oracle Solaris OS Overview" on page 3
- "Oracle VM Server for SPARC Overview" on page 4

Hardware Management Pack Overview

The Sun Server Hardware Management Pack (HMP) from Oracle provides tools you can use to manage and configure Oracle servers from the host operating system. To use these tools, you must install the HMP software on your server. After installing the software, you will be able to perform the following server management tasks described in the following table.

TABLE: Hardware Management Pack -- Server Managment Tasks

Server Management Task From Host OS[*]	Hardware Management Pack Implementation	Tool
Monitor Oracle hardware with host IP address	Use the Hardware Management Agent and the associated Simple Network Management Protocol (SNMP) plug-ins at the operating-system level to enable in-band monitoring of your Oracle hardware. This in-band monitoring functionality enables you to use your host operating system IP address to monitor your Oracle servers without the need of connecting the Oracle ILOM management port to your network.	Host OS-level management tool
Monitor storage devices, including RAID arrays	Use the Server Storage Management Agent at the operating-system level to enable in-band monitoring of the storage devices configured on your Oracle servers. The Server Storage Management Agent provides an operating-system daemon that gathers information about your server's storage devices, such as hard disk drives (HDDs) and RAID arrays, and sends this information to the Oracle ILOM service processor. The Storage Monitoring features in Oracle ILOM enable you to view and monitor the information provided by the Server Storage Management Agent. You can access the Storage Monitoring features in Oracle ILOM from the command-line interface (CLI).	Oracle ILOM 3.0 CLI Storage Monitoring features
Query, update, and validate firmware versions on supported SAS storage devices	Use the `fwupdate` CLI tool from the host operating system to query, update, and validate firmware versions on supported storage devices such as SAS host bus adapters (HBAs), embedded SAS storage controllers, LSI SAS storage expanders, and disk drives.	Host OS-level `fwupdate` CLI
Restore, set, and view Oracle ILOM configuration settings	Use the `ilomconfig` CLI tool from the host operating system to restore Oracle ILOM configuration settings,as well as to view and set Oracle ILOM properties associated with network management, clock configuration, and user management.	Host OS-level `ilomconfig` CLI
View or create RAID volumes on storage drives	Use the `raidconfig` CLI tool from the host operating system to view and create RAID volumes on storage drives that are attached to RAID controllers, including storage arrays.	Host OS-level `raidconfig` CLI
Use IPMItool to access and manage Oracle servers	Use the open source command-line IPMItool from the host operating system to access and manage your Oracle servers via the IPMI protocol.	Host OS-level command-line IMPItool

* Supported host operating systems include: Solaris, Linux, Windows, and VMware

Source for Downloading Hardware Management Pack Software

You can download the Hardware Management Pack software from the following location:

http://support.oracle.com

Hardware Management Pack Documentation

You can download Hardware Management Pack documentaton from the following location:

http://www.oracle.com/pls/topic/lookup?ctx=ohmp

For additional details about how to use the Storage Monitoring features in Oracle ILOM, see the *Oracle Integrated Lights Out Manager (ILOM) 3.0 Concepts Guide* and the *Oracle Integrated Lights Out Manager (ILOM) 3.0 Daily Management CLI Procedures Guide*.

For additional details about accessing and managing your server via SNMP or IPMI, see the *Oracle Integrated Lights Out Manager (ILOM) 3.0 SNMP, IPMI, CIM, WSMAN, Protocol Management Reference*.

Links to these Oracle ILOM manuals are provided on the web site shown above. You can find the complete set of Oracle ILOM documentation at this location:

http://www.oracle.com/pls/topic/lookup?ctx=ilom30

Accessing the Server

These topics include information on establishing low-level communication with the server using the Oracle ILOM tool and the system console.

▼ Log In to Oracle ILOM

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

Note – For a SPARC T3 server module, you can either log in directly to the modular server's SP or you can start Oracle ILOM through the chassis' CMM. Refer to the modular server's installation guide for instructions on both approaches.

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

The Oracle 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)

Integrated Lights Out Manager

Version 3.0.12.1 r57146

Copyright 2010 Oracle and/or its affiliates. All rights reserved.

->
```

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

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

Related Information

▼ Log In to the System Console

● **At the Oracle ILOM -> prompt, type:**

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

where *option* can be:

- -f | force – Enables a user with a Console (c) role to take the console from any current user and force that user into view mode.

■ -script – Bypasses the prompt for a yes or no confirmation.

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

Related Information

▼ 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.**

To ensure getting to the ok prompt, set the following Oracle ILOM property before performing the procedures described in the table:

```
-> set /HOST/bootmode script="setenv auto-boot? false"
```

Caution – When possible, reach the ok prompt by performing a graceful shutdown of the OS. Using any other method might cause the loss of server state information.

Server State	What To Do
OS running and responsive	**Shut down the server using one of these methods:** From a shell or command tool window, type an appropriate command (for example, the shutdown or init 0 command) as described in the Oracle Solaris system administration documentation. Then take one of the following steps: • At the ILOM -> prompt, type: -> **stop /SYS** • Use the server Power button.
OS unresponsive	**Disable autoboot and thereset the host.** (Provided the operating system software is not running and the server is already under OpenBoot firmware control.) At the Oracle ILOM prompt, type: -> **set /HOST send_break_action=break** Press Enter. Then type: -> **start /HOST/console**
OS unresponsive and need to prevent auto-boot	**Shut down the server from ILOM and disable autoboot.** At the ILOM -> prompt, type: -> **set /HOST/bootmode script="setenv auto-boot? false"** Press Enter. Then type: -> **reset /SYS** -> **start /HOST/console**

Related Information

- "Monitoring Faults" on page 65

▼ Display the Oracle ILOM -> Prompt

- **Use one of the following ways to display the Oracle ILOM -> prompt:**

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

  - Log in to Oracle ILOM directly from a device connected to the SER MGT port or the NET MGT port.

  - Log in to Oracle ILOM through an SSH connection. See "Log In to Oracle ILOM" on page 9.

Related Information

▼ Use a Local Graphics Monitor

You can redirect the system console to a local graphic monitor. You *cannot* use a local graphics monitor to perform the initial server installation, nor can you use a local graphics monitor to view POST messages.

To use a local graphics monitor:

1. **Connect the monitor video cable to a video port on the server.**

   Tighten the thumbscrews to secure the connection. Refer to your system's documentation for any special connection instructions that might apply to your server.

2. **Connect the monitor power cord to an AC outlet.**

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

4. **Connect the USB mouse cable to another USB port on the server.**

5. **Display the** ok **prompt.**

   See "Display the ok Prompt" on page 11.

6. **At the** ok **prompt, set the following OpenBoot PROM configuration variables:**

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

7. **Make the changes take effect:**

   ```
   ok reset-all
   ```

   The server stores the parameter changes and boots automatically.

Note – Instead of using the reset-all command to store the parameter changes, you can also power cycle the server using the Power button.

You can now type system commands and view system messages using your local graphics monitor. To activate the GUI interface, continue to the next step.

8. **Activate the Oracle Solaris OS GUI interface.**

   Once the Oracle Solaris OS is installed and booted, type the following commands to display the GUI login screen.

   ```
   # ln -s /dev/fbs/ast0 /dev/fb
   ```

   ```
   # fbconfig -xserver Xorg
   ```

   ```
   # reboot
   ```

Related Information

- "Display the ok Prompt" on page 11.
- "Oracle ILOM Remote Console" on page 14

Oracle ILOM Remote Console

Oracle ILOM Remote Console is a Java application that enables you to remotely redirect and control the following devices on the host server. This group of devices is commonly abbreviated as KVMS.

- Keyboard
- Video console display
- Mouse
- Serial console display
- Storage devices or images (CD/DVD)

Oracle ILOM Remote Console is documented in the *Oracle Integrated Lights Out Manager (ILOM) 3.0 Daily Management Web Interface Procedures Guide* ("Managing Remote Hosts Redirection and Securing the Oracle ILOM Remote Console").

Related Information

- "Oracle ILOM In-band (Sideband) Management" on page 48

Controlling the Server

These topics include procedures for contolling basic server operations.

-
-
-
-
-

▼ Power On the Server

1. **Log in to ILOM.**

   "Log In to Oracle ILOM" on page 9.

   ***

   Note – If you have a modular system, make sure you are logged in to the desired server module.

   ***

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

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

->
```

***

Note – To skip being prompted for confirmation, use the `start -script /SYS` command.

***

Related Information

- "Power Off the Server" on page 16

▼ Power Off the Server

1. Shut down the Oracle Solaris OS.

At the Oracle 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?
# o
```

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

```
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 typing this command.

Related Information

- "Power On the Server" on page 15
- "Reset the Server From the Oracle Solaris OS" on page 17
- "Reset the Server From Oracle ILOM" on page 17

▼ Reset the Server From the Oracle Solaris OS

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

● **To reset the server from the Oracle Solaris prompt, type one of the following commands:**

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

or

```
# reboot
```

Related Information

■ "Power Off the Server" on page 16

■ "Power On the Server" on page 15

■ "Reset the Server From Oracle ILOM" on page 17

▼ Reset the Server From Oracle ILOM

The ILOM reset command generates a graceful or forced hardware reset of the server. By default, the reset command gracefully resets the server.

● **Type one of the following commands to reset the server.**

  ■ Perform a graceful reset from ILOM:

```
-> reset /SYS
```

  ■ If a graceful reset is not possible, perform a forced hardware reset from ILOM:

```
-> reset -force /SYS
```

Related Information

■ "Power Off the Server" on page 16

■ "Power On the Server" on page 15

▼ Reset the SP to Default Values

If your SP becomes corrupt, or you want to reset the SP to the factory default values, change the /SP reset_to_defaults setting then power off the host to implement the changes. This is new behavior. Previously you did not have to power off the host to reset default values to the SP. You need administrator permissions to perform this task.

1. To reset the SP to the default values, type:

```
-> set /SP reset_to_defaults=value
```

where *value* can be:

- all – Removes all of the SP configuration data.
- factory – Removes all SP configuration data as well as all log files.

2. Power off and restart the host to complete the setting change.

```
-> stop /SYS
-> reset /SP
```

Related Information

- "Power Off the Server" on page 16

Configuring Hardware RAID

These topics describe how to configure and manage RAID disk volumes using the server's onboard SAS2 disk controllers.

Hardware RAID Support

In the SPARC and Netra SPARC T3 series servers and blades, hardware RAID support is provided by either onboard SAS2 RAID controllers or pluggable RAID expansion modules (REMs). Both types of controller enable you to create logical disk volumes consisting of two or more redundant disk drives. The following table identifies the type of hardware RAID controller used by each T3 platform.

TABLE: Hardware RAID Support for T3 Series Servers and Server Modules

Platform	Controller Description	Supported RAID Levels
SPARC T3-1	Two onboard SAS2 controllers	0, 1, 1e
SPARC T3-2	One onboard SAS2 controller	0, 1, 1e
SPARC T3-4	Two pluggable SAS2 REMs	0, 1
SPARC T3-1B	Two pluggable SAS2 REMs	0, 1

Note – Up to two logical volume can be configured on each RAID controller.

There are three sets of RAID management utilities available for use with the onboard SAS2 controllers provided on SPARC T3-1 and T3-2 servers:

- Fcode utility -- You can use Fcode commands to display targets and manage logical volumes on your server. These commands are accessed through OBP.

  The examples shown in this manual rely on Fcode commands.

- Oracle Hardware Management Pack -- You can use the raidconfig commands contained in this software's Oracle Server CLI Tools component to create and manage RAID volumes on your server. To use these commands, download and install the Oracle Hardware Management Pack version 2.2 or later. Refer to the *Oracle Hardware Management Pack 2.2 Installation Guide* for instructions.

  You can find documentation for the Oracle Hardware Management Pack software at this location:

  http://www.oracle.com/pls/topic/lookup?ctx=mgtpk22

- LSI SAS2 2008 RAID Management Utility (sas2ircu) -- You can use sas2ircu commands (version 6.250.02.00 or later) to manage RAID functionality for the onboard SAS2 controllers.

  The sas2ircu utility for Oracle's T3-1 and T3-2 servers, together with user documentation, are available on the LSI website:

  http://www.lsi.com/sep/Pages/oracle/index.aspx

Note – You can download documentation for the REM controllers from:
http://docs.oracle.com/cd/E19946-01/index.html

Caution – Firmware updates for the onboard SAS controllers *must* be acquired from Oracle's support site http://support.oracle.com. Firmware applied to the onboard SAS controllers from any other location or any vendor other than Oracle is not supported.

Related Information

- "Create a Hardware RAID Volume" on page 30
- "Prepare to Use the FCode Utility" on page 28

Important Hardware RAID Guidelines

The following points are important to understand when configuring RAID volumes on a SPARC T3 series server:

- Before configuring and using RAID disk volumes on a SPARC T3 series server, ensure that you have installed the latest patches or SRU available for your operating system. Review the most recent product notes provided for your system for further guidance in keeping your OS up to date.
- Volume migration (relocating all RAID volume disk members from one T3 series server to another) is not supported. If you must perform this operation, contact your authorized Oracle service provider.

Caution – When you create a RAID volume using the on-board disk controllers, all data currently on the disks you include in the volume will be lost.

Disk Zones for SPARC T3-1 Servers With Sixteen-Disk Backplanes

Whenever a SPARC T3-1 sixteen-disk backplane is managed by the onboard SAS-2 controllers, the backplane *must* be partitioned into two logical zones, with eight disk slots per zone:

- Zone A -- Contains backplane slots 0 through 7, which are visible only to onboard SAS-2 controller 0.
- Zone B -- Contains backplane slots 8 through 15, which are visible only to onboard SAS-2 controller 1.

Disk zoning requires that the LSI expander firmware on the backplane include patch 147034-01 (at least). This patch creates the required disk zones.

Note – For additional information about patch 147034-01, refer to its README document, *README.147034* on the My Oracle Support site.

Tip – When zoning is enabled, devalias for slots 8-15 will be incorrect unless the system firmware is updated to either 8.0.5.b (or a higher 8.0 level) or to 8.1.0 (or higher). For more information see "Minimum System Firmware for Valid devalias in Upper Disk Locations" on page 28 for more information.

These disk zone definitions are persistent. Their configuration cannot be modified and they remain in effect across power cycles and reboot operations. They must be active whenever the SAS-2 onboard controllers are used to manage disks in a sixteen-disk capable backplane.

However, if you use an internal PCIe RAID HBA instead of the onboard SAS-2 controllers, you must disable disk zoning. You do this with the following command:

```
# zoningcli disable zoning
```

The zoningcli command is contained in versions 2.1.1 and 2.2 of the Oracle Hardware Management Pack. See "Hardware Management Pack Overview" on page 5 for more information.

Most SPARC T3-1 servers with sixteen-disk capable backplanes are shipped with zoning enabled. There are two exceptions to this default rule:

- Zoning is disabled at the factory when a SPARC T3-1 system containing a sixteen-disk backplane is built with an internal PCIe RAID HBA.
- SPARC T3-1 systems manufactured with 16-disk backplanes before disk zoning became the default were shipped without patch 147034-01. For these systems, the patch must be installed in the field to support disk zoning.

Note – Because disks are not visible to controllers across zone boundaries, a SAS-2 controller cannot create a RAID volume that contains disks that are included in the other zone.

You can use the `zoningcli` utility to enable or disable zoning in a sixteen-disk capable backplane. See "Enabling and Disabling Disk Zoning In the Field" on page 27 for more information.

Displaying Disk Zone Information

These topics describe two methods for accessing disk zoning information.

- "Disk Zone Status (`zoningcli list`)" on page 23
- "Disk Zone Status (`probe-scsi-all`)" on page 24

Disk Zone Status (`zoningcli list`)

If Hardware Management Pack software version 2.1.1 (or later) is installed on your system, you can use the `zoningcli list` command to determine whether or not zoning is enabled. Zoning status is shown in the first line of the output in each of the following examples.

Zoning is disabled:

```
# zoningcli list expander
Expander: SUN     SAS2 X16DBP     zoning: disable    <====
=================================================
PHY   SAS ADDRESS     ZG   ZG Persist  att-dev att-id
=================================================
00 5000c50017b0c149 010     1          01     00
01 5000c5000576a349 010     1          01     00
[...]
25 0000000000000000 001     1          00     00
=================================================
```

Zoning is enabled:

```
# zoningcli list expander
Expander: SUN     SAS2 X16DBP     zoning: enable    <====
=================================================
PHY   SAS ADDRESS     ZG   ZG Persist  att-dev att-id
=================================================
00 5000c50017b0c149 010     1          01     00
```

```
01 5000c5000576a349 010        1             01      00
[...]
25 0000000000000000 001        1             00      00
=================================================
```

If the `zoningcli` utility is not available on your system, you can use the OBP command `probe-scsi-all` to determine whether or not zoning is enabled. See "Disk Zone Status (`probe-scsi-all`)" on page 24.

Disk Zone Status (`probe-scsi-all`)

The following `probe-scsi-all` output examples show the difference in how storage devices are listed when zoning is enabled and when it is disabled. There are 10 disks represented in both examples. They are identified as PhyNum 0-9.

Note – The entries for PhyNum 1-6 are omitted in these examples to reduce the size of the examples. The missing entries would not alter the illustration of the concept.

Zoning is disabled:

When zoning is disabled, each disk is visible to both controllers. The following example shows this by listing every disk with controller 0 and with controller 1.

Tip – Note that SASDeviceName and SASAddress for each PhyNum listed under controller 1 matches the corresponding PhyNum's DeviceName and SASAddress listed under controller 0. This is true for PhyNum 1-6, which were omitted to reduce the size of the table.

```
{0} ok probe-scsi-all
/pci@400/pci@2/pci@0/pci@f/pci@0/usb@0,2/hub@2/hub@3/storage@2
  Unit 0    Removable Read Only device    AMI      Virtual CDROM    1.00

/pci@400/pci@2/pci@0/pci@4/scsi@0                               <==== Controller 1

FCode Version 1.00.56, MPT Version 2.00, Firmware Version 5.00.13.00

Target a
  Unit 0    Removable Read Only device    TEAC     DV-W28SS-R      1.0C
  SATA device   PhyNum 6
Target b
  Unit 0    Disk    SEAGATE   ST930003SSUN300G 0868    585937500 Blocks, 300 GB
  SASDeviceName 5000c50017b0c14b   SASAddress 5000c50017b0c149   PhyNum 0
```

```
         [PhyNum 1-6 are omitted to conserve space.]

Target 12
  Unit 0    Disk    SEAGATE   ST973402SSUN72G  0400    143374738 Blocks, 73 GB
  SASDeviceName 5000c50003d5c6a3  SASAddress 5000c50003d5c6a1   PhyNum 7
Target 13
  Unit 0    Disk    SEAGATE   ST914603SSUN146G 0768    286739329 Blocks, 146 GB
  SASDeviceName 5000c50012ef2247  SASAddress 5000c50012ef2245   PhyNum 8
Target 14
  Unit 0    Disk    SEAGATE   ST973402SSUN72G  0400    143374738 Blocks, 73 GB
  SASDeviceName 5000c50003d49c77  SASAddress 5000c50003d49c75   PhyNum 9
Target 15
  Unit 0    Encl Serv device    SUN      SAS2 X16DBP      0302
  SASAddress 500605b0000272bd  PhyNum 18

/pci@400/pci@1/pci@0/pci@4/scsi@0                          <==== Controller 0

FCode Version 1.00.56, MPT Version 2.00, Firmware Version 5.00.13.00

Target a
  Unit 0    Disk    SEAGATE   ST930003SSUN300G 0868    585937500 Blocks, 300 GB
  SASDeviceName 5000c50017b0c14b  SASAddress 5000c50017b0c149   PhyNum 0

          [PhyNum 1-6 are omitted to conserve space.]

Target 11
  Unit 0    Disk    SEAGATE   ST973402SSUN72G  0400    143374738 Blocks, 73 GB
  SASDeviceName 5000c50003d5c6a3  SASAddress 5000c50003d5c6a1   PhyNum 7
Target 12
  Unit 0    Disk    SEAGATE   ST914603SSUN146G 0768    286739329 Blocks, 146 GB
  SASDeviceName 5000c50012ef2247  SASAddress 5000c50012ef2245   PhyNum 8
Target 13
  Unit 0    Disk    SEAGATE   ST973402SSUN72G  0400    143374738 Blocks, 73 GB
  SASDeviceName 5000c50003d49c77  SASAddress 5000c50003d49c75   PhyNum 9
Target 14
  Unit 0    Encl Serv device    SUN      SAS2 X16DBP      0302
  SASAddress 500605b0000272bd  PhyNum 18

{0} ok
```

Zoning is enabled:

In the following example, zoning is enabled and each disk is shown connected to a single SAS-2 controller. PhyNum 0-7 are connected to controller 0 and PhyNum 8-9 are connected to controller 1.

```
{0} ok probe-scsi-all
/pci@400/pci@2/pci@0/pci@f/pci@0/usb@0,2/hub@2/hub@3/storage@2
  Unit 0   Removable Read Only device    AMI     Virtual CDROM   1.00

/pci@400/pci@2/pci@0/pci@4/scsi@0                        <==== Controller 1

FCode Version 1.00.56, MPT Version 2.00, Firmware Version 5.00.13.00

Target a
  Unit 0   Disk   SEAGATE  ST914603SSUN146G 0768    286739329 Blocks, 146 GB
  SASDeviceName 5000c50012ef2247  SASAddress 5000c50012ef2245  PhyNum 8
Target b
  Unit 0   Disk   SEAGATE  ST973402SSUN72G  0400    143374738 Blocks, 73 GB
  SASDeviceName 5000c50003d49c77  SASAddress 5000c50003d49c75  PhyNum 9
Target c
  Unit 0   Encl Serv device   SUN      SAS2 X16DBP       0305
  SASAddress 50800200000272bd  PhyNum 18
Target d
  Unit 0   Removable Read Only device   TEAC    DV-W28SS-R      1.0C
  SATA device   PhyNum 6

/pci@400/pci@1/pci@0/pci@4/scsi@0                        <==== Controller 0

Target a
  Unit 0   Disk   SEAGATE  ST930003SSUN300G 0868    585937500 Blocks, 300 GB
  SASDeviceName 5000c50017b0c14b  SASAddress 5000c50017b0c149  PhyNum 0

         [PhyNum 1-6 are omitted to conserve space.]

Target 11
  Unit 0   Disk   SEAGATE  ST973402SSUN72G  0400    143374738 Blocks, 73 GB
  SASDeviceName 5000c50003d5c6a3  SASAddress 5000c50003d5c6a1  PhyNum 7
Target 12
  Unit 0   Encl Serv device   SUN      SAS2 X16DBP       0305
  SASAddress 50800200000272bd  PhyNum 18

{0} ok
```

Enabling and Disabling Disk Zoning In the Field

These topics describe how to enable and disable disk zones in the field.

Note – You must have version 2.1.1 of the Oracle Hardware Management Pack (or later) to have access to the `zoningcli` command.

- "Enable Disk Zones (`zoningcli` Command)" on page 27
- "Disable Disk Zones (`zoningcli` Command)" on page 27
- "Update Firmware to Enable Zoning" on page 27

▼ Enable Disk Zones (`zoningcli` Command)

- **If zoning has been disabled, you can enable it by running the following command:**

```
# zoningcli enable zoning
```

▼ Disable Disk Zones (`zoningcli` Command)

- **If zoning has been enabled, you can disable it by running the following command:**

```
# zoningcli disable zoning
```

▼ Update Firmware to Enable Zoning

If your SPARC T3-1 server requires disk zoning, but does not have the required minimum system firmware level and/or the LSI firmware on the backplane does not have patch 147034-01, you must perform the following step for the controllers to function properly.

Caution – Be certain to back up any data stored on the disks before you install this patch. You can restore the files after the patch is installed.

● **Apply patch 147034-01 to the LSI firmware on the disk backplane.**

This patch causes the backplane to be partitioned into the two eight-disk zones described in "Disk Zones for SPARC T3-1 Servers With Sixteen-Disk Backplanes" on page 21.

Tip – When zoning is enabled, devalias for slots 8-15 will be incorrect unless the system firmware is updated to either 8.0.5.b (or a higher 8.0 level) or to 8.1.0 (or higher). See "Minimum System Firmware for Valid devalias in Upper Disk Locations" on page 28 for more information.

Minimum System Firmware for Valid devalias in Upper Disk Locations

For devalias to be valid for disk slots 8-15, the system firmware level must be at 8.0.5.b (or a higher 8.0 level) or at 8.1.0.c (or higher). If your server's system firmware does not meet the minimum requirement, you must use the full disk path to identify individual disks in the range 8-15.

For example, if running with the minimum system firmware level and the correct devalias, you can use the following command line to boot from the disk in slot 12:

```
# boot disk12
```

If your disk slot is in the range 8-15 and the system firmware level does not meet the requirement described above, you will need to specify the full device path for the boot disk. This example shows the path for disk 12:

```
# boot /pci@400/pci@2/pci@0/pci@4/scsi@0/disk@p10c
```

▼ Prepare to Use the FCode Utility

1. **Open an** xterm **or** gnome **terminal window.**

FCode commands produce large amounts of detailed output. The xterm or gnome terminal windows provide scroll bar functionality, which helps view such output.

2. **Disable auto-boot in OBP and enter the OBP environment after a power on or reset.**

3. **Use the** `show-devs` **command to list the device paths on the server.**

```
{0} ok show-devs
...
/pci@400/pci@1/pci@0/pci@4/scsi@0
/pci@400/pci@2/pci@0/pci@4/scsi@0
...
```

Note – For a server module, the device path might be /pci@400/pci@1/pci@0/pci@2/LSI,sas@0.

---

4. **Use the** `select` **command to choose the controller where you want to create a hardware RAID volume.**

```
{0} ok select /pci@400/pci@2/pci@0/pci@4/scsi@0
```

Instead of using the entire device path for the controller, you can use a preconfigured alias for a controller. For example:

```
{0} ok select scsi0
```

To view the preconfigured aliases on your server, use the `devalias` command. See "Locate Device Paths" on page 35.

Display the SAS addresses of any connected drives using the `show-children` command.

Related Information

- "FCode Utility Commands" on page 30
- "Display the ok Prompt" on page 11

FCode Utility Commands

FCode Command	Description
show-children	Lists all connected physical drives and logical volumes.
show-volumes	Lists all connected logical volumes in detail.
create-raid0-volume	Creates a RAID 0 volume (minimum two targets).
create-raid1-volume	Creates a RAID 1 volume (exactly two targets).
create-raid1e-volume	Creates a RAID 1e volume (minimum three targets).
delete-volume	Deletes a RAID volume.
activate-volume	Re-activate a RAID volume after the motherboard has been replaced.

Related Information

- "Create a Hardware RAID Volume" on page 30
- "Prepare to Use the FCode Utility" on page 28

▼ Create a Hardware RAID Volume

1. **Prepare to create a RAID volume.**

   See "Prepare to Use the FCode Utility" on page 28.

2. **Use the** show-children **command to list the physical drives on the selected controller.**

```
{0} ok show-children

FCode Version 1.00.54, MPT Version 2.00, Firmware Version 5.00.17.00

Target 9
  Unit 0   Disk    SEAGATE  ST930003SSUN300G 0B70    585937500 Blocks, 300 GB
  SASDeviceName 5000c5001771776f  SASAddress 5000c5001771776d  PhyNum 0
Target a
  Unit 0   Disk    SEAGATE  ST930003SSUN300G 0B70    585937500 Blocks, 300 GB
```

```
  SASDeviceName 5000c5001d0c38c7  SASAddress 5000c5001d0c38c5  PhyNum 1
Target b
  Unit 0   Disk   SEAGATE  ST930003SSUN300G 0B70    585937500 Blocks, 300 GB
  SASDeviceName 5000c5001d097407  SASAddress 5000c5001d097405  PhyNum 2
Target c
  Unit 0   Disk   SEAGATE  ST930003SSUN300G 0B70    585937500 Blocks, 300 GB
  SASDeviceName 5000c5001d09a51f  SASAddress 5000c5001d09a51d  PhyNum 3
{0} ok
```

3. **Use the** `create-raid0-volume`, `create-raid1-volume`, **or**
   `create-raid1e-volume` **command to create a logical drive from the physical
   disks.**

   For example, to create a RAID 0 volume with targets 9 and a, type the targets first
   then type the `create` command:

   ```
   {0} ok 9 a create-raid0-volume
   ```

   For example, to create a RAID 1e volume with targets a, b, and c, type:

   ```
   {0} ok a b c create-raid1e-volume
   ```

4. **To verify creation of the volume, type:**

   ```
   {0} ok show-volumes
   ```

5. **Type** `unselect-dev` **to deselect the controller.**

   ```
   {0} ok unselect-dev
   ```

Related Information

- "FCode Utility Commands" on page 30
- "Display the ok Prompt" on page 11

Hot Spare Drives in RAID Volumes (LSI)

You can configure two global hot spare drives to protect data on mirrored RAID
volumes. If one of the drives in a RAID 1 or RAID 1E mirrored volume fails, the
onboard RAID controller will replace the failed drive automatically with a hot spare
drive and then resynchronize the mirrored data.

Use the `sas2ircu` LSI utility to add global hot spare drives. Refer to the *SAS2 Integrated RAID Solution User Guide* for more information about adding hot spare drives.

Determining If a Drive Has Failed

These topics explain various ways to determine if a drive contained in a RAID volume has failed:

■ "Front Service Required Drive LEDs" on page 32
■ "Error Messages (System Console and Log Files)" on page 33
■ "Display Status (`show-volumes` Command, OBP)" on page 34
■ "Display Status (`sas2ircu` Utility, LSI)" on page 34

Front Service Required Drive LEDs

When a fault occurs on a drive in a SPARC T3 server, the amber Service Required LED will light on the front of the drive. This amber LED will allow you to locate the faulted drive in the system. Additionally, the Service Action Required LEDs on the front and rear panels also light when the system detects a hard drive fault. Refer to your service manual for the location and description of these LEDs.

Error Messages (System Console and Log Files)

When a fault occurs on a drive, error messages will be displayed on the system console. This is an example of a system console display indicating that volume 905 has been degraded with the loss of PhysDiskNum 1:

```
Mar 16 16:28:26 hostname scsi: /pci@400/pci@2/pci@0/pci@e/scsi@0 (mpt_sas0):
Mar 16 16:28:26 hostname      PhysDiskNum 1 with DevHandle 0xc in slot 0 for
enclosure with handle 0x0 is now offline
Mar 16 16:28:26 hostname scsi: /pci@400/pci@2/pci@0/pci@e/scsi@0 (mpt_sas0):
Mar 16 16:28:26 hostname      PhysDiskNum 1 with DevHandle 0xc in slot 0 for
enclosure with handle 0x0 is now , active, out of sync
Mar 16 16:28:26 hostname scsi: WARNING: /pci@400/pci@2/pci@0/pci@e/scsi@0
(mpt_sas0):
Mar 16 16:28:26 hostname      Volume 905 is degraded
Mar 16 16:28:26 hostname scsi: /pci@400/pci@2/pci@0/pci@e/scsi@0 (mpt_sas0):
```

```
Mar 16 16:28:26 hostname      Volume 0 is now degraded
Mar 16 16:28:26 hostname scsi: WARNING: /pci@400/pci@2/pci@0/pci@e/scsi@0
(mpt_sas0):
Mar 16 16:28:26 hostname      Volume 905 is degraded
```

You can also view these messages by reviewing the /var/adm/messages files:

```
# more /var/adm/messages*
. . .
Mar 16 16:45:19 hostname SC Alert: [ID 295026 daemon.notice] Sensor | minor:
Entity Presence : /SYS/SASBP/HDD3/PRSNT : Device Absent
Mar 16 16:45:19 hostname scsi: [ID 107833 kern.notice]
/pci@400/pci@2/pci@0/pci@e/scsi@0 (mpt_sas0):
Mar 16 16:45:19 hostname PhysDiskNum 0 with DevHandle 0xd in slot 0 for enclosure
with handle 0x0 is now offline
Mar 16 16:45:19 hostname scsi: [ID 107833 kern.notice]
/pci@400/pci@2/pci@0/pci@e/scsi@0 (mpt_sas0):
Mar 16 16:45:19 hostname PhysDiskNum 0 with DevHandle 0xd in slot 0 for enclosure
with handle 0x0 is now , active, out of sync
Mar 16 16:45:19 hostname scsi: [ID 107833 kern.warning] WARNING:
/pci@400/pci@2/pci@0/pci@e/scsi@0 (mpt_sas0):
Mar 16 16:45:19 hostname Volume 905 is degraded
Mar 16 16:45:19 hostname scsi: [ID 107833 kern.notice]
/pci@400/pci@2/pci@0/pci@e/scsi@0 (mpt_sas0):
Mar 16 16:45:19 hostname Volume 0 is now degraded
Mar 16 16:45:19 hostname scsi: [ID 107833 kern.warning] WARNING:
/pci@400/pci@2/pci@0/pci@e/scsi@0 (mpt_sas0):
Mar 16 16:45:19 hostname Volume 905 is degraded
```

Refer to the *View the System Message Log Files* topic in your service manual for more information about examining these messages.

▼ Display Status (show-volumes Command, OBP)

You can halt the system and use the show-volumes OBP command to determine if a drive has failed.

1. **Halt the system and display the OBP ok prompt.**

   See "Display the ok Prompt" on page 11.

2. **Select the SAS controller device.**

```
ok select /pci@400/pci@2/pci@0/pci@e/scsi@0
```

See "Prepare to Use the FCode Utility" on page 28 for more information.

3. **Type the** `show-volumes` **command to display the RAID volumes and their associated disks.**

In the following example, the secondary disk in a RAID 1 volume is offline.

```
ok show-volumes
Volume 0 Target 389   Type RAID1 (Mirroring)
  Name raid1test   WWID 04eec3557b137f31
  Degraded    Enabled
  2 Members                                              2048 Blocks, 1048 KB
  Disk 1
    Primary   Optimal
    Target c      HITACHI    H101414SCSUN146G SA25
  Disk 0
    Secondary   Offline  Out Of Sync
    Target 0        SEAGATE
```

4. **Type the** `unselect-dev` **command to deselect the SAS controller device.**

```
ok unselect-dev
```

▼ Display Status (`sas2ircu` Utility, LSI)

● **Use the LSI** `sas2ircu` **utility to display the status of the RAID volume and its associated drives.**

Refer to the *SAS2 Integrated RAID Solution User Guide* for more information about displaying and interpreting device status using the `sas2ircu` utility.

RAID Drive Replacement Strategies

Follow the strategies described below when replacing a failed drive in a RAID volume.

RAID Volume Level	Strategy
RAID 0	If a drive fails in a RAID 0 volume, you will lose all data on that volume. Replace the failed drive with a new drive of the same capacity, recreate the RAID 0 volume, and restore the data from a backup.

RAID Volume Level	Strategy
RAID 1	Remove the failed drive and replace it with a new drive of the same capacity. The new drive will be automatically configured and synced with the RAID volume.
RAID 1E	Remove the failed drive and replace it with a new drive of the same capacity. The new drive will be automatically configured and synced with the RAID volume.

Note – The `cfgadm` instructions in the service manual are for individual drives that are not part of RAID volumes. When a drive is part of a RAID volume, you do not need to unconfigure it before hot swapping it with a new drive.

▼ Locate Device Paths

Use this procedure to locate device paths specific to your server.

1. **Display the** ok **prompt.**

   See "Display the ok Prompt" on page 11.

2. **From the** ok **prompt, type:**

```
{0} ok devalias
screen                 /pci@400/pci@2/pci@0/pci@0/pci@0/display@0
mouse                  /pci@400/pci@2/pci@0/pci@f/pci@0/usb@0,2/hub@2/device@4
/mouse@1
rcdrom                  /pci@400/pci@2/pci@0/pci@f/pci@0/usb@0,2/hub@2/hub@3
/storage@2/disk@0
rkeyboard               /pci@400/pci@2/pci@0/pci@f/pci@0/usb@0,2/hub@2/device@4
/keyboard@0
rscreen                /pci@400/pci@2/pci@0/pci@0/pci@0/display@0:r1280x1024x60
net3                    /pci@400/pci@2/pci@0/pci@7/network@0,1
net2                    /pci@400/pci@2/pci@0/pci@7/network@0
net1                    /pci@400/pci@2/pci@0/pci@6/network@0,1
net0                    /pci@400/pci@2/pci@0/pci@6/network@0
net                     /pci@400/pci@2/pci@0/pci@6/network@0
disk7                   /pci@400/pci@2/pci@0/pci@4/scsi@0/disk@p3
disk6                   /pci@400/pci@2/pci@0/pci@4/scsi@0/disk@p2
disk5                   /pci@400/pci@2/pci@0/pci@4/scsi@0/disk@p1
disk4                   /pci@400/pci@2/pci@0/pci@4/scsi@0/disk@p0
cdrom                   /pci@400/pci@2/pci@0/pci@4/scsi@0/disk@p6
```

```
scsi1                  /pci@400/pci@2/pci@0/pci@4/scsi@0
disk3                  /pci@400/pci@1/pci@0/pci@4/scsi@0/disk@p3
disk2                  /pci@400/pci@1/pci@0/pci@4/scsi@0/disk@p2
disk1                  /pci@400/pci@1/pci@0/pci@4/scsi@0/disk@p1
disk0                  /pci@400/pci@1/pci@0/pci@4/scsi@0/disk@p0
disk                   /pci@400/pci@1/pci@0/pci@4/scsi@0/disk@p0
scsi0                  /pci@400/pci@1/pci@0/pci@4/scsi@0
scsi                   /pci@400/pci@1/pci@0/pci@4/scsi@0
virtual-console        /virtual-devices@100/console@1
name                   aliases
{0} ok
```

Related Information

■ "Display the ok Prompt" on page 11

■ "Prepare to Use the FCode Utility" on page 28

Changing Server Identification Information

These topics describe how to store information (for purposes such as inventory control or site resource management) on the SP and FRU PROMs using the Oracle ILOM CLI interface.

- "Change Customer FRUdata Information" on page 37
- "Change System Identifier Information" on page 38

▼ Change Customer FRUdata Information

Use the /SP customer_frudata property to store information in all FRU PROMs. This field can be used to identify a particular system to a third-party application, or for any other identifcation needs in your environment.

- **At the Oracle ILOM -> prompt, type:**

```
-> set /SP customer_frudata="data"
```

Note – You must enclose the data string (*data*) in quote marks.

Related Information

- "Log In to Oracle ILOM" on page 9
- "Change System Identifier Information" on page 38

▼ Change System Identifier Information

Use the /SP system_identifier property to store customer identification information. This string is encoded in all trap messages generated by SNMP. Assigning a unique system identifier can be useful in distinguishing which system generates which SNMP message.

● **At the Oracle ILOM -> prompt, type:**

```
-> set /SP system_identifier="data"
```

Note – The data string (*data*) must be enclosed in quote marks.

Related Information

■ "Log In to Oracle ILOM" on page 9
■ "Change Customer FRUdata Information" on page 37

Configuring Policy Settings

These topics describe managing configuration policies using Oracle ILOM.

- "Specify Cooldown Mode" on page 39
- "Restore Host Power State at Restart" on page 40
- "Specify Host Power State at Restart" on page 41
- "Disable or Re-Enable Power-On Delay" on page 41
- "Specify Parallel Boot of the SP and Host" on page 42
- "Configure Host Behavior With the Keyswitch State" on page 42

▼ Specify Cooldown Mode

The HOST_COOLDOWN property, when set to enabled, will cause the server to enter cooldown mode when the host is powering off. Upon server poweroff, cooldown mode directs Oracle ILOM to monitor certain components to ensure that they are below a minimum temperature so as not to cause harm to the user. Once the components are below the minimum temperature, then the power is removed from the server, or if it takes longer then 4 minutes the host will turn off.

Note – The HOST_COOLDOWN property does not apply to server modules.

- **At the Oracle ILOM -> prompt, type:**

```
-> set /SP/policy HOST_COOLDOWN=value
```

where *value* can be:

- enabled – Server cools certain components before powering off.
- disabled – Component temperatures are not monitored during power off.

Related Information

■ "Power Off the Server" on page 16

▼ Restore Host Power State at Restart

Use the /SP/policy HOST_LAST_POWER_STATE property to control the behavior of the server after an unexpected power outage. When external power is restored, the Oracle ILOM service processor starts to run automatically. Normally, the host power is not turned on until you use Oracle ILOM to turn it on.

Oracle ILOM records the current power state of the server in nonvolatile storage. If the HOST_LAST_POWER_STATE policy is enabled, Oracle ILOM can restore the host to the previous power state. This policy is useful in the event of a power failure, or if you physically move the server to a different location.

For example, if the host server is running when power is lost and the /SP/policy HOST_LAST_POWER_STATE property is set to disabled, the host server remains off when power is restored. If the /SP/policy HOST_LAST_POWER_STATE property is set to enabled, the host server restarts when the power is restored.

● **At the Oracle ILOM -> prompt, type:**

```
-> set /SP/policy HOST_LAST_POWER_STATE=enabled
```

where *value* can be:

■ enabled – When power is restored, returns the server to the state it was in before the power was removed.

■ disabled – Keeps the server off when power is applied (the default).

If you enable HOST_LAST_POWER_STATE, you should also configure /SP/policy HOST_POWER_ON_DELAY. For further information, see "Disable or Re-Enable Power-On Delay" on page 41.

Related Information

■ "Disable or Re-Enable Power-On Delay" on page 41

■ "Specify Host Power State at Restart" on page 41

▼ Specify Host Power State at Restart

Use `/SP/policy HOST_AUTO_POWER_ON` to power on the host when external power is applied to the server. If this policy is set to enabled, the service processor sets `HOST_LAST_POWER_STATE` to disabled.

● **At the Oracle ILOM `->` prompt, type:**

```
-> set /SP/policy HOST_AUTO_POWER_ON=value
```

where *value* can be:

- `enabled` – When power is applied, automatically powers on the host when the SP has been booted.
- `disabled` – Keeps the host power off when power is applied (the default).

Related Information

- "Restore Host Power State at Restart" on page 40
- "Disable or Re-Enable Power-On Delay" on page 41

▼ Disable or Re-Enable Power-On Delay

Use the `/SP/policy HOST_POWER_ON_DELAY` property to cause the server to wait for a short time before powering on automatically. The delay is a random interval of one to five seconds. Delaying the server power on helps minimize current surges on the main power source. This power-on delay is important when multiple servers in racks power on after a power outage.

● **At the Oracle ILOM `->` prompt, type:**

```
-> set /SP/policy HOST_POWER_ON_DELAY=value
```

where *value* can be:

- `enabled` – Causes the server to wait for a short time before powering on automatically.
- `disabled` – Server powers on automatically without a delay (the default).

Related Information

- "Specify Host Power State at Restart" on page 41

▼ Specify Parallel Boot of the SP and Host

The /SP/policy PARALLEL_BOOT property, when enabled, allows the host to boot and power on in parallel with the SP if an auto-power policy (HOST_AUTO_POWER_ON or HOST_LAST_POWER_STATE) was on or a user presses the power button while the SP is in the process of booting. Oracle ILOM has to be running in order to allow the host to power on when the power button is pressed or the the auto-power policies are set. When this property is set to disabled, the SP boots first, then the host boots.

Note – Parallel boot is not supported on server modules.

● **At the Oracle ILOM -> prompt, type:**

```
-> set /SP/policy PARALLEL_BOOT=value
```

where *value* can be:

- enabled – SP and host boot concurrently.
- disabled – SP and host boot serially.

Related Information

- "Power On the Server" on page 15
- "Power Off the Server" on page 16

▼ Configure Host Behavior With the Keyswitch State

Use the /SYS keyswitch_state property to control the position of the virtual keyswitch.

● **At the Oracle ILOM** –> **prompt, type:**

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

where *value* can be:

- `normal` – The server can power itself on and start the boot process (the default).
- `standby` – Powers off host, disables power on.
- `diag` – Host power on is allowed, it overrides the settings of `/HOST/diag` *target*, resulting in Max POST being executed.
- `locked` – Host power on is allowed, however, you are prohibited from updating any of the flash devices or setting `/HOST send_break_action= break`.

Related Information

Configuring Network Addresses

These topics describe managing network addresses with Oracle ILOM.

SP Network Address Options

You can access the SP on your system in multiple ways. Consider the following options and choose the access method that is best for your environment.

You can physically connect to the SP using a serial connection or a network connection. The network connection can be configured to use a static IP address or DHCP (the default). Optionally, the T3 Series servers can use an in-band network connection to the SP, rather than the default out-of-band network management port.

For more information on each option, see the following documentation:

- To use a serial connection to the SP, see:

  Connect a Terminal or Emulator to the SER MGT Port, in the installation guide for your server or *Communicating With the Server Module During Startup* in the installation guide for your server module.

- To assign a static IP address to the SP, see:

  Assign a Static IP Address to the SP, in the installation guide for your server.

- To use an in-band connection to the SP, see:

  "Oracle ILOM In-band (Sideband) Management" on page 48

Related Information

- Oracle Integrated Lights Out Manager (ILOM) 3.0 Documentation
- "Oracle ILOM Overview" on page 1

▼ Disable or Re-Enable Network Access to the SP

Use the /SP/network state property to enable or disable the service processor's network interface.

● **At the Oracle ILOM** -> **prompt, type:**

```
-> set /SP/network state=value
```

where *value* can be:

- enabled (the default)
- disabled

Related Information

- "SP Network Address Options" on page 45

▼ Display the DHCP Server IP Address

To display the IP address of the DHCP server that provided the dynamic IP address requested by the service processor, view the dhcp_server_ip property.

● **At the Oracle ILOM** -> **prompt, type:**

```
-> show /SP/network

  /SP/network
    Targets:
        interconnect
        ipv6
        test
```

```
      Properties:
          commitpending = (Cannot show property)
          dhcp_server_ip = 10.8.31.5
          ipaddress = 10.8.31.188
          ipdiscovery = dhcp
          ipgateway = 10.8.31.248
          ipnetmask = 255.255.252.0
          macaddress = 00:14:4F:7E:83:4F
          managementport = /SYS/MB/SP/NETMGMT
          outofbandmacaddress = 00:21:28:8F:13:5F
          pendingipaddress = 10.8.31.188
          pendingipdiscovery = dhcp
          pendingipgateway = 10.8.31.248
          pendingipnetmask = 255.255.252.0
          sidebandmacaddress = 00:21:28:8F:13:5E
          state = enabled

      Commands:
          cd
          set
          show
```

Note – The list of properties might vary based on your server.

Related Information

- "Display the Host MAC Address" on page 47

▼ Display the Host MAC Address

The /HOST macaddress property is automatically configured by the server software, so you cannot set or change the property. The value is read and determined from the server's removable system configuration card (SCC PROM) or from the server module's ID PROM and then stored as a property in Oracle ILOM.

/HOST macaddress is the MAC address for the net0 port. The MAC addresses for each additional port increments from the /HOST macaddress. For example, net1 is equal to the value of /HOST macaddress plus one (1).

● **At the Oracle ILOM -> prompt, type:**

```
-> show /HOST macaddress
```

Using an In-band Connection to the SP

These topics describe how to use an in-band, or sideband connection, to the SP.

Oracle ILOM In-band (Sideband) Management

By default, you connect to the server's SP using the out-of-band network management port (NET MGT). The Oracle ILOM sideband management feature enables you to select either the NET MGT port or one of the server's Gigabit Ethernet ports (NET*n*), which are in-band ports, to send and receive Oracle ILOM commands to and from the server SP. In-band ports are also called sideband ports.

The advantage of using a sideband management port to manage the server's SP is that one fewer cable connection and one fewer network switch port are needed. In configurations where numerous servers are being managed, such as data centers, sideband management can represent a significant savings in hardware and network utilization.

Note – In-band connection is not recommended for server modules.

When sideband management is enabled in Oracle ILOM, the following conditions might occur:

■ Connectivity to the server SP might be lost when the SP management port configuration is changed while you are connected to the SP using a network connection, such as SSH, web, or Oracle ILOM Remote Console.

■ In-chip connectivity between the SP and the host operating system might not be supported by the on-board host Gigabit Ethernet controller. If this condition occurs, use a different port or route to transmit traffic between the source and destination targets instead of using L2 bridging/switching.

■ Server host power cycles might cause a brief interruption of network connectivity for server Gigabit Ethernet ports (NET*n*) that are configured for sideband management. If this condition occurs, configure the adjacent switch/bridge ports as host ports.

Related Information

- "Configure SP In-band (Sideband) Access" on page 49
- "SP Network Address Options" on page 45

▼ Configure SP In-band (Sideband) Access

This procedure describes how to access the SP from an in-band (or sideband) management using a host network port.

If you perform this procedure using a network connection, you might lose connectivity to the server. Using a serial connection for this procedure eliminates the possibility of losing connectivity during sideband management configuration changes.

1. Log in to Oracle ILOM.

See "Log In to Oracle ILOM" on page 9.

2. If you logged in using the serial port, you can assign a static IP address.

For instructions, see the information about assigning an IP address in the installation guide for your server.

3. View the current network settings:

```
-> show /SP/network

/SP/network
    Targets:
        interconnect
        ipv6
        test

    Properties:
        commitpentding = (Cannot show property)
        dhcp_server_ip = none
        ipaddress = 129.148.62.55
        ipdiscovery = static
        ipgateway = 129.148.62.225
        ipnetmask = 255.255.255.0
        macaddress = 00:11:3T:8O:44:B7
        managementport= /SYS/MB/SP/NETMGMT
        outofbandmacaddress = 00:11:3T:8O:44:B7
        pendingipaddress = 129.148.62.55
        pendingipdiscovery = static
        pendingipgateway = 129.148.62.225
        pendingipnetmask = 255.255.255.0
```

```
         pendingmanagementport = /SYS/MB/SP/NETMGMT
         sidebandmacaddress = 00:11:3T:8O:44:B7
         state = enabled

    Commands:
         cd
         set
         show
```

4. Set the SP management port to a sideband port (where *n* is 0-3):

```
-> set /SP/network pendingmanagementport=/SYS/MB/NETn

-> set commitpending=true
```

5. Verify the change:

```
-> show /SP/network

/SP/network
    Targets:
         interconnect
         ipv6
         test

    Properties:
         commitpentding = (Cannot show property)
         dhcp_server_ip = none
         ipaddress = 129.148.62.55
         ipdiscovery = static
         ipgateway = 129.148.62.225
         ipnetmask = 255.255.255.0
         macaddress = 00:11:3T:8O:44:B7
         managementport= /SYS/MB/SP/NET0
         outofbandmacaddress = 00:11:3T:8O:44:B7
         pendingipaddress = 129.148.62.55
         pendingipdiscovery = static
         pendingipgateway = 129.148.62.225
         pendingipnetmask = 255.255.255.0
         pendingmanagementport = /SYS/MB/SP/NET0
         sidebandmacaddress = 00:11:3T:8O:44:B7
         state = enabled

    Commands:
```

```
cd
set
show
```

Related Information

- "Oracle ILOM In-band (Sideband) Management" on page 48
- "SP Network Address Options" on page 45

Configuring Boot Mode

Use the Oracle ILOM boot mode properties to specify how the host boots when correcting a problem with OpenBoot or Oracle VM Server for SPARC settings.

Boot Mode Overview

Boot mode (`bootmode`) properties enable you to override the default method the server uses when it boots. This ability is useful to override particular OpenBoot or Oracle VM Server settings that might be incorrect, to set up OpenBoot variables using a script, or to perform similar tasks.

For example, if the OpenBoot settings have become corrupt, you can set the `bootmode` state property to `reset_nvram` then reset the server to its factory default OpenBoot settings.

Service personnel might instruct you to use the `bootmode` script property for problem resolution. The full extent of script capabilities are not documented and exist primarily for debugging.

Because `bootmode` is intended to be used to correct a problem with the OpenBoot or Oracle VM Server settings, the `bootmode` takes effect for a single boot only. Additionally, to prevent an administrator from setting a `bootmode` state property and forgetting about it, a `bootmode` state property expires if the host is not reset within 10 minutes of the `bootmode` state property being set.

Related Information

- "Reset the Server From the Oracle Solaris OS" on page 17
- "Reset the Server From Oracle ILOM" on page 17
- "OpenBoot Overview" on page 3
- "Oracle VM Server for SPARC Overview" on page 4

▼ Configure the Host Boot Mode of Oracle VM Server for SPARC

Note – You must use a valid Oracle VM Server configuration name for this task.

1. Determine the valid Oracle VM Server configurations on your SP, at the Oracle ILOM -> prompt, type:

```
-> show /HOST/domain/configs
```

2. Set the boot mode configuration, at the Oracle ILOM -> prompt, type:

```
-> set /HOST/bootmode config=configname
```

where the `config` property takes a *configname* value that is a valid named logical domain configuration.

For example, if you created a Oracle VM Server configuration called `ldm-set1`:

```
-> set bootmode config=ldm-set1
```

To return the boot mode `config` to the factory default configuration, specify `factory-default`.

For example:

```
-> set bootmode config=factory-default
```

Related Information

- "Reset the Server From Oracle ILOM" on page 17

▼ Change the Host Boot Mode Behavior at Reset

The `/HOST/bootmode` state property controls how OpenBoot NVRAM variables are used. Normally the current settings of these variables are retained. Setting `/HOST/bootmode state=reset_nvram` changes the OpenBoot NVRAM variables to their default settings at the next reset.

● **At the Oracle ILOM** `->` **prompt, type:**

```
-> set /HOST/bootmode state=value
```

where *value* is one of the following:

- `normal` – At next reset, retains current NVRAM variable settings.
- `reset_nvram` – At next reset, returns OpenBoot variables to default settings.

Note – state=reset_nvram will return to normal after the next server reset or 10 minutes (see `expires` property in "Display Host Boot Mode Expiration Date" on page 57). `config` and `script` properties do not expire and will be cleared upon the next server reset or manually by setting *value* to "".

Related Information

▼ Manage the Host Boot Mode Script

● **At the Oracle ILOM -> prompt, type:**

```
-> set /HOST/bootmode script=value
```

where `script` controls the host server OpenBoot PROM firmware method of booting.

`script` does not affect the current `/HOST/bootmode` setting.

value can be up to 64 bytes in length.

You can specify a `/HOST/bootmode` setting and set the script within the same command. For example:

```
-> set /HOST/bootmode state=reset_nvram script="setenv diag-switch? true"
```

After the server resets and OpenBoot PROM reads the values stored in the script, the OpenBoot PROM sets the OpenBoot PROM variable `diag-switch?` to the user-requested value of `true`.

Note – If you set `/HOST/bootmode  script=""`, Oracle ILOM sets the `script` to empty.

Related Information

- "Reset the Server From Oracle ILOM" on page 17
- "Boot Mode Overview" on page 53
- "Oracle VM Server for SPARC Overview" on page 4

▼ Display Host Boot Mode Expiration Date

● **At the Oracle ILOM** -> **prompt, type:**

```
-> show /HOST/bootmode expires
     Properties:
        expires = Thu Oct 14 18:24:16 2010
```

where `expires` is the date and time when the current boot mode will expire.

Related Information

■ "Reset the Server From Oracle ILOM" on page 17

■ "Boot Mode Overview" on page 53

■ "Oracle VM Server for SPARC Overview" on page 4

▼ Override OpenBoot PROM Settings to Reset the Server

Use this procedure to override the OpenBoot PROM settings and initiate reboot of the control domain, which results in the host booting to the `ok` prompt.

● **At the Oracle ILOM** -> **prompt, type:**

```
-> set /HOST/domain/control auto-boot=disabled
   reset /HOST/domain/control [-force] [-script]
```

The host reboots and stops at the `ok` prompt.

Related Information

■ "Change the Host Boot Mode Behavior at Reset" on page 55

■ "Configuring Boot Mode" on page 53

■ "Boot Mode Overview" on page 53

■ "Oracle VM Server for SPARC Overview" on page 4

Configuring Server Behavior at Restart

Use the following procedures to configure how Oracle ILOM should behave during the following restart scenarios.

- "Specify Behavior When the Host Resets" on page 59
- "Specify Behavior When the Host Stops Running" on page 60
- "Set the Boot Timeout Interval" on page 60
- "Specify Behavior at Boot Timeout" on page 61
- "Specify Behavior if Restart Fails" on page 61
- "Specify Maximum Restart Attempts" on page 62

▼ Specify Behavior When the Host Resets

Specify if the host should continue to boot if an error is encountered.

● **Set this property:**

```
-> set /HOST autorunonerror=value
```

where *value* can be:

- `false` – The host continues to boot if an error is encountered.
- `true` – The host *does not* continue to boot if an error is encountered.

Related Information

- "Configuring Policy Settings" on page 39

▼ Specify Behavior When the Host Stops Running

Specify what Oracle ILOM should do when the host leaves the RUNNING state (when the watchdog timer expires).

● **Set this property:**

```
-> set /HOST autorestart=value
```

where *value* can be:

- none – Oracle ILOM takes no action other than to issue a warning.
- reset – Oracle ILOM attempts to reset the server when the Oracle Solaris watchdog timer expires (the default).
- dumpcore – Oracle ILOM attempts to force a core dump of the OS when the watchdog timer expires.

Related Information

-

▼ Set the Boot Timeout Interval

● **Set the time delay between a request to boot the host and booting the host:**

```
-> set /HOST boottimeout=seconds
```

The default value of boottimeout is 0 (zero seconds) or no timeout. Possible values are in the range from 0 to 36000 seconds.

Related Information

-

▼ Specify Behavior at Boot Timeout

Specify what Oracle ILOM should do if the host fails to boot before the boot timeout interval.

● **Specify behavior at the completion of** `boottimeout`:

```
-> set /HOST bootrestart=value
```

where *value* can be:
- `none` (the default)
- `reset`

Related Information

▼ Specify Behavior if Restart Fails

Specify what Oracle ILOM should do if the host fails to reach the Oracle Solaris `running` state.

● **At the Oracle ILOM** –> **prompt, type:**

```
-> set /HOST bootfailrecovery=value
```

where *value* can be:
- `powercycle`
- `poweroff` (the default)

Related Information

▼ Specify Maximum Restart Attempts

Specify how many times Oracle ILOM should attempt to restart the host.

● **At the Oracle ILOM -> prompt, type:**

```
-> set /HOST maxbootfail=attempts
```

The default value of maxbootfail is 3 (three attempts).

If the host does not boot successfully within the number of tries indicated by maxbootfail, the host is powered off or powercycled (depending upon the setting of bootfailrecovery). In either case, boottimeout is set to 0 (zero seconds), disabling further attempts to restart the host.

Related Information

■ "Specify Behavior if Restart Fails" on page 61

Configuring Devices

These topics contain information about configuring devices in the server.

▼ Unconfigure a Device Manually

The Oracle ILOM firmware provides a component_state=disabled command, which enables you to unconfigure server 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 server's machine description prior to transfer of control to other layers of system firmware, such as OpenBoot PROM.

● **At the Oracle ILOM** -> **prompt, type:**

```
-> set component-name component_state=disabled
```

Related Information

▼ Reconfigure a Device Manually

The Oracle ILOM firmware provides a component_state=enabled command, which enables you to reconfigure server devices manually. Use this command to mark the specified device as enabled.

● **At the Oracle ILOM -> prompt, type:**

```
-> set component-name component_state=enabled
```

Related Information

- "Log In to Oracle ILOM" on page 9
- "Unconfigure a Device Manually" on page 63
- "Display Server Components" on page 75
- "Locate Device Paths" on page 35

Monitoring the Server

The server provides many ways to indicate faulty behavior, including LEDs, Oracle ILOM, and POST. For specific information about LEDs, and for complete troubleshooting information, refer to the service manual for your server.

Monitoring Faults

These topics contain a summary of diagnostic tools and basic information about finding server faults using pre-OS tools, including Oracle ILOM and POST. For complete troubleshooting information, see the service manual for your server.

Diagnostics Overview

You can use a variety of diagnostic tools, commands, and indicators to monitor and troubleshoot a server. See the service manual for your server for complete information about these diagnostic tools:

- **LEDs** – Provide a quick visual notification of the status of the server and of some of the FRUs.

- **Oracle ILOM** – This firmware runs on the service processor. In addition to providing the interface between the hardware and OS, Oracle ILOM also tracks and reports the health of key server components. Oracle ILOM works closely with POST and Oracle Solaris Predictive Self-Healing technology to keep the server running even when there is a faulty component.

- **Power-on self-test** – POST performs diagnostics on server components upon server reset to ensure the integrity of those components. POST is configureable and works with Oracle ILOM to take faulty components offline if needed.

- **Oracle Solaris OS Predictive Self-Healing** – This technology continuously monitors the health of the CPU, memory, and other components, and works with Oracle ILOM to take a faulty component offline if needed. The PSH technology enables servers to accurately predict component failures and mitigate many serious problems before they occur.

- **Log files and command interface** – Provide the standard Oracle Solaris OS log files and investigative commands that can be accessed and displayed on the device of your choice.

- **SunVTS –** An application that exercises the server, provides hardware validation, and discloses possible faulty components with recommendations for repair.

The LEDs, Oracle ILOM, PSH, and many of the log files and console messages are integrated. For example, when the Oracle Solaris software detects a fault, it displays the fault, and logs it, and passes information to Oracle ILOM, where it is logged.

Related Information

- "Discover Faults (Oracle ILOM)" on page 67
- "Discover Faults Using POST" on page 69
- "Discover Faults (Oracle ILOM Fault Management Shell)" on page 67
- See the *server service manual section detecting and managing faults*

▼ Discover Faults (Oracle ILOM)

● **At the Oracle ILOM** -> **prompt, type:**

```
-> show faulty
```

This command displays the target, the property, and the value of the fault.

For example:

```
-> show faulty
Target            | Property           | Value
------------------+--------------------+-------------------------
/SP/faultmgmt/0   | fru                | /SYS
/SP/faultmgmt/1   | fru                | /SYS/MB/CMP0/BOBO/CH1/D0
/SP/faultmgmt/1/  | fru_part_number    | 18JS25672PDZ1G1F1
 faults/0         |                    |

->
```

Related Information

▼ Discover Faults (Oracle ILOM Fault Management Shell)

The Oracle ILOM Fault Management shell provides a way to use the Oracle Solaris Fault Manager commands (fmadm, fmstat) from within Oracle ILOM, and to view both host and Oracle ILOM faults.

1. To start the captive shell, at the Oracle ILOM -> **prompt, type:**

```
-> start /SP/faultmgmt/shell
Are you sure you want to start /SP/Faultmgt/shell (y/n)? y

faultmgmtsp>
```

2. For a list of current server faults, type:

```
faultmgmtsp> fmadm faulty
---------------- ------------------------------------ -------------- --------
Time                  UUID                                 msgid       Severity
---------------- ------------------------------------ -------------- --------
2010-09-03/20:46:23 fa4a2f86-5156-4243-8e88-d6516db12970 SPT-8000-DH   Critical

Fault class : fault.chassis.voltage.fail

FRU         : /SYS/MB
              (Part Number: 541-4197-04)
              (Serial Number: 1005LCB-1025D40059)

Description : A chassis voltage supply is operating outside of the
              allowable range.

Response    : The system will be powered off.  The chassis-wide service
              required LED will be illuminated.

Impact      : The system is not usable until repaired.  ILOM will not allow
              the system to be powered on until repaired.

Action      : The administrator should review the ILOM event log for
              additional information pertaining to this diagnosis.  Please
              refer to the Details section of the Knowledge Article for
              additional information.
```

Note – If the server detects the replacement of the faulty FRU, the repair does not require a user command, and the fault will be cleared automatically.

3. Discover more information about a specific fault.

Find the fault MSG-ID (SPT-8000-42 in the preceeding example), and type it in the Search Knowledge Base search window at http://support.oracle.com.

4. To repair the fault, see:

"Repair a Fault (Oracle ILOM Fault Management Shell)" on page 71.

5. To leave the Fault Management shell and return to Oracle ILOM, type:

```
faultmgmtsp> exit
->
```

Related Information

- "Oracle Solaris 10 OS Feature Spotlight: Predictive Self Healing" at
  `www.oracle.com/technetwork/systems/dtrace/self-healing/index.h`
  `tml`
- "Log In to Oracle ILOM" on page 9
- "Discover Faults (Oracle ILOM)" on page 67
- "Repair a Fault (Oracle ILOM Fault Management Shell)" on page 71

▼ 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 server reset.

1. Log in to Oracle ILOM.

See "Log In to Oracle ILOM" on page 9.

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

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

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

3. To return to your normal diagnostic settings *after* running POST, at the Oracle ILOM -> prompt, type:

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

Related Information

- "Discover Faults (Oracle ILOM)" on page 67
- "Locate the Server" on page 76
- "Clear a Fault" on page 72

▼ Display Console History

This topic describes displaying the host server console output buffers.

There are two console history buffers that can contain up to 1 Mbyte of information. The `/HOST/console/history` target writes all types of log information. The `/HOST/console/bootlog` target writes boot information and initialization data into the console buffer until Oracle ILOM is notified by the server that the Oracle Solaris OS is up and running. This buffer is kept until the host is booted again.

Note – You must have Oracle ILOM Administrator level user permission to use this command.

---

1. **To manage the** `/HOST/console/history` **log, at the Oracle ILOM** `->` **prompt, type:**

```
-> set /HOST/console/history property=option [...]
-> show /HOST/console/history
```

where *property* can be:

- `line_count` – This option accepts a value within the range of 1 to 2048 lines. Specify "" for an unlimited number of lines. The default is all lines.
- `pause_count` – This option accepts a value of 1 to any valid integer or "" for infinite number of lines. The default is not to pause.
- `start_from` – The options are:
  - `end` – The last line (most recent) in the buffer (the default).
  - `beginning` – The first line in the buffer.

    If you type the show `/HOST/console/history` command without first setting any arguments with the `set` command, Oracle ILOM displays all lines of the console log, starting from the end.

Note – Timestamps recorded in the console log reflect server time. These timestamps reflect local time, and the Oracle ILOM console log uses UTC (Coordinated Universal Time). The Oracle Solaris OS system time is independent of the Oracle ILOM time.

---

2. **To view the** `/HOST/console/bootlog`, **at the Oracle ILOM** `->` **prompt, type:**

```
-> show /HOST/console/bootlog property
```

where *property* can be:

- `line_count` – This option accepts a value within the range of 0 to 2048 lines. Specify "0" for an unlimited number of lines. The default is all lines.
- `pause_count` – This option accepts a value of within the range of 0 to 2048 lines. Specify "0" for an unlimited number of lines. The default is not to pause.

- start_from – The options are:
  - end – The last line (most recent) in the buffer (the default).
  - beginning – The first line in the buffer.

Note – Timestamps recorded in the console log reflect server time. These timestamps reflect local time, and the Oracle ILOM console log uses UTC (Coordinated Universal Time). The Oracle Solaris OS system time is independent of the Oracle ILOM time.

Related Information

- "Specify Host Power State at Restart" on page 41

▼ Repair a Fault (Oracle ILOM Fault Management Shell)

You can use the fmadm repair command to fix faults diagnosed by Oracle ILOM. (Faults diagnosed by Oracle ILOM, rather than the host, have message IDs starting with "SPT.")

The only time you should use the fmadm repair command in the Oracle ILOM Fault Management shell for a host-diagnosed fault is when the fault is repaired and Oracle ILOM is unaware of the repair. For example, Oracle ILOM might be down when the fault is repaired. In that case, the host would no longer display the fault, yet the fault is still displayed in Oracle ILOM. Use the the fmadm repair command to clear the fault.

1. Locate the fault:

```
faultmgmtsp> fmadm faulty
----------------- ------------------------------------- -------------- --------
Time                UUID                                      msgid        Severity
----------------- ------------------------------------- -------------- --------
2010-09-03/20:46:23 fa4a2f86-5156-4243-8e88-d6516db12970 SPT-8000-DH   Critical

Fault class : fault.chassis.voltage.fail

FRU          : /SYS/MB
               (Part Number: 541-4197-04)
               (Serial Number: 1005LCB-1025D40059)

Description : A chassis voltage supply is operating outside of the
              allowable range.

Response    : The system will be powered off.  The chassis-wide service
```

```
              required LED will be illuminated.

Impact     : The system is not usable until repaired.  ILOM will not allow
             the system to be powered on until repaired.

Action     : The administrator should review the ILOM event log for
             additional information pertaining to this diagnosis.  Please
             refer to the Details section of the Knowledge Article for
             additional information.

faultmgmtsp> fmadm repair fa4a2f86-5156-4243-8e88-d6516db12970
faultmgmtsp>
```

2. To repair an Oracle ILOM-detected fault, use the fmadm repair **command:**

```
faultmgmtsp> fmadm repair fa4a2f86-5156-4243-8e88-d6516db12970
faultmgmtsp>
```

Note – You can use either the NAC name (for instance, /SYS/MB) or the UUID (for instance, fa4a2f86-5156-4243-8e88-d6516db12970) of the fault with the fmadm repair command.

3. To leave the Fault Management shell and return to Oracle ILOM, type:

```
faultmgmtsp> exit
->
```

Related Information

■ "Discover Faults (Oracle ILOM Fault Management Shell)" on page 67

▼ Clear a Fault

● **At the Oracle ILOM** -> **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

■ "Discover Faults (Oracle ILOM)" on page 67

Enabling Automatic System Recovery

These topics include information about configuring your server to automatically recover from minor faults.

Note – This section refers to the automatic system recovery feature, not the similarly named auto service request feature.

Automatic System Recovery Overview

The server provides for ASR from failures in memory modules or PCI cards.

ASR functionality enables the server 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 server operation. As long as the server is capable of operating without the failed component, the ASR features enable the server to reboot automatically, without operator intervention.

Note – ASR is not activated until you enable it. See "Enable ASR" on page 74.

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

Related Information

▼ Enable ASR

1. At the Oracle ILOM `->` 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 server permanently stores the parameter changes and boots automatically when the OpenBoot configuration variable `auto-boot?` is set to `true` (its default value).

Related Information

- "Automatic System Recovery Overview" on page 73
- "Disable ASR" on page 74
- "View Information on Components Affected by ASR" on page 75

▼ 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 server permanently stores the parameter change.

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

Related Information

- "Enable ASR" on page 74
- "View Information on Components Affected by ASR" on page 75
- "Automatic System Recovery Overview" on page 73

▼ View Information on Components Affected by ASR

- **At the Oracle ILOM -> 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 Overview" on page 73
- "Enable ASR" on page 74
- "Disable ASR" on page 74
- "Unconfigure a Device Manually" on page 63
- "Reconfigure a Device Manually" on page 63

▼ Display Server Components

View real-time information about the components installed in your server using the Oracle ILOM `show components` command.

● **At the Oracle ILOM -> prompt, type:**

```
-> show components
Target               | Property                | Value
---------------------+-------------------------+-------------------
/SYS/MB/RISER0/      | component_state         | Enabled
PCIE0                |                         |
/SYS/MB/RISER0/      | component_state         | Disabled
PCIE3                |                         |
/SYS/MB/RISER1/      | component_state         | Enabled
PCIE1                |                         |
/SYS/MB/RISER1/      | component_state         | Enabled
/SYS/MB/NET0         | component_state         | Enabled
/SYS/MB/NET1         | component_state         | Enabled
/SYS/MB/NET2         | component_state         | Enabled
```

Note – Components will vary based on your server.

Related Information

■ "Locate Device Paths" on page 35

▼ Locate the Server

In case you need to service a component, lighting the system locator LED assists in easily identifying the correct server. You do not need administrator permissions to use the set /SYS/LOCATE and show /SYS/LOCATE commands.

1. Log in to Oracle ILOM.

   See "Log In to Oracle ILOM" on page 9.

2. Manage the Locator LED with the following commands.

   ■ To turn on the Locator LED, from the Oracle ILOM service processor command prompt, type:

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

   ■ To turn off the Locator LED, from the Oracle ILOM service processor command prompt, type:

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

■ To display the state of the Locator LED, from the Oracle ILOM service
  processor command prompt, type:

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

Related Information

- "Monitoring Faults" on page 65
- "Configuring Devices" on page 63

Updating the Firmware

These topics describe how to update the system firmware and view current versions of firmware for Oracle's SPARC T3 series servers.

-
-
-
-

▼ Display the Firmware Version

The `/HOST sysfw_version` property displays information about the system firmware version on the host.

- **View the current setting for this property. At the Oracle ILOM -> prompt, type:**

```
-> show /HOST sysfw_version
```

Related Information

-

▼ Update the Firmware

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

   See the installation guide for your server 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

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:**

```
-> stop /SYS
```

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

```
-> 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.xxx/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://129.99.99.99/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 server automatically resets, runs diagnostics, and returns to the login prompt on the serial console.

```
U-Boot 1.x.x

Custom AST2100 U-Boot 3.0 (Aug 21 2010 - 10:46:54) r58174
***
Net:   faradaynic#0, faradaynic#1
Enter Diagnostics Mode
['q'uick/'n'ormal(default)/e'x'tended(manufacturing mode)] .....
0
Diagnostics Mode - NORMAL
<DIAGS> Memory Data Bus Test ... PASSED
<DIAGS> Memory Address Bus Test ... PASSED
I2C Probe Test - SP
     Bus      Device                          Address Result
     ===      ===========================     ======= ======
      6                  SP FRUID (U1101)      0xA0    PASSED
      6                 DS1338(RTC) (U1102)    0xD0    PASSED

<DIAGS> PHY #0 R/W Test ... PASSED
<DIAGS> PHY #0 Link Status ... PASSED
<DIAGS> ETHERNET PHY #0, Internal Loopback Test ... PASSED
## Booting image at 110a2000 ... ***

Mounting local filesystems...
Mounted all disk partitions.

Configuring network interfaces...FTGMAC100: eth0:ftgmac100_open
```

```
Starting system log daemon: syslogd and klogd.
Starting capidirect daemon: capidirectd . Done
Starting Event Manager: eventmgr  . Done
Starting ipmi log manager daemon: logmgr . Done
Starting IPMI Stack: . Done
Starting sshd.
Starting SP fishwrap cache daemon: fishwrapd . Done
Starting Host deamon: hostd . Done
Starting Network Controller Sideband Interface Daemon: ncsid . Done
Starting Platform Obfuscation Daemon: pod . Done
Starting lu main daemon: lumain . Done
Starting Detection/Diagnosis After System Boot: dasboot  Done
Starting Servicetags discoverer: stdiscoverer.
Starting Servicetags listener: stlistener.
Starting Dynamic FRUID Daemon: dynafrud  Done

hostname login:
```

Related Information

■ "Display the Firmware Version" on page 79

▼ Display OpenBoot Version

The /HOST obp_version property displays information about the version of OpenBoot on the host.

● **View the current setting for this property:**

```
-> show /HOST obp_version
```

Related Information

■ "Update the Firmware" on page 79
■ "OpenBoot Overview" on page 3

▼ Display POST Version

The /HOST post_version property displays information about the version of POST on the host.

● **View the current setting for this property:**

```
-> show /HOST post_version
```

Related Information

- "Update the Firmware" on page 79

Identifying WWN-Designated SAS2 Devices

These topics describe how to update the Oracle Solaris OS on Oracle's SPARC T3 series servers.

World Wide Name Syntax

The Oracle Solaris OS now uses the World Wide Name (WWN) syntax in place of the locally unique t*n* (target ID) field in logical device names. This change affects how a target storage device is identified when downloading the operating system over a network. The following points are key to understanding the impact of this change:

- Before the change to WWN nomenclature, the Oracle Solaris OS would ordinarily identify the default boot device as `c0t0d0`.
- With the change, the device identifier for the default boot device is now referred to as `c0t`*WWN*`d0`, where *WWN* is a hexadecimal value that is unique to this device throughout the world.

- This WWN value does not map in a predictable way to the physical address of the device to which it refers.

To reliably specify a particular storage device for an OS download operation, you must determine the correspondence between the WWN value assigned to that device and its physical location.

You can find this correspondence using either OBP or Oracle Solaris commands:

- When in OBP, run `probe-scsi-all`. See "Identify a Disk Slot Using `probe-scsi-all` (OBP)" on page 96 for instructions.
- When in Oracle Solaris, run `format`, followed by `prtconf -v`. See "Identify a Disk Slot Using `prtconf` (Oracle Solaris, Onboard Controllers)" on page 98 or "Identify a Disk Slot Using `prtconf` (Oracle Solaris, Single Initiator)" on page 100 for instructions.

These commands generate information about the SAS controllers and the storage devices that are connected to them. This information includes both logical and physical names, which you can analyze to determine these logical and physical address relationships.

Note – Refer to your service manual for a description of the physical drive slot organization on your system.

Related Information

`probe-scsi-all` Output Example (SPARC T3-1, Eight-Disk Backplane)

When analyzing `probe-scsi-all` output, look for this SAS device information:

- `SASDeviceName` -- This is the WWN that the Oracle Solaris OS recognizes.
- `SASAddress` -- This is the WWN that the OBP recognizes.
- `PhyNum` -- This is a hexadecimal value that identifies the port that a disk is connected to.
- `VolumeDeviceName` -- When a RAID volume has been configured, this is the RAID volume WWN value recognized by the Oracle Solaris OS.
- `VolumeWWID` -- When a RAID volume has been configured, this is the RAID volume WWN value referenced by OBP.

The following table shows the PhyNum-to-disk slot correspondence for an eight-disk backplane configuration.

TABLE: SAS Controller Port Mapping for Eight-Disk Backplanes

SAS Controller	PhyNum	Disk Slot	SAS Controller	PhyNum	Disk Slot
0	0	0	1	0	4
	1	1		1	5
	2	2		2	6
	3	3		3	7

The following example illustrates `probe-scsi-all` output for a SPARC T3-1 server with seven hard drives in an eight-disk capable backplane configuration. In this example, the hard drives are connected to the two SAS controllers in this manner:

- Four hard drives are connected to SAS controller 0. These are targets 9, a, b, and c.
- Three hard drives and a SATA DVD device are connected to SAS controller 1. These are targets 9, b, c, and a, respectively.

Because SAS controller 0 controls the hard drives installed in backplane slots 0-3, the default boot device for this sample configuration is PhyNum 0 in the group of hard drives under controller 0. It has a SASDeviceName value of 5000cca00a75dcac and a SASAddress value of 5000cca00a75dcad.

```
ok probe-scsi-all
/pci@400/pci@2/pci@0/pci@f/pci@0/usb@0,2/hub@2/hub@3/storage@2
Unit 0    Removable Read Only device    AMI      Virtual CDROM    1.00

/pci@400/pci@2/pci@0/pci@4/scsi@0                          <---- SAS controller 1

FCode Version 1.00.54, MPT Version 2.00, Firmware Version 5.00.17.00

Target 9
  Unit 0   Disk   SEAGATE  ST930003SSUN300G 0B70    585937500 Blocks, 300 GB
  SASDeviceName 5000c5001d33fba7  SASAddress 5000c5001d33fba5  PhyNum 0
Target a
  Unit 0   Removable Read Only device   TEAC    DV-W28SS-R     1.0C
  SATA device   PhyNum 6
Target b
  Unit 0   Disk   HITACHI  H103030SCSUN300G A2A8    585937500 Blocks, 300 GB
  SASDeviceName 5000cca00a76e380  SASAddress 5000cca00a76e381  PhyNum 1
Target c
  Unit 0   Disk   SEAGATE  ST930003SSUN300G 0B70    585937500 Blocks, 300 GB
  SASDeviceName 5000cca00a76ddcc  SASAddress 5000cca00a76ddcd  PhyNum 3

/pci@400/pci@1/pci@0/pci@4/scsi@0                          <---- SAS controller 0
FCode Version 1.00.54, MPT Version 2.00, Firmware Version 5.00.17.00

Target 9
  Unit 0   Disk   HITACHI  H103030SCSUN300G A2A8    585937500 Blocks, 300 GB
  SASDeviceName 5000cca00a75dcac  SASAddress 5000cca00a75dcad  PhyNum 0
Target a
  Unit 0   Disk   HITACHI  H103030SCSUN300G A2A8    585937500 Blocks, 300 GB
  SASDeviceName 5000cca00a7680d4  SASAddress 5000cca00a7680d5  PhyNum 2
Target b
  Unit 0   Disk   SEAGATE  ST930003SSUN300G 0B70    585937500 Blocks, 300 GB
  SASDeviceName 5000c5001d33eb5f  SASAddress 5000c5001d33eb5d  PhyNum 3
Target c
  Unit 0   Disk   SEAGATE  ST930003SSUN300G 0B70    585937500 Blocks, 300 GB
  SASDeviceName 5000c5001d40bf9b  SASAddress 5000c5001d40bf99  PhyNum 1
```

Related Information

- "World Wide Name Syntax" on page 85
- "probe-scsi-all Output Example (SPARC T3-1, Sixteen-Disk Backplane)" on page 89

probe-scsi-all Output Example (SPARC T3-1, Sixteen-Disk Backplane)

When analyzing probe-scsi-all output, look for this SAS device information:

- SASDeviceName -- This is the WWN that the Oracle Solaris OS recognizes.
- SASAddress -- This is the WWN that the OBP recognizes.
- PhyNum -- This is a hexadecimal value that identifies the port that a disk is connected to.
- VolumeDeviceName -- When a RAID volume has been configured, this is the RAID volume WWN value recognized by the Oracle Solaris OS.
- VolumeWWID -- When a RAID volume has been configured, this is the RAID volume WWN value referenced by OBP.

The following table shows the PhyNum-to-disk slot correspondence for a sixteen-disk backplane connected to onboard SAS controllers 0 and 1.

Note – The default 16-disk backplane configuration is divided into two logically isolated zones, with disk slots 0-7 visible to SAS controller 0 and disk slots 8-15 visible to SAS controller 1.

TABLE: SAS Controller Port Mapping for Sixteen-Disk Backplanes

SAS Controller	PhyNum (Hex)	Disk Slot (Dec)	SAS Controller	PhyNum (Hex)	Disk Slot (Dec)
0	0	0	1	8	8
	1	1		9	9
	2	2		A	10

TABLE: SAS Controller Port Mapping for Sixteen-Disk Backplanes *(Continued)*

SAS Controller	PhyNum (Hex)	Disk Slot (Dec)	SAS Controller	PhyNum (Hex)	Disk Slot (Dec)
	3	3	B		11
	4	4	C		12
	5	5	D		13
	6	6	E		14
	7	7	F		15

The following example illustrates `probe-scsi-all` output for a SPARC T3-1 server with nine hard drives in a sixteen-disk capable backplane configuration.

Note – Controller 1 manages the SATA DVD.

The default boot location (physical slot 0) is listed as `PhyNum 0` in the SAS controller 0 section. It has a `SASDeviceName` value of `5000cca00a59278c` and a `SASAddress` value of `5000cca00a59278d`.

```
ok probe-scsi-all
/pci@400/pci@2/pci@0/pci@f/pci@0/usb@0,2/hub@2/hub@3/storage@2
   Unit 0    Removable Read Only device     AMI     Virtual CDROM   1.00

/pci@400/pci@2/pci@0/pci@4/scsi@0           <---- SAS controller 1

FCode Version 1.00.54, MPT Version 2.00, Firmware Version 5.00.17.00

Target a
  Unit 0   Disk   HITACHI  H103030SCSUN300G A2A8    585937500 Blocks, 300 GB
  SASDeviceName 5000cca00a746490  SASAddress 5000cca00a746491  PhyNum 8
Target b
  Unit 0   Disk    SEAGATE  ST930003SSUN300G 0B70    585937500 Blocks, 300 GB
  SASDeviceName 5000c5001d0d1283  SASAddress 5000c5001d0d1281  PhyNum 9
Target c
  Unit 0   Disk   HITACHI  H103030SCSUN300G A2A8    585937500 Blocks, 300 GB
  SASDeviceName 5000cca00a01014c  SASAddress 5000cca00a01014d  PhyNum a
Target d
  Unit 0   Disk    SEAGATE  ST930003SSUN300G 0B70    585937500 Blocks, 300 GB
  SASDeviceName 5000c5001d0ca947  SASAddress 5000c5001d0ca945  PhyNum b
Target e
  Unit 0   Disk    SEAGATE  ST930003SSUN300G 0B70    585937500 Blocks, 300 GB
  SASDeviceName 5000c5001d0d2857  SASAddress 5000c5001d0d2855  PhyNum c
Target f
  Unit 0   Disk   HITACHI  H103030SCSUN300G A2A8    585937500 Blocks, 300 GB
  SASDeviceName 5000cca00a02f5d0  SASAddress 5000cca00a02f5d1  PhyNum d
```

```
Target 10
  Unit 0   Disk   SEAGATE  ST930003SSUN300G 0B70    585937500 Blocks, 300 GB
  SASDeviceName 5000c5001d0c3d9b  SASAddress 5000c5001d0c3d99  PhyNum e
Target 11
  Unit 0   Disk   SEAGATE  ST930003SSUN300G 0B70    585937500 Blocks, 300 GB
  SASDeviceName 5000c5001d097737  SASAddress 5000c5001d097735  PhyNum f
Target 12
  Unit 0   Encl Serv device   SUN      SAS2 X16DBP     0305
  SASAddress 50800200000272bd  PhyNum 18


/pci@400/pci@1/pci@0/pci@4/scsi@0          <---- SAS controller 0

FCode Version 1.00.54, MPT Version 2.00, Firmware Version 5.00.17.00

Target a
  Unit 0   Disk   HITACHI  H103030SCSUN300G A2A8    585937500 Blocks, 300 GB
  SASDeviceName 5000cca00a59278c  SASAddress 5000cca00a59278d  PhyNum 0
Target b
  Unit 0   Disk   SEAGATE  ST930003SSUN300G 0768    585937500 Blocks, 300 GB
  SASDeviceName 5000c50017c7e6fb  SASAddress 5000c50017c7e6f9  PhyNum 1
Target c
  Unit 0   Disk   HITACHI  H103030SCSUN300G A2A8    585937500 Blocks, 300 GB
  SASDeviceName 5000cca00a4bfcb8  SASAddress 5000cca00a4bfcb9  PhyNum 2
Target d
  Unit 0   Disk   SEAGATE  ST930003SSUN300G 0768    585937500 Blocks, 300 GB
  SASDeviceName 5000c50017c7a3a3  SASAddress 5000c50017c7a3a1  PhyNum 3
Target e
  Unit 0   Disk   HITACHI  H103030SCSUN300G A2A8     585937500 Blocks, 300 GB
  SASDeviceName 5000cca00a02cc18  SASAddress 5000cca00a02cc19  PhyNum 4
Target f
  Unit 0   Disk   HITACHI  H103030SCSUN300G A2A8     585937500 Blocks, 300 GB
  SASDeviceName 5000cca00a76dec0  SASAddress 5000cca00a76dec1  PhyNum 5
Target 10
  Unit 0   Disk   HITACHI  H103030SCSUN300G A2A8     585937500 Blocks, 300 GB
  SASDeviceName 5000cca00a773eac  SASAddress 5000cca00a773ead  PhyNum 6
Target 11
  Unit 0   Disk   SEAGATE  ST930003SSUN300G 0768    585937500 Blocks, 300 GB
  SASDeviceName 5000c5001d09654f  SASAddress 5000c5001d09654d  PhyNum d
Target 12
  Unit 0   Encl Serv device   SUN      SAS2 X16DBP     0305
  SASAddress 50800200000272bd  PhyNum 18


ok
```

Related Information

probe-scsi-all Output Example (SPARC T3-4)

When analyzing probe-scsi-all output, look for this SAS device information:

- SASDeviceName -- This is the WWN that the Oracle Solaris OS recognizes.
- SASAddress -- This is the WWN that the OBP recognizes.
- PhyNum -- This is a hexadecimal value that identifies the port that a disk is connected to.
- VolumeDeviceName -- When a RAID volume has been configured, this is the RAID volume WWN value recognized by the Oracle Solaris OS.
- VolumeWWID -- When a RAID volume has been configured, this is the RAID volume WWN value referenced by OBP.

A SPARC T3-4 server has two onboard SAS controllers, each connected to a separate four-disk capable backplane. The following table shows the PhyNum-to-disk slot correspondence for these backplanes.

TABLE: SAS Controller Port Mapping for an Eight-Disk Backplane

SAS Controller	PhyNum	Disk Slot	SAS Controller	PhyNum	Disk Slot
0	0	0	1	0	4
	1	1		1	5
	2	2		2	6
	3	3		3	7

Note – OBP uses a different device path for SAS controller 1 in SPARC T3-4 servers, depending on whether the server has four-processors or two processors. The path for SAS Controller 0 is the same for both processor configurations.

SAS Controller Device Path: SPARC T3-4 Server With Four Processors

The following example illustrates `probe-scsi-all` output for a SPARC T3-4 with four processors and eight drives.

Because SAS controller 0 controls the hard drives installed in backplane slots 0-3, the default boot device for this example is PhyNum 0 in the group of hard drives under controller 0. It has a `SASDeviceName` value of 5000cca00a75dcac and a `SASAddress` value of 5000cca00a75dcad.

```
ok probe-scsi-all
/pci@700/pci@1/pci@0/pci@0/LSI,sas@0                       <---- SAS controller 1

FCode Version 1.00.54, MPT Version 2.00, Firmware Version 5.00.17.00

Target 9
  Unit 0   Disk   HITACHI  H103030SCSUN300G A2A8    585937500 Blocks, 300 GB
  SASDeviceName 5000cca00abc5cc8  SASAddress 5000cca00abc5cc9  PhyNum 0
Target a
  Unit 0   Disk   HITACHI  H103030SCSUN300G A2A8    585937500 Blocks, 300 GB
  SASDeviceName 5000cca00abaf620  SASAddress 5000cca00abaf621  PhyNum 1
Target b
  Unit 0   Disk   HITACHI  H103030SCSUN300G A2A8    585937500 Blocks, 300 GB
  SASDeviceName 5000cca00abcec4c  SASAddress 5000cca00abcec4d  PhyNum 2
Target c
  Unit 0   Disk   HITACHI  H103030SCSUN300G A2A8    585937500 Blocks, 300 GB
  SASDeviceName 5000cca00abc5218  SASAddress 5000cca00abc5219  PhyNum 3
```

```
/pci@400/pci@1/pci@0/pci@8/pci@0/usb@0,2/hub@2/hub@3/storage@2
Unit 0    Removable Read Only device    AMI    Virtual CDROM   1.00

/pci@400/pci@1/pci@0/pci@0/LSI,sas@0                    <---- SAS controller 0

FCode Version 1.00.54, MPT Version 2.00, Firmware Version 5.00.17.00

Target 9
  Unit 0   Disk   HITACHI  H103030SCSUN300G A2A8    585937500 Blocks, 300 GB
  SASDeviceName 5000cca00abcede0  SASAddress 5000cca00abcede1   PhyNum 0
Target a
  Unit 0   Disk   HITACHI  H103030SCSUN300G A2A8    585937500 Blocks, 300 GB
  SASDeviceName 5000cca00abc51a8  SASAddress 5000cca00abc51a9   PhyNum 1
Target b
  Unit 0   Disk   HITACHI  H103030SCSUN300G A2A8    585937500 Blocks, 300 GB
  SASDeviceName 5000cca00abce89c  SASAddress 5000cca00abce89d   PhyNum 2
Target c
  Unit 0   Disk   HITACHI  H103030SCSUN300G A2A8    585937500 Blocks, 300 G
  SASDeviceName 5000cca00abc5354  SASAddress 5000cca00abc5355   PhyNum 3
```

SAS Controller Device Path: SPARC T3-4 Server With Two Processors

The following example illustrates `probe-scsi-all` output for a SPARC T3-4 with two processors and eight drives.

Because SAS controller 0 controls the hard drives installed in backplane slots 0-3, the default boot device for this example is PhyNum 0 in the group of hard drives under controller 0. It has a SASDeviceName value of 5000cca00a75dcac and a SASAddress value of 5000cca00a75dcad.

```
ok probe-scsi-all
/pci@500/pci@1/pci@0/pci@0/LSI,sas@0                    <---- SAS controller 1

FCode Version 1.00.54, MPT Version 2.00, Firmware Version 5.00.17.00

Target 9
  Unit 0   Disk   HITACHI  H103030SCSUN300G A2A8    585937500 Blocks, 300 GB
  SASDeviceName 5000cca00abc5cc8  SASAddress 5000cca00abc5cc9   PhyNum 0
Target a
  Unit 0   Disk   HITACHI  H103030SCSUN300G A2A8    585937500 Blocks, 300 GB
  SASDeviceName 5000cca00abaf620  SASAddress 5000cca00abaf621   PhyNum 1
Target b
  Unit 0   Disk   HITACHI  H103030SCSUN300G A2A8    585937500 Blocks, 300 GB
  SASDeviceName 5000cca00abcec4c  SASAddress 5000cca00abcec4d   PhyNum 2
Target c
```

```
   Unit 0   Disk   HITACHI  H103030SCSUN300G A2A8    585937500 Blocks, 300 GB
   SASDeviceName 5000cca00abc5218  SASAddress 5000cca00abc5219  PhyNum 3

/pci@400/pci@1/pci@0/pci@8/pci@0/usb@0,2/hub@2/hub@3/storage@2
Unit 0   Removable Read Only device    AMI     Virtual CDROM   1.00

/pci@400/pci@1/pci@0/pci@0/LSI,sas@0                    <---- SAS controller 0

FCode Version 1.00.54, MPT Version 2.00, Firmware Version 5.00.17.00

Target 9
   Unit 0   Disk   HITACHI  H103030SCSUN300G A2A8    585937500 Blocks, 300 GB
   SASDeviceName 5000cca00abcede0  SASAddress 5000cca00abcede1  PhyNum 0
Target a
   Unit 0   Disk   HITACHI  H103030SCSUN300G A2A8    585937500 Blocks, 300 GB
   SASDeviceName 5000cca00abc51a8  SASAddress 5000cca00abc51a9  PhyNum 1
Target b
   Unit 0   Disk   HITACHI  H103030SCSUN300G A2A8    585937500 Blocks, 300 GB
   SASDeviceName 5000cca00abce89c  SASAddress 5000cca00abce89d  PhyNum 2
Target c
   Unit 0   Disk   HITACHI  H103030SCSUN300G A2A8    585937500 Blocks, 300 G
   SASDeviceName 5000cca00abc5354  SASAddress 5000cca00abc5355  PhyNum 3
```

Related Information

- "World Wide Name Syntax" on page 85
- "probe-scsi-all Output Example (SPARC T3-1, Eight-Disk Backplane)" on page 87
- "probe-scsi-all Output Example (SPARC T3-1, Sixteen-Disk Backplane)" on page 89
- "Identify a Disk Slot Using probe-scsi-all (OBP)" on page 96
- "Identify a Disk Slot Using prtconf (Oracle Solaris, Onboard Controllers)" on page 98
- "Identify a Disk Slot Using prtconf (Oracle Solaris, Single Initiator)" on page 100
- "WWN Syntax in an OS Installation on a Specific Device" on page 103
- "WWN Syntax in an OS Installation on a RAID Volume" on page 105

▼ Identify a Disk Slot Using `probe-scsi-all` (OBP)

Use `probe-scsi-all` to display the following SAS device information:

- `SASDeviceName` -- This is the WWN that the Oracle Solaris OS recognizes.
- `SASAddress` -- This is the WWN that the OBP recognizes.
- `PhyNum` -- This is a hexadecimal value that identifies the port that a disk is connected to.
- `VolumeDeviceName` -- When a RAID volume has been configured, this is the RAID volume WWN value recognized by the Oracle Solaris OS.
- `VolumeWWID` -- When a RAID volume has been configured, this is the RAID volume's WWN value referenced by OBP.

The following example is based on a SPARC T3-2 server with six hard drives. Four hard drives are connected to the SAS controller as individual storage devices. These are `Target 9`, `d`, `e`, and `f`. Two hard drives are configured as a RAID volume. This disk array is `Target 389`.

- **Run** `probe-scsi-all`.

  In the following example, the hard drive installed in physical slot 0 has the `PhyNum` value 0. It is designated `Target 9` and has a `SASDeviceName` value of `5000c5001cb4a637`.

```
ok probe-scsi-all
/pci@400/pci@2/pci@0/pci@4/scsi@0                           <---- SAS controller

FCode Version 1.00.54, MPT Version 2.00, Firmware Version 5.00.17.00

Target 9
  Unit 0   Disk    SEAGATE   ST930003SSUN300G 0868 585937500 Blocks, 300 GB
  SASDeviceName  5000c5001cb4a637  SASAddress  5000c5001cb4a635  PhyNum 0
Target a
  Unit 0   Removable Read Only device   TEAC     DV-W28SS-R      1.0C
  SATA device  PhyNum 7
Target d
  Unit 0   Disk    HITACHI   H103030SCSUN300G A2A8   585937500 Blocks, 300 GB
  SASDeviceName  5000c5001cb477cb  SASAddress  5000c5001cb477c9  PhyNum 1
Target e
  Unit 0   Disk    SEAGATE   ST930003SSUN300G 0B70   585937500 Blocks, 300 GB
  SASDeviceName  5000c5001cb47f93  SASAddress  5000c5001cb47f91  PhyNum 2
Target f
  Unit 0   Disk    SEAGATE   ST930003SSUN300G 0B70   585937500 Blocks, 300 GB
```

```
   SASDeviceName  5000c5001cb47f7f  SASAddress  5000c5001cb47f7d  PhyNum 3
Target 389 Volume 0
  Unit 0   Disk   LSI     Logical Volume   3000   583983104 Blocks,   298 GB
  VolumeDeviceName 3ce534e42c02a3c0   VolumeWWID 0ce534e42c02a3c0

/pci@400/pci@1/pci@0/pci@b/pci@0/usb@0,2/hub@2/hub@3/storage@2
  Unit 0   Removable Read Only device    AMI    Virtual CDROM    1.00
```

Note – The `probe-scsi-all` output example shown above can be applied to the interpretation of `probe-scsi-all` output for SPARC T3-1B systems.

Note – The `probe-scsi-all` output for SPARC T3-1 and SPARC T3-4 servers has special characteristics due to their use of two on-board SAS controllers. In addition, on SPARC T3-1 servers with sixteen-disk capable backplanes, the two controllers connect to the hard drives through a SAS expander unit on the backplane. These differences are explained in "`probe-scsi-all` Output Example (SPARC T3-1, Eight-Disk Backplane)" on page 87 and "`probe-scsi-all` Output Example (SPARC T3-1, Sixteen-Disk Backplane)" on page 89.

Related Information

- "World Wide Name Syntax" on page 85
- "`probe-scsi-all` Output Example (SPARC T3-1, Eight-Disk Backplane)" on page 87
- "`probe-scsi-all` Output Example (SPARC T3-1, Sixteen-Disk Backplane)" on page 89
- "`probe-scsi-all` Output Example (SPARC T3-4)" on page 92
- "Identify a Disk Slot Using `prtconf` (Oracle Solaris, Onboard Controllers)" on page 98
- "Identify a Disk Slot Using `prtconf` (Oracle Solaris, Single Initiator)" on page 100
- "WWN Syntax in an OS Installation on a Specific Device" on page 103
- "WWN Syntax in an OS Installation on a RAID Volume" on page 105

▼ Identify a Disk Slot Using `prtconf` (Oracle Solaris, Onboard Controllers)

The procedure described below is intended for SPARC T3-1 servers with sixteen-disk capable backplanes. These systems include a SAS expander between the two on-board SAS controllers and the hard drive connectors. The 16 disk backplane is zoned such that the `phy-num` value associated with each drive is relative to its controller.

Note – For systems with an eight-disk capable backplane and two on-board controllers, each controller provides an interface to four of the eight drive slots. As a result, the `phy-num` value associated with each drive is relative to its controller. A example of the single-initiator configuration procedure is provided in "Identify a Disk Slot Using `prtconf` (Oracle Solaris, Single Initiator)" on page 100

---

1. **Run the** `format` **command.**

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

AVAILABLE DISK SELECTIONS:
    0. c0t5000C5001D3FF2D7d0 <SUN300G cyl 46873 alt 2 hd 20 sec 625>
       /scsi_vhci/disk@g5000c5001d3ff2d7
    1. c0t5000C5001D40D7F7d0 <SUN300G cyl 46873 alt 2 hd 20 sec 625>
       /scsi_vhci/disk@g5000c5001d40d7f7
[...]
```

The remaining steps in this example will identify the physical slot that corresponds to the device name `c0t5000C5001D3FF2D7d0`.

2. **Run** `prtconf -v` **and search for the device link** `c0t5000C5001D3FF2D7d0`**.**

```
Device Minor Nodes:
    dev=(32,0)
        dev_path=/scsi_vhci/disk@g5000c5001d3ff2d7:a
            spectype=blk type=minor
            dev_link=/dev/dsk/c0t5000C5001D3FF2D7d0s0      <<=== Device path
            dev_link=/dev/sd3a
        dev_path=/scsi_vhci/disk@g5000c5001d3ff2d7:a,raw
            spectype=chr type=minor
            dev_link=/dev/rdsk/c0t5000C5001D3FF2D7d0s0      <<=== Device path
            dev_link=/dev/rsd3a
```

3. **Search the** `prtconf` **output for a** `name='wwn'` **entry that has the WWN value** `5000c5001d3ff2d7`**.**

   Note the `target-port` value listed with that WWN value. In this example it is `5000c5001d3ff2d5`. This is the target ID of the physical disk.

```
Paths from multipath bus adapters:
    mpt_sas#2 (online)
        name='wwn' type=string items=1
            value='5000c5001d3ff2d7'              <<=== Hard drive WWN ID
        name='lun' type=int items=1
            value=00000000
        name='target-port' type=string items=1
            value='5000c5001d3ff2d5'              <<=== Hard drive target ID
        name='obp-path' type=string items=1
            value='/pci@400/pci@2/pci@0/pci@4/scsi@0/disk@w5000c5001d3ff2d5,0'
        name='phy-num' type=int items=1
            value=00000004                        <<=== Hard drive slot number
        name='path-class' type=string items=1
            value='primary'
    mpt_sas#5 (online)
        name='wwn' type=string items=1
            value='5000c5001d3ff2d7'
        name='lun' type=int items=1
            value=00000000
        name='target-port' type=string items=1
            value='5000c5001d3ff2d5'
        name='obp-path' type=string items=1
            value='/pci@400/pci@1/pci@0/pci@4/scsi@0/disk@w5000c5001d3ff2d5,0'
        name='phy-num' type=int items=1
            value=00000004
        name='path-class' type=string items=1
            value='primary'
```

Note – In a SPARC T3-1 system with a sixteen-disk backplane, both controllers list the logical name (the `wwn` value), a corresponding port name (target-port value), and physical disk location (`phy-num` value) for every connected hard drive.

4. **The value shown for the** `name='phy-num'` **entry indicates the physical slot containing the hard drive.**

   In this example, the target device is in slot 4.

Related Information

- "World Wide Name Syntax" on page 85
- "`probe-scsi-all` Output Example (SPARC T3-1, Eight-Disk Backplane)" on page 87

▼ Identify a Disk Slot Using `prtconf` (Oracle Solaris, Single Initiator)

The procedure described below is intended for SPARC T3-1 and SPARC T3-4 servers in a single-initiator, eight-disk backplane configuration.

1. Run the `format` **command.**

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

AVAILABLE DISK SELECTIONS:
    0. c0t5000CCA00ABBAEB8d0 <SUN300G cyl 46873 alt 2 hd 20 sec 625>
       /scsi_vhci/disk@g5000cca00abbaeb8
    1. c0t5000C5001D40D7F7d0 <SUN300G cyl 46873 alt 2 hd 20 sec 625>
       /scsi_vhci/disk@g5000c5001d40d7f7
[...]
```

The remaining steps in this example will identify the physical slot that corresponds to the device name c0t5000CCA00ABBAEB8d0.

2. Run `prtconf  -v` **and search for the device link** c0t5000CCA00ABBAEB8d0**.**

```
Device Minor Nodes:
    dev=(32,0)
        dev_path=/scsi_vhci/disk@g5000cca00abbaeb8:a
            spectype=blk type=minor
            dev_link=/dev/dsk/c0t5000CCA00ABBAEB8d0s0      <<=== Device path
            dev_link=/dev/sd3a
        dev_path=/scsi_vhci/disk@g5000cca00abbaeb8:a,raw
            spectype=chr type=minor
            dev_link=/dev/rdsk/c0t5000CCA00ABBAEB8d0s0     <<=== Device path
            dev_link=/dev/rsd3a
```

3. **Search the** `prtconf` **output for a** `name='wwn'` **entry that has the WWN value** `5000cca00abbaeb8`**.**

   Note the `obp-path` value listed under WWN `5000cca00abbaeb8`.

   Refer to the following table to find the controller.

SPARC T3-1	
Controller 0	`/pci@400/pci@1/pci@0/pci@4/scsi@0/disk@w5000cca00abbaeb9,0`
Controller 1	`/pci@400/pci@2/pci@0/pci@4/scsi@0/disk@w5000cca00abbaeb9,0`
SPARC T3-4 (4 processor)	
Controller 0	`/pci@400/pci@1/pci@0/pci@0/LSI,sas@0/disk@w5000cca00abbaeb9,0`
Controller 1	`/pci@700/pci@1/pci@0/pci@0/LSI,sas@0/disk@w5000cca00abbaeb9,0`
SPARC T3-4 (2 processor)	
Controller 0	`/pci@400/pci@1/pci@0/pci@0/LSI,sas@0/disk@w5000cca00abbaeb9,0`
Controller 1	`/pci@500/pci@1/pci@0/pci@0/LSI,sas@0/disk@w5000cca00abbaeb9,0`

For SPARC T3-1 servers,the controller is identified in the second field: `pci@1` = controller 0 and `pci@2` = controller 1.

For SPARC T3-4 servers, the controller is identified in the first field. For systems with a four-processor configuration, `pci@400` = controller 0 and `pci@700` = controller 1. For systems with a two-processor configuration, `pci@400` = controller 0 and `pci@500` = controller 1.

The following output example shows the `obp-path` for a SPARC T3-1 server.

```
Paths from multipath bus adapters:
    mpt_sas#5 (online)
        name='wwn' type=string items=1
            value='5000cca00abbaeb8'               <<=== Hard drive WWN ID
        name='lun' type=int items=1
            value=00000000
        name='target-port' type=string items=1
            value='5000cca00abbaeb9'               <<=== Hard drive Target ID
        name='obp-path' type=string items=1
            value='/pci@400/pci@1/pci@0/pci@4/scsi@0/disk@w5000cca00abbaeb9,0'
        name='phy-num' type=int items=1
            value=00000000
        name='path-class' type=string items=1
            value='primary'
```

In this SPARC T3-1 example, the obp-path is:

```
/pci@400/pci@1/pci@0/pci@4/scsi@0/disk@w5000cca00abbaeb9,0
```

Based on the previous table, this disk is on controller 0.

The following output example shows the `obp-path` for a SPARC T3-4 server.

```
Paths from multipath bus adapters:
    mpt_sas#5 (online)
        name='wwn' type=string items=1
            value='5000cca00abbaeb8'               <<=== Hard drive WWN ID
        name='lun' type=int items=1
            value=00000000
        name='target-port' type=string items=1
            value='5000cca00abbaeb9'               <<=== Hard drive Target ID
        name='obp-path' type=string items=1
            value='/pci@400/pci@1/pci@0/LSI,sas@0/disk@w5000cca00abbaeb9,0'
        name='phy-num' type=int items=1
            value=00000000
        name='path-class' type=string items=1
            value='primary'
```

In this SPARC T3-4 example, the obp-path is:

```
/pci@400/pci@1/pci@0/pci@0/LSI,sas@0/disk@w5000cca00abbaeb9,0
```

Based on the previous table, this disk is on controller 0.

4. **This `phy-num` value corresponds to physical disk slot `0`, as shown in the following port mapping table.**

SAS Controller	PhyNum	Disk Slot	SAS Controller	PhyNum	Disk Slot
0	0	0	1	0	4
	1	1		1	5
	2	2		2	6
	3	3		3	7

Related Information

WWN Syntax in an OS Installation on a Specific Device

The following Oracle Solaris Jumpstart profile example shows how to use the WWN syntax when installing the operating system on a specific disk drive. In this example, the device name contains the WWN value `5000CCA00A75DCAC`.

Note – Oracle Solaris syntax rules require all alpha characters to be capitalized.

```
#
install_type flash_install
boot_device c0t5000CCA00A75DCACd0s0     preserve
```

```
archive_location nfs
129.148.94.249:/export/install/media/solaris/builds/s10u9/flar/latest.flar

# Disk layouts
#
partitioning explicit
filesys rootdisk.s0         free /
filesys rootdisk.s1         8192 swap
```

Related Information

- "World Wide Name Syntax" on page 85
- "`probe-scsi-all` Output Example (SPARC T3-1, Eight-Disk Backplane)" on page 87
- "`probe-scsi-all` Output Example (SPARC T3-1, Sixteen-Disk Backplane)" on page 89
- "`probe-scsi-all` Output Example (SPARC T3-4)" on page 92
- "Identify a Disk Slot Using `probe-scsi-all` (OBP)" on page 96
- "Identify a Disk Slot Using `prtconf` (Oracle Solaris, Onboard Controllers)" on page 98
- "Identify a Disk Slot Using `prtconf` (Oracle Solaris, Single Initiator)" on page 100
- "WWN Syntax in an OS Installation on a RAID Volume" on page 105

WWN Syntax in an OS Installation on a RAID Volume

The following Oracle Solaris Jumpstart profile example shows how to use the WWN syntax when installing the OS on a RAID volume. When installing software on a RAID volume, use the virtual device's VolumeDeviceName instead of an individual device name. In this example, the RAID volume name is 3ce534e42c02a3c0.

```
#
install_type flash_install
boot_device 3ce534e42c02a3c0       preserve

archive_location nfs
129.148.94.249:/export/install/media/solaris/builds/s10u9/flar/latest.flar

# Disk layouts
```

```
#
partitioning explicit
filesys rootdisk.s0        free /
filesys rootdisk.s1        8192 swap
```

Related Information

Glossary

B

BMC   baseboard management controller

C

CMA   cable management arm

D

DHCP   Dynamic Host Configuration Protocol

DTE   data terminal equipment

E

ESD   electrostatic discharge

H

HBA   host bus adapter

I

ILOM   Oracle Integrated Lights Out Manager

IP   Internet Protocol

N

NET MGT   network management port

NIC   network interface card or controller

O

Oracle Solaris OS   Oracle Solaris Operation System

POST

POST   power-on self-test

Q

QSFP   quad small form-factor pluggable

S

SAS	serial attached SCSI
SER MGT	serial management port
SP	service processor
SSD	sold-state drive
SSH	Secure Shell

U

UI	user interface
UUID	Universal Unique Identifier

W

WWID	world-wide identifier. A unique number that identifies a SAS target.

Index