



Sun SPARC[®] Enterprise T5120 and T5220 Servers Administration Guide

Sun Microsystems, Inc.
www.sun.com

Part No. 820-2179-10
August 2007, Revision A

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

Copyright 2007 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, U.S.A. All rights reserved.

Fujitsu Limited provided technical input and review on portions of this material.

Sun Microsystems, Inc. and Fujitsu Limited each own or control intellectual property rights relating to products and technology described in this document, and such products, technology and this document are protected by copyright laws, patents and other intellectual property laws and international treaties. The intellectual property rights of Sun Microsystems, Inc. and Fujitsu Limited in such products, technology and this document include, without limitation, one or more of the United States patents listed at <http://www.sun.com/patents> and one or more additional patents or patent applications in the United States or other countries.

This document and the product and technology to which it pertains are distributed under licenses restricting their use, copying, distribution, and decompilation. No part of such product or technology, or of this document, may be reproduced in any form by any means without prior written authorization of Fujitsu Limited and Sun Microsystems, Inc., and their applicable licensors, if any. The furnishing of this document to you does not give you any rights or licenses, express or implied, with respect to the product or technology to which it pertains, and this document does not contain or represent any commitment of any kind on the part of Fujitsu Limited or Sun Microsystems, Inc., or any affiliate of either of them.

This document and the product and technology described in this document may incorporate third-party intellectual property copyrighted by and/or licensed from suppliers to Fujitsu Limited and/or Sun Microsystems, Inc., including software and font technology.

Per the terms of the GPL or LGPL, a copy of the source code governed by the GPL or LGPL, as applicable, is available upon request by the End User. Please contact Fujitsu Limited or Sun Microsystems, Inc.

This distribution may include materials developed by third parties.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and in other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, Java, Netra, Solaris, Sun StorEdge, SPARC Enterprise, docs.sun.com, OpenBoot, SunVTS, Sun Fire, SunSolve, CoolThreads, J2EE, and Sun are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

Fujitsu and the Fujitsu logo are registered trademarks of Fujitsu Limited.

All SPARC trademarks are used under license and are registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon architecture developed by Sun Microsystems, Inc.

SPARC64 is a trademark of SPARC International, Inc., used under license by Fujitsu Microelectronics, Inc. and Fujitsu Limited.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

United States Government Rights - Commercial use. U.S. Government users are subject to the standard government user license agreements of Sun Microsystems, Inc. and Fujitsu Limited and the applicable provisions of the FAR and its supplements.

Disclaimer: The only warranties granted by Fujitsu Limited, Sun Microsystems, Inc. or any affiliate of either of them in connection with this document or any product or technology described herein are those expressly set forth in the license agreement pursuant to which the product or technology is provided. EXCEPT AS EXPRESSLY SET FORTH IN SUCH AGREEMENT, FUJITSU LIMITED, SUN MICROSYSTEMS, INC. AND THEIR AFFILIATES MAKE NO REPRESENTATIONS OR WARRANTIES OF ANY KIND (EXPRESS OR IMPLIED) REGARDING SUCH PRODUCT OR TECHNOLOGY OR THIS DOCUMENT, WHICH ARE ALL PROVIDED AS IS, AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING WITHOUT LIMITATION ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. Unless otherwise expressly set forth in such agreement, to the extent allowed by applicable law, in no event shall Fujitsu Limited, Sun Microsystems, Inc. or any of their affiliates have any liability to any third party under any legal theory for any loss of revenues or profits, loss of use or data, or business interruptions, or for any indirect, special, incidental or consequential damages, even if advised of the possibility of such damages.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.



Adobe PostScript

Copyright 2007 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, Etats-Unis. Tous droits réservés.

Entrée et revue tecnica fournies par Fujitsu Limited sur des parties de ce matériel.

Sun Microsystems, Inc. et Fujitsu Limited détiennent et contrôlent toutes deux des droits de propriété intellectuelle relatifs aux produits et technologies décrits dans ce document. De même, ces produits, technologies et ce document sont protégés par des lois sur le copyright, des brevets, d'autres lois sur la propriété intellectuelle et des traités internationaux. Les droits de propriété intellectuelle de Sun Microsystems, Inc. et Fujitsu Limited concernant ces produits, ces technologies et ce document comprennent, sans que cette liste soit exhaustive, un ou plusieurs brevets déposés aux États-Unis et indiqués à l'adresse <http://www.sun.com/patents> de même qu'un ou plusieurs brevets ou applications brevetées supplémentaires aux États-Unis et dans d'autres pays.

Ce document, le produit et les technologies afférents sont exclusivement distribués avec des licences qui en restreignent l'utilisation, la copie, la distribution et la décompilation. Aucune partie de ce produit, de ces technologies ou de ce document ne peut être reproduite sous quelque forme que ce soit, par quelque moyen que ce soit, sans l'autorisation écrite préalable de Fujitsu Limited et de Sun Microsystems, Inc., et de leurs éventuels bailleurs de licence. Ce document, bien qu'il vous ait été fourni, ne vous confère aucun droit et aucune licence, expresses ou tacites, concernant le produit ou la technologie auxquels il se rapporte. Par ailleurs, il ne contient ni ne représente aucun engagement, de quelque type que ce soit, de la part de Fujitsu Limited ou de Sun Microsystems, Inc., ou des sociétés affiliées.

Ce document, et le produit et les technologies qu'il décrit, peuvent inclure des droits de propriété intellectuelle de parties tierces protégés par copyright et/ou cédés sous licence par des fournisseurs à Fujitsu Limited et/ou Sun Microsystems, Inc., y compris des logiciels et des technologies relatives aux polices de caractères.

Par limites du GPL ou du LGPL, une copie du code source régi par le GPL ou LGPL, comme applicable, est sur demande vers la fin utilisateur disponible; veuillez contacter Fujitsu Limited ou Sun Microsystems, Inc.

Cette distribution peut comprendre des composants développés par des tierces parties.

Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux États-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, Java, Netra, Solaris, Sun StorEdge, SPARC Enterprise, docs.sun.com, OpenBoot, SunVTS, Sun Fire, SunSolve, CoolThreads, J2EE, et Sun sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc. aux États-Unis et dans d'autres pays.

Fujitsu et le logo Fujitsu sont des marques déposées de Fujitsu Limited.

Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux États-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

SPARC64 est une marque déposée de SPARC International, Inc., utilisée sous le permis par Fujitsu Microelectronics, Inc. et Fujitsu Limited.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui, en outre, se conforment aux licences écrites de Sun.

Droits du gouvernement américain - logiciel commercial. Les utilisateurs du gouvernement américain sont soumis aux contrats de licence standard de Sun Microsystems, Inc. et de Fujitsu Limited ainsi qu'aux clauses applicables stipulées dans le FAR et ses suppléments.

Avis de non-responsabilité: les seules garanties octroyées par Fujitsu Limited, Sun Microsystems, Inc. ou toute société affiliée de l'une ou l'autre entité en rapport avec ce document ou tout produit ou toute technologie décrit(e) dans les présentes correspondent aux garanties expressément stipulées dans le contrat de licence régissant le produit ou la technologie fourni(e). SAUF MENTION CONTRAIRE EXPRESSEMENT STIPULÉE DANS CE CONTRAT, FUJITSU LIMITED, SUN MICROSYSTEMS, INC. ET LES SOCIÉTÉS AFFILIÉES REJETTENT TOUTE REPRÉSENTATION OU TOUTE GARANTIE, QUELLE QU'EN SOIT LA NATURE (EXPRESSE OU IMPLICITE) CONCERNANT CE PRODUIT, CETTE TECHNOLOGIE OU CE DOCUMENT, LESQUELS SONT FOURNIS EN L'ÉTAT. EN OUTRE, TOUTES LES CONDITIONS, REPRÉSENTATIONS ET GARANTIES EXPRESSES OU TACITES, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE À LA QUALITÉ MARCHANDE, À L'APTITUDE À UNE UTILISATION PARTICULIÈRE OU À L'ABSENCE DE CONTREFAÇON, SONT EXCLUES, DANS LA MESURE AUTORISÉE PAR LA LOI APPLICABLE. Sauf mention contraire expressément stipulée dans ce contrat, dans la mesure autorisée par la loi applicable, en aucun cas Fujitsu Limited, Sun Microsystems, Inc. ou l'une de leurs filiales ne sauraient être tenues responsables envers une quelconque partie tierce, sous quelque théorie juridique que ce soit, de tout manque à gagner ou de perte de profit, de problèmes d'utilisation ou de perte de données, ou d'interruptions d'activités, ou de tout dommage indirect, spécial, secondaire ou consécutif, même si ces entités ont été préalablement informées d'une telle éventualité.

LA DOCUMENTATION EST FOURNIE "EN L'ÉTAT" ET TOUTES AUTRES CONDITIONS, DECLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES, DANS LA MESURE AUTORISÉE PAR LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE À LA QUALITÉ MARCHANDE, À L'APTITUDE À UNE UTILISATION PARTICULIÈRE OU À L'ABSENCE DE CONTREFAÇON.

Contents

Preface xiii

1. Configuring the System Console 1

Communicating With the System 1

What the System Console Does 3

Using the System Console 3

Default System Console Connection Through the Serial Management and Network Management Ports 4

Alternative System Console Configuration 7

Accessing the System Console Through a Graphics Monitor 8

Accessing the Service Processor 8

Using the Serial Management Port 8

▼ To Use the Serial Management Port 8

Activating the Network Management Port 9

▼ To Activate the Network Management Port 10

Accessing the System Console Through a Terminal Server 11

▼ To Access The System Console Through a Terminal Server 11

Accessing the System Console Through a Tip Connection 13

▼ To Access the System Console Through the Tip Connection 14

Modifying the `/etc/remote` File 15

▼ To Modify the <code>/etc/remote</code> File	15
Accessing the System Console Through an Alphanumeric Terminal	16
▼ To Access the System Console Through an Alphanumeric Terminal	17
Accessing the System Console Through a Local Graphics Monitor	17
▼ To Access the System Console Through a Local Graphics Monitor	18
Switching Between the Service Processor and the System Console	19
ILOM -> Prompt	21
Access Through Multiple Controller Sessions	22
Reaching the -> Prompt	22
OpenBoot ok Prompt	23
Reaching the ok Prompt	24
Graceful Shutdown	24
ILOM set <code>/HOST send_break_action=break,</code> start <code>/SP/console</code> Commands, or Break Key	25
Manual System Reset	25
Getting to the ok Prompt	26
▼ To Obtain the ok Prompt	27
For More Information	27
System Console OpenBoot Configuration Variable Settings	28
2. Managing RAS Features and System Firmware	29
ILOM and The Service Processor	29
Logging In To ILOM	30
▼ To Log In To ILOM	31
▼ To View Environmental Information	31
Interpreting System LEDs	32
Controlling the Locator LED	33
OpenBoot Emergency Procedures	35

OpenBoot Emergency Procedures for SPARC Enterprise T5120 and T5220 Systems	35
Stop-N Functionality	36
▼ To Restore OpenBoot Configuration Defaults	36
Stop-F Functionality	37
Stop-D Functionality	37
Automatic System Recovery	37
Auto-Boot Options	38
Error Handling Summary	38
Reset Scenarios	39
Automatic System Recovery User Commands	40
Enabling and Disabling Automatic System Recovery	41
▼ To Enable Automatic System Recovery	41
▼ To Disable Automatic System Recovery	41
Obtaining Automatic System Recovery Information	42
Unconfiguring and Reconfiguring Devices	43
▼ To Unconfigure a Device Manually	43
▼ To Reconfigure a Device Manually	44
Displaying System Fault Information	44
▼ To Display Current Valid System Faults	44
▼ To Clear a Fault	45
Storing FRU Information	45
▼ To Store Information in Available FRU PROMs	45
Multipathing Software	46
For More Information	46
3. Managing Disk Volumes	47
Requirements	47
Disk Volumes	47

RAID Technology	48
Integrated Stripe Volumes (RAID 0)	48
Integrated Mirror Volumes (RAID 1)	49
Hardware Raid Operations	50
Physical Disk Slot Numbers, Physical Device Names, and Logical Device Names for Non-RAID Disks	51
▼ To Create a Hardware Mirrored Volume	51
▼ To Create a Hardware Mirrored Volume of the Default Boot Device	54
▼ To Create a Hardware Striped Volume	55
▼ To Configure and Label a Hardware RAID Volume for Use in the Solaris Operating System	57
▼ To Delete a Hardware RAID Volume	60
▼ To Perform a Mirrored Disk Hot-Plug Operation	62
▼ To Perform a Nonmirrored Disk Hot-Plug Operation	63
4. Logical Domains Software	67
About Logical Domains Software	67
Logical Domain Configurations	68
Logical Domains Software Requirements	68
A. OpenBoot Configuration Variables	71
Index	75

Figures

FIGURE 1-1	Directing the System Console Input and Output	4
FIGURE 1-2	Rear I/O Panel of the SPARC Enterprise T5120 Chassis.	6
FIGURE 1-3	Rear I/O Panel of the SPARC Enterprise T5220 Chassis.	6
FIGURE 1-4	Patch Panel Connection Between a Terminal Server and a SPARC Enterprise T5120 or T5220 Server	12
FIGURE 1-5	Tip Connection Between a SPARC Enterprise T5120 or T5220 Server and Another System	14
FIGURE 1-6	Separate System Console and Service Processor Channels	20
FIGURE 2-1	Locator Button (top button) on SPARC Enterprise T5120 Chassis	34
FIGURE 2-2	Locator Button (top button) on SPARC Enterprise T5220 Chassis	34
FIGURE 3-1	Graphical Representation of Disk Striping	49
FIGURE 3-2	Graphical Representation of Disk Mirroring	50

Tables

TABLE 1-1	Ways of Communicating With the System	2
TABLE 1-2	Pin Crossovers for Connecting to a Typical Terminal Server	12
TABLE 1-3	Ways of Accessing the <code>ok</code> Prompt	27
TABLE 1-4	OpenBoot Configuration Variables That Affect the System Console	28
TABLE 2-1	Standard LED Behaviors and Values	32
TABLE 2-2	System LED Behaviors with Assigned Meanings	32
TABLE 2-3	Virtual Keyswitch Setting for Reset Scenario	40
TABLE 2-4	ILOM Property Settings for Reset Scenario	40
TABLE 2-5	Device Identifiers and Devices	43
TABLE 3-1	Disk Slot Numbers, Logical Device Names, and Physical Device Names	51
TABLE 4-1	Logical Domain Roles	68
TABLE A-1	OpenBoot Configuration Variables Stored on the System Configuration Card	71

Preface

The *Sun SPARC Enterprise T5120 & T5220 Server System Administration Guide* is for experienced system administrators. The guide includes general descriptive information about the Sun SPARC Enterprise™ T5120 and T5220 servers, and detailed instructions for configuring and administering the servers. To use the information in this document, you must have working knowledge of computer network concepts and terms, and advanced familiarity with the Solaris™ Operating System (Solaris OS).

Note – For information about changing the hardware configuration of your server, or about running diagnostics, see the service manual for your server.

How This Book Is Organized

The *Sun SPARC Enterprise T5120 & T5220 Server System Administration Guide* is divided into the following chapters:

- Chapter 1 describes the system console and how to access it.
- Chapter 2 describes the tools used to configure system firmware, including system controller environmental monitoring, Automatic System Recovery (ASR), and multipathing software. In addition, the chapter describes how to unconfigure and reconfigure a device manually.
- Chapter 3 describes redundant array of independent disks (RAID) concepts, and how to configure and manage RAID disk volumes using your server's onboard serial attached SCSI (SAS) disk controller.
- Chapter 4 describes Logical Domain software.

This document also includes the following reference appendix:

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

Using UNIX Commands

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

- Software documentation that you received with your system
- Solaris OS documentation

Shell Prompts

Shell	Prompt
C shell	<i>machine-name%</i>
C shell superuser	<i>machine-name#</i>
Bourne shell and Korn shell	\$
Bourne shell and Korn shell superuser	#

Typographic Conventions

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

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

Related Documentation

The documents listed as online are available at:

<http://www.sun.com/documentation>

Title	Description	Part Number
<i>Sun SPARC Enterprise T5120 and T5220 Servers Product Notes</i>	Information about the latest product updates and issues	820-2176
<i>Sun SPARC Enterprise T5120 and T5220 Servers Product Overview Guide</i>	Product features	820-2183
<i>Sun SPARC Enterprise T5120 and T5220 Servers Site Planning Guide</i>	Server specifications for site planning	820-2177
<i>Sun SPARC Enterprise T5120 and T5220 Servers Installation Guide</i>	Detailed rackmounting, cabling, power-on, and configuring information	820-2178
<i>Integrated Lights Out Management 2.0 (ILOM 2.0) Supplement for Sun SPARC Enterprise T5120 and T5220 Servers</i>	How to use the Integrated Lights Out Manager (ILOM) software	Varies based on version
<i>Sun SPARC Enterprise T5120 and T5220 Servers Service Manual</i>	How to run diagnostics to troubleshoot the server, and how to remove and replace parts in the server	820-2181
<i>Sun SPARC Enterprise T5120 and T5220 Servers Safety and Compliance manual</i>	Safety and compliance information about this server	

Documentation, Support, and Training

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

Sun Welcomes Your Comments

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

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

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

Sun SPARC Enterprise T5120 and T5220 Server Administration Guide, part number 820-2179-10

Configuring the System Console

This chapter explains what the system console is, describes the different ways of configuring it on a SPARC Enterprise T5120 or T5220 server, and helps you understand the relationship between the system console and the service processor. This chapter contains the following sections:

- [“Communicating With the System” on page 1](#)
- [“Accessing the Service Processor” on page 8](#)
- [“Switching Between the Service Processor and the System Console” on page 19](#)
- [“ILOM -> Prompt” on page 21](#)
- [“OpenBoot ok Prompt” on page 23](#)
- [“System Console OpenBoot Configuration Variable Settings” on page 28](#)

Note – For information about changing the hardware configuration of your server, or about running diagnostics, see the service manual for your server.

Communicating With the System

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

The serial management port (SER MGT) is the default port for accessing the system console upon initial system installation. After installation, you can configure the system console to accept input from and send output to different devices. [TABLE 1-1](#) lists these devices and where they are discussed in this document.

TABLE 1-1 Ways of Communicating With the System

Devices Available	During Installation	After Installation	Further Information
A terminal server attached to the serial management port (SER MGT)	X	X	“Accessing the Service Processor” on page 8
	X	X	“Accessing the System Console Through a Terminal Server” on page 11
	X	X	“System Console OpenBoot Configuration Variable Settings” on page 28
An alphanumeric terminal or similar device attached to the serial management port (SER MGT)	X	X	“Accessing the Service Processor” on page 8
	X	X	“Accessing the System Console Through an Alphanumeric Terminal” on page 16
	X	X	“System Console OpenBoot Configuration Variable Settings” on page 28
A Tip line attached to the serial management port (SER MGT)	X	X	“Accessing the Service Processor” on page 8
	X	X	“Accessing the System Console Through a Tip Connection” on page 13
		X	“Modifying the <code>/etc/remote</code> File” on page 15
	X	X	“System Console OpenBoot Configuration Variable Settings” on page 28
An Ethernet line connected to the network management port (NET MGT)		X	“Activating the Network Management Port” on page 9
A local graphics monitor (graphics accelerator card, graphics monitor, mouse, and keyboard).		X	“Accessing the System Console Through a Local Graphics Monitor” on page 17
		X	“System Console OpenBoot Configuration Variable Settings” on page 28

What the System Console Does

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

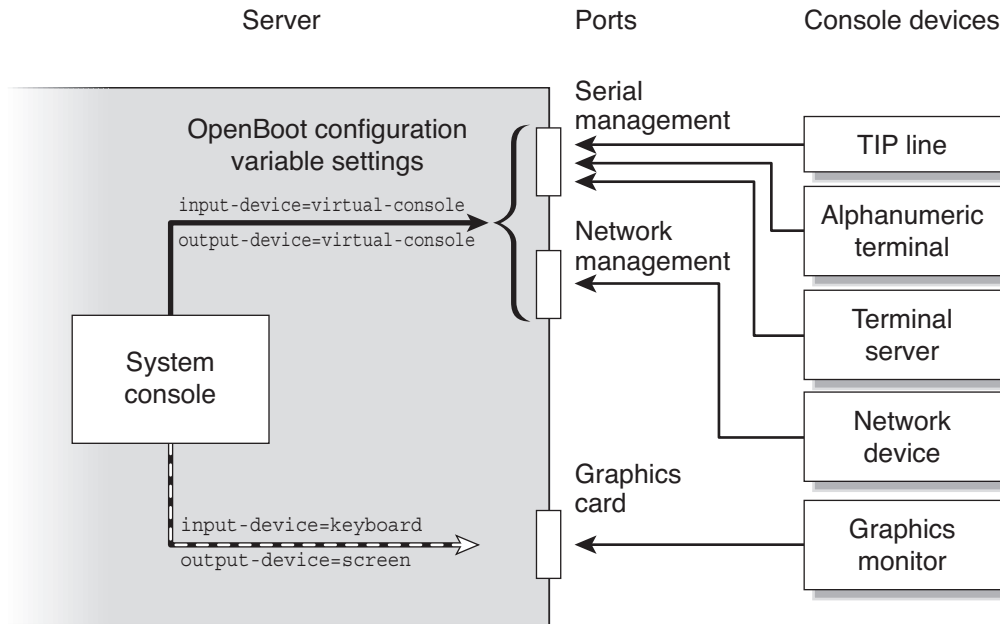
Once the operating system is booted, the system console displays UNIX system messages and accepts UNIX commands.

Using the System Console

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

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

FIGURE 1-1 Directing the System Console Input and Output



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

On your server, the system console comes preconfigured to allow input and output only by means of the service processor. The service processor must be accessed either through the serial management port (SER MGT) or the network management port (NET MGT). By default, the network management port is configured to retrieve network configuration using Dynamic Host Configuration Protocol (DHCP) and to allow connections using Secure Shell (SSH). You can modify the network management port configuration after connecting to ILOM through either the serial or network management ports.

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

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

These devices provide for secure access at the installation site.

FIGURE 1-2 Rear I/O Panel of the SPARC Enterprise T5120 Chassis.




Figure Legend Management Ports, ttya port, and Locator LED on the SPARC Enterprise T5120.

1	Locator LED	3	DB-9 (ttya)
2	SER MGT Port	4	NET MGT Port

FIGURE 1-3 Rear I/O Panel of the SPARC Enterprise T5220 Chassis.




Figure Legend Management Ports, ttya port, and Locator LED on the SPARC Enterprise T5220.

1	Locator LED	3	DB-9 (ttya)
2	SER MGT Port	4	NET MGT Port

Figure Legend Management Ports and Locator LED of the SPARC Enterprise T5220

The service processor serial management port is the default console connection.

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

The serial management port is not a general-purpose serial port. If you want to use a general-purpose serial port with your server—to connect a serial printer, for instance—use the standard 9-pin serial port on the back panel of the SPARC Enterprise T5120 and T5220. The Solaris OS sees this port as `ttya`.

- For instructions on accessing the system console through a terminal server, see [“Accessing the System Console Through a Terminal Server” on page 11](#).
- For instructions on accessing the system console through an alphanumeric terminal, see [“Accessing the System Console Through an Alphanumeric Terminal” on page 16](#).
- For instructions on accessing the system console through a Tip line, see [“Accessing the System Console Through a Tip Connection” on page 13](#).

Alternative System Console Configuration

In the default configuration, service processor alerts and system console output appear interspersed in the same window. *After initial system installation*, you can redirect the system console to take its input from and send its output to a graphics card’s port.

The best practice is to leave the console port in its default configuration for the following reasons:

- In a default configuration, the serial management and network management ports enable you to open up to eight additional windows through which you can view, but not affect, system console activity. You cannot open these connections if the system console is redirected to a graphics card’s port.
- In a default configuration, the serial management and network management ports enable you to switch between viewing system console and service processor output on the same device by typing a simple escape sequence or command. The escape sequence and command do not work if the system console is redirected to a graphics card’s port.
- The service processor keeps a log of console messages, but some messages are not logged if the system console is redirected to a graphic card’s port. The omitted information could be important if you need to contact customer service about a problem.

You change the system console configuration by setting OpenBoot configuration variables. See [“System Console OpenBoot Configuration Variable Settings” on page 28](#).

Accessing the System Console Through a Graphics Monitor

The SPARC Enterprise T5120 or T5220 server is shipped without a mouse, keyboard, monitor, or frame buffer for the display of bitmapped graphics. To install a graphics monitor on the server, you must install a graphics accelerator card into a PCI slot, and attach a monitor, mouse, and keyboard to the appropriate front or rear USB ports.

After starting the system, you might need to install the correct software driver for the PCI card you have installed. For detailed hardware instructions, see [“Accessing the System Console Through a Local Graphics Monitor” on page 17](#).

Note – POST diagnostics cannot display status and error messages to a local graphics monitor.

Accessing the Service Processor

The following sections describe methods of accessing the service processor.

Using the Serial Management Port

This procedure assumes that the system console uses the serial management and network management ports (the default configuration).

When you are accessing the system console using a device connected to the serial management port, you first access the ILOM service processor and its `->` prompt. After connecting to the ILOM service processor, you can switch to the system console.

For more information about the ILOM service processor, refer to the ILOM user's guide and the ILOM supplement for your server.

▼ To Use the Serial Management Port

1. **Ensure that the serial port on your connecting device is set to the following parameters:**
 - 9600 baud
 - 8 bits

- No parity
- 1 stop bit
- No handshaking

2. Establish an ILOM service processor session.

See the ILOM user's guide for instructions.

3. To connect to the system console, at the ILOM command prompt, type:

```
-> start /SP/console
```

The `start /SP/console` command switches you to the system console.

4. To switch back to the -> prompt, type the #. (Pound-Period) escape sequence.

```
ok #.
```

Characters are not echoed to the screen.

For instructions on how to use the ILOM service processor, refer to the ILOM user's guide and the ILOM supplement for your server.

Activating the Network Management Port

The network management port is configured by default to retrieve network settings using DHCP and allow connections using SSH. You might need to modify these settings for your network. If you are unable to use DHCP and SSH on your network, you must connect to the service processor using the serial management port to reconfigure the network management port. See [“Using the Serial Management Port” on page 8](#)

Note – The default username when connecting to the service processor for the first time is `root`. The default password is `changeme`. You should assign a new password during initial system configuration. For more information, refer to your server installation guide, the ILOM user's guide, and the ILOM supplement for your server.

You can assign the network management port a static IP address or you can configure the port to obtain an IP address using DHCP from another server. The network management port can be configured to accept connections from SSH clients.

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

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

▼ To Activate the Network Management Port

1. **Connect an Ethernet cable to the network management port.**
2. **Log in to the ILOM service processor through the serial management port.**
See the ILOM user's guide for instructions.
3. **Type one of the following commands:**
 - **If your network uses static IP addresses, type:**

```
-> set /SP/network state=enabled
Set 'state' to 'enabled'

-> set /SP/network pendingipaddress=10.8.31.163
Set 'pendingipaddress' to '10.8.31.163'

-> set /SP/network pendingipdiscovery=static
Set 'pendingipdiscovery' to 'static'

-> set /SP/network pendingipnetmask=255.255.252.0
Set 'pendingipnetmask' to '255.255.252.0'

-> set /SP/network pendingipgateway=10.8.31.248
Set 'pendingipgateway' to '10.8.31.248'

-> set /SP/network commitpending=true
Set 'commitpending' to 'true'

->
```

Note – If you have configured your server to use static IP addresses, but wish to reset your network to use Dynamic Host Configuration Protocol (DHCP), type the following commands:

```
-> set /SP/network pendingipdiscovery=dhcp
Set 'pendingipdiscovery' to 'dhcp'

-> set /SP/network commitpending=true
Set 'commitpending' to 'true'

->
```

4. Issue the following command to verify network settings:

```
-> show /SP/network
```

To connect through the network management port, use `ssh` to the IP address you specified in [Step 3](#).

Accessing the System Console Through a Terminal Server

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

▼ To Access The System Console Through a Terminal Server


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

The serial management port on the SPARC Enterprise T5120 or T5220 server is a data terminal equipment (DTE) port. The pinouts for the serial management port correspond with the pinouts for the RJ-45 ports on the Serial Interface Breakout Cable supplied by Cisco for use with the Cisco AS2511-RJ terminal server. If you use a terminal server made by another manufacturer, check that the serial port pinouts of the SPARC Enterprise T5120 or T5220 server matches those of the terminal server you plan to use.

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

- Connect a serial interface breakout cable directly to the SPARC Enterprise T5120 or T5220 server. See “Accessing the Service Processor” on page 8.
- Connect a serial interface breakout cable to a patch panel and use the straight-through patch cable (supplied by your server’s manufacturer) to connect the patch panel to the server.

FIGURE 1-4 Patch Panel Connection Between a Terminal Server and a SPARC Enterprise T5120 or T5220 Server



If the pinouts for the serial management port *do not* correspond with the pinouts for the RJ-45 ports on the terminal server, you need to make a crossover cable that connects each pin on the SPARC Enterprise T5120 or T5220 server serial management port to the corresponding pin in the terminal server’s serial port.

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

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

SPARC Enterprise T5120 or T5220 Serial Port (RJ-45 Connector) Pin	Terminal Server Serial Port Pin
Pin 1 (RTS)	Pin 1 (CTS)
Pin 2 (DTR)	Pin 2 (DSR)
Pin 3 (TXD)	Pin 3 (RXD)

TABLE 1-2 Pin Crossovers for Connecting to a Typical Terminal Server *(Continued)*

SPARC Enterprise T5120 or T5220 Serial Port (RJ-45 Connector) Pin	Terminal Server Serial Port Pin
Pin 4 (Signal Ground)	Pin 4 (Signal Ground)
Pin 5 (Signal Ground)	Pin 5 (Signal Ground)
Pin 6 (RXD)	Pin 6 (TXD)
Pin 7 (DSR /DCD)	Pin 7 (DTR)
Pin 8 (CTS)	Pin 8 (RTS)

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

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


For example, for a SPARC Enterprise T5120 or T5220 server connected to port 10000 on a terminal server whose IP address is 192.20.30.10, you would type:

```
% ssh 192.20.30.10 10000
```

Accessing the System Console Through a Tip Connection

Use this procedure to access the SPARC Enterprise T5120 or T5220 server system console by connecting the serial management port (SER MGT) to the serial port of another system ([FIGURE 1-5](#)).

FIGURE 1-5 Tip Connection Between a SPARC Enterprise T5120 or T5220 Server and Another System



▼ To Access the System Console Through the Tip Connection

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

The cable and adapter connect between another system's serial port (typically `ttyb`) and the serial management port on the back panel of the SPARC Enterprise T5120 or T5220 server.

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

Most releases of Solaris OS software shipped since 1992 contain an `/etc/remote` file with the appropriate `hardwire` entry. However, if the system is running an older version of Solaris OS software, or if the `/etc/remote` file has been modified, you might need to edit it. See [“Modifying the `/etc/remote` File” on page 15](#) for details.

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

```
% tip hardware
```

The system responds by displaying:

```
connected
```

The shell tool is now a Tip window directed to the SPARC Enterprise T5120 or T5220 server through the system's serial port. This connection is established and maintained even when the server is completely powered off or just starting up.

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

Modifying the `/etc/remote` File

This procedure might be necessary if you are accessing the SPARC Enterprise T5120 or T5220 server using a Tip connection from a system running an older version of the Solaris OS software. You might also need to perform this procedure if the `/etc/remote` file on the system has been altered and no longer contains an appropriate `hardware` entry.

▼ To Modify the `/etc/remote` File

1. Log in as superuser to the system console of a system that you intend to use to establish a Tip connection to your server.
2. Determine the release level of Solaris OS software installed on the system.
Type:

```
# uname -r
```

The system responds with a release number.

3. Take one of the following actions, depending on the number displayed.

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

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

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

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

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

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

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

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

The `/etc/remote` file is now properly configured. Continue establishing a Tip connection to the SPARC Enterprise T5120 or T5220 server system console. See [“Accessing the System Console Through a Tip Connection” on page 13](#).

If you have redirected the system console to `ttyb` and want to change the system console settings back to use the serial management and network management ports, see [“System Console OpenBoot Configuration Variable Settings” on page 28](#).

Accessing the System Console Through an Alphanumeric Terminal

Use this procedure when you are accessing the SPARC Enterprise T5120 or T5220 server system console by connecting the serial port of an alphanumeric terminal to the serial management port (SER MGT) of the server.

▼ To Access the System Console Through an Alphanumeric Terminal

1. **Attach one end of the serial cable to the alphanumeric terminal's serial port.**
Use a null modem serial cable or an RJ-45 serial cable and null modem adapter. Connect this cable to the terminal's serial port connector.
2. **Attach the opposite end of the serial cable to the serial management port on the SPARC Enterprise T5120 or T5220 server.**
3. **Connect the alphanumeric terminal's power cord to an AC outlet.**
4. **Set the alphanumeric terminal to receive:**
 - 9600 baud
 - 8 bits
 - No parity
 - 1 stop bit
 - No handshake protocol

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

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

For more information about connecting to and using the ILOM service processor, refer to the ILOM guide and the ILOM supplement for your server.

Accessing the System Console Through a Local Graphics Monitor

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

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

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

- Supported USB mouse

▼ To Access the System Console Through a Local Graphics Monitor

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

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

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

Tighten the thumbscrews to secure the connection.

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

4. Connect the USB keyboard cable to one USB port and the USB mouse cable to the other USB port on the SPARC Enterprise T5120 or T5220 server back panel (FIGURE 1-3).

5. Get to the `ok` prompt.

For more information, see [“Getting to the `ok` Prompt” on page 26](#).

6. Set OpenBoot configuration variables appropriately.

From the existing system console, type:

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

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

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

```
ok reset-all
```

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

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

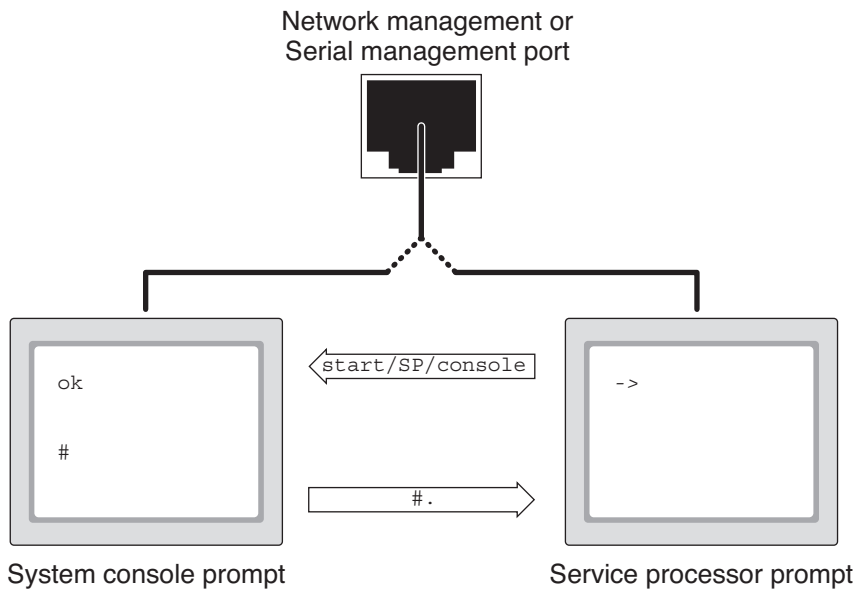
You can now issue system commands and view system messages using your local graphics monitor. Continue with your installation or diagnostic procedure, as needed.

If you want to redirect the system console back to the serial management and network management ports, see [“System Console OpenBoot Configuration Variable Settings” on page 28](#).

Switching Between the Service Processor and the System Console

The service processor features two management ports, labeled SER MGT and NET MGT, located on the server’s back panel. If the system console is directed to use the serial management and network management ports (the default configuration), these ports provide access to both the system console and the ILOM command-line interface (the ILOM service processor prompt), each on a separate channel (see [FIGURE 1-6](#)).

FIGURE 1-6 Separate System Console and Service Processor Channels



If the system console is configured to be accessible from the serial management and network management ports, when you connect through one of these ports you can access either the ILOM command-line interface or the system console. You can switch between the ILOM service processor prompt and the system console at any time, but you cannot access both at the same time from a single terminal window or shell tool.

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

- The # or % prompt indicates that you are at the system console and that the Solaris OS is running.
- The ok prompt indicates that you are at the system console and that the server is running under OpenBoot firmware control.
- The -> prompt indicates that you are at the service processor.

Note – If no text or prompt appears, it might be because no console messages were recently generated by the system. Pressing the terminal’s Enter or Return key should produce a prompt.

- To reach the system console from the service processor, type the `start /SP/console` command at the -> prompt.
- To reach the service processor from the system console, type the service processor escape sequence. By default, the escape sequence is #. (Pound-Period).

For more information about communicating with the service processor and system console, see:

- [“Communicating With the System” on page 1](#)
- [“ILOM -> Prompt” on page 21](#)
- [“OpenBoot ok Prompt” on page 23](#)
- [“Accessing the Service Processor” on page 8](#)
- The ILOM user’s guide and the ILOM supplement for your server

ILOM -> Prompt

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

Note – To view ILOM service processor boot messages, you must establish a connection using a serial device (such as an Alphanumeric terminal) to the serial management port *before* connecting AC power cords to the SPARC Enterprise T5120 or T5220 server.

You can log in to the ILOM service processor at any time, regardless of system power state, as long as AC power is connected to the system and you have a way of interacting with the system. You can also access the ILOM service processor prompt (->) from the OpenBoot ok prompt or from the Solaris # or % prompt, provided the system console is configured to be accessible through the serial management and network management ports.

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

Note – When you access the ILOM service processor for the first time, the default username is `root` and the default password is `changeme`.

For more information on navigating to or from the ILOM prompt, see the following:

[“Getting to the ok Prompt” on page 26](#)

[“Switching Between the Service Processor and the System Console” on page 19](#)

Access Through Multiple Controller Sessions

Up to five ILOM sessions can be active concurrently, one session through the serial management port and up to four SSH sessions through the network management port. Users of each of these sessions can issue commands at the -> prompt. However, only one user at a time can access the system console, and then only if the system console is configured to be accessible through the serial and network management ports. For more information, see:

[“Accessing the Service Processor” on page 8](#)

[“Activating the Network Management Port” on page 9](#)

Any additional ILOM sessions afford passive views of system console activity, until the active user of the system console logs out.

Reaching the -> Prompt

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

- If the system console is directed to the serial management and network management ports, you can type the ILOM escape sequence (`#.`).

- You can log in directly to the service processor from a device connected to the serial management port. See [“Accessing the Service Processor” on page 8](#).
 - You can log in directly to the service processor using a connection through the network management port. See [“Activating the Network Management Port” on page 9](#).
-

OpenBoot ok Prompt

A SPARC Enterprise T5120 or T5220 server with the Solaris OS installed operates at different *run levels*. For a full description of run levels, refer to the Solaris system administration documentation.

Most of the time, you operate a SPARC Enterprise T5120 or T5220 server at run level 2 or run level 3, which are multiuser states with access to full system and network resources. Occasionally, you might operate the system at run level 1, which is a single-user administrative state. However, the lowest operational state is run level 0. At this state, it is safe to turn off power to the system.

When a SPARC Enterprise T5120 or T5220 server is at run level 0, the ok prompt appears. This prompt indicates that the OpenBoot firmware is in control of the system.

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

- By default, before the operating system is installed the system comes up under OpenBoot firmware control.
- When the `auto-boot?` OpenBoot configuration variable is set to `false`, the system boots to the ok prompt.
- When the operating system is halted, the system transitions to run level 0 in an orderly way.
- When the operating system crashes, the system reverts to OpenBoot firmware control.
- During the boot process, when there is a serious hardware problem that prevents the operating system from running, the system reverts to OpenBoot firmware control.
- When a serious hardware problem develops while the system is running, the operating system transitions smoothly to run level 0.
- When you deliberately place the system under firmware control in order to execute firmware-based commands.

It is the last of these scenarios which most often concerns you as an administrator, since there will be times when you need to reach the `ok` prompt. Several ways to do this are outlined in [“Reaching the `ok` Prompt” on page 24](#). For detailed instructions, see [“Getting to the `ok` Prompt” on page 26](#).

Reaching the `ok` Prompt

There are several ways to reach the `ok` prompt, depending on the state of the system and the means by which you are accessing the system console. In decreasing order of desirability, these are:

- Graceful shutdown
- ILOM service processor `set /HOST send_break_action=break` and `start /SP/console` command pair
- Break key
- Manual system reset

A discussion of each method follows. For step-by-step instructions, see [“Getting to the `ok` Prompt” on page 26](#).

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

Graceful Shutdown

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

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

You can also perform a graceful system shutdown from the ILOM service processor command prompt using the `stop /SYS` command.

ILOM set /HOST send_break_action=break, start /SP/console Commands, or Break Key

Typing set /HOST send_break_action=break from the -> prompt forces a running SPARC Enterprise T5120 or T5220 server to drop to a menu:
c)ontinue, s)ync, r)eboot, h)alt?

from which you can type c to get OpenBoot firmware control.

When init 0 is used, the server drops to this menu:
r)eboot, o)k prompt, h)alt?

from which you can type o to get OpenBoot firmware control.

If the operating system is already halted, you can use the start /SP/console command instead of set /HOST send_break_action=break to reach the ok prompt.

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

When it is impossible or impractical to shut down the system gracefully, you can get to the ok prompt, if you have an alphanumeric terminal attached to the SPARC Enterprise T5120 or T5220 server, by pressing the Break key.

Note – These methods of reaching the ok prompt will only work if the system console has been redirected to the appropriate port. For details, see [“System Console OpenBoot Configuration Variable Settings” on page 28](#).

Manual System Reset



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

Use the ILOM service processor reset /SYS command, or start /SYS and stop /SYS commands, to reset the server. Reaching the ok prompt by performing a manual system reset or by power-cycling the system should be the method of last resort. Using these commands results in the loss of all system coherence and state

information. A manual system reset could corrupt the server's file systems, although the `fsck` command usually restores them. Use this method only when nothing else works.



Caution – Accessing the `ok` prompt suspends the Solaris OS.

When you access the `ok` prompt from a functioning SPARC Enterprise T5120 or T5220 server, you are suspending the Solaris OS and placing the system under firmware control. Any processes that were running under the operating system are also suspended, and *the state of such processes might not be recoverable*.

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

Getting to the `ok` Prompt

This procedure provides several ways of reaching the `ok` prompt. The methods are not equally desirable. For details about when to use each method, see [“OpenBoot `ok` Prompt” on page 23](#).



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

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

▼ To Obtain the ok Prompt

1. **Decide which method you need to use to reach the ok prompt.**
See “OpenBoot ok Prompt” on page 23 for details.
2. **Follow the appropriate instructions in TABLE 1-3.**

TABLE 1-3 Ways of Accessing the ok Prompt

Access Method	What to Do
Graceful shutdown of the Solaris OS	From a shell or command tool window, issue an appropriate command (for example, the <code>shutdown</code> or <code>init</code> command) as described in Solaris system administration documentation.
Break key	From an alphanumeric terminal configured to access the system console, press the Break key.
ILOM commands	From the <code>-></code> prompt, type the <code>set /HOST send_break_action=break</code> command. Then issue the <code>start /SP/console</code> command, provided the operating system software is not running and the server is already under OpenBoot firmware control.
Manual system reset	From the <code>-></code> prompt, type: <code>-> set /HOST/bootmode script="setenv auto-boot? false"</code> Press Enter. then type: <code>-> reset /SYS</code>

For More Information

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

System Console OpenBoot Configuration Variable Settings

On the SPARC Enterprise T5120 and T5220 servers, the system console is directed to the serial management and network management ports (SER MGT and NET MGT) by default. However, you can redirect the system console to a local graphics monitor, keyboard, and mouse. You can also redirect the system console back to the serial management and network management ports.

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

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

OpenBoot Configuration Variable Name	Setting for Sending System Console Output to:	
	Serial and Network Management Ports	Local Graphics Monitor/USB Keyboard and Mouse
output-device	virtual-console	screen
input-device	virtual-console	keyboard

Note – POST output will still be directed to the serial management port, as POST has no mechanism to direct its output to a graphics monitor.

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

The `->` prompt and POST messages are only available through the serial management port and network management port. The ILOM service processor `start /SP/console` command is ineffective when the system console is redirected to a local graphics monitor.

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

Managing RAS Features and System Firmware

This chapter describes how to manage reliability, availability, and serviceability (RAS) features and system firmware, including ILOM on the service processor, and automatic system recovery (ASR). In addition, this chapter describes how to unconfigure and reconfigure a device manually, and introduces multipathing software.

This chapter contains the following sections:

- [“ILOM and The Service Processor” on page 29](#)
- [“OpenBoot Emergency Procedures” on page 35](#)
- [“Automatic System Recovery” on page 37](#)
- [“Unconfiguring and Reconfiguring Devices” on page 43](#)
- [“Displaying System Fault Information” on page 44](#)
- [“Multipathing Software” on page 46](#)
- [“Storing FRU Information” on page 45](#)

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

ILOM and The Service Processor

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

After you log in to your ILOM account, the ILOM service processor command prompt (->) appears, and you can enter ILOM service processor commands. If the command you want to use has multiple options, you can either enter the options individually or grouped together, as shown in the following example.

```
-> stop -force -script /SYS  
-> start -script /SYS
```

Logging In To ILOM

All environmental monitoring and control is handled by ILOM on the ILOM service processor. The ILOM service processor command prompt (->) provides you with a way of interacting with ILOM. For more information about the -> prompt, see [“ILOM -> Prompt” on page 21](#).

For instructions on connecting to the ILOM service processor, see:

- [“Accessing the Service Processor” on page 8](#)
- [“Activating the Network Management Port” on page 9](#)

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

▼ To Log In To ILOM

1. At the ILOM login prompt, enter the login name and press Return.

The default login name is `root`.

```
Integrated Lights Out Manager 2.0
Please login: root
```

2. At the password prompt, enter the password and press Return to get to the -> prompt.

```
Please Enter password:
->
```

Note – The default user is `root` and the password is `changeme`. For more information, refer to the installation guide, the ILOM user's guide, and the ILOM supplement for your server.



Caution – In order to provide optimum system security, change the default system password during initial setup.

Using the ILOM service processor, you can monitor the system, turn the Locator LED on and off, or perform maintenance tasks on the ILOM service processor itself. For more information, refer to the ILOM user's guide and the ILOM supplement for your server.

▼ To View Environmental Information

1. Log in to the ILOM service processor.
2. Use the following command to display a snapshot of the server's environmental status.

```
show /SP/faultmgmt
```

Note – You do not need ILOM Administrator permissions to use this command.

Interpreting System LEDs

The behavior of LEDs on the SPARC Enterprise T5120 and T5220 servers conform to the American National Standards Institute (ANSI) Status Indicator Standard (SIS). These standard LED behaviors are described in [TABLE 2-1](#).

TABLE 2-1 Standard LED Behaviors and Values

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

The system LEDs have assigned meanings, described in [TABLE 2-2](#).

TABLE 2-2 System LED Behaviors with Assigned Meanings

Color	Behavior	Definition	Description
White	Off	Steady state	
	Fast blink	4-Hz repeating sequence, equal intervals On and Off.	This indicator helps you to locate a particular enclosure, board, or subsystem. Example: the Locator LED.
Blue	Off	Steady state	
	Steady On	Steady state	If blue is on, a service action can be performed on the applicable component with no adverse consequences. Example: the OK-to-Remove LED.
Yellow / Amber	Off	Steady state	
	Slow Blink	1-Hz repeating sequence, equal intervals On and Off.	This indicator signals new fault conditions. Service is required. Example: the Service Required LED.
	Steady On	Steady state	The amber indicator stays on until the service action is completed and the system returns to normal function.

TABLE 2-2 System LED Behaviors with Assigned Meanings (*Continued*)

Color	Behavior	Definition	Description
Green	Off	Steady state	
	Standby Blink	Repeating sequence consisting of a brief (0.1 sec.) ON flash followed by a long OFF period (2.9 sec.)	The system is running at a minimum level and is ready to be quickly revived to full function. Example: the System Activity LED.
	Steady On	Steady state	Status normal. System or component functioning with no service actions required.
	Slow Blink		A transitory (temporary) event is taking place for which direct proportional feedback is not needed or not feasible.

Controlling the Locator LED

You control the Locator LED from the -> prompt or with the locator button on the front of the chassis.

FIGURE 2-1 Locator Button (top button) on SPARC Enterprise T5120 Chassis





FIGURE 2-2 Locator Button (top button) on SPARC Enterprise T5220 Chassis



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

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

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

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

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

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

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

OpenBoot Emergency Procedures

The introduction of Universal Serial Bus (USB) keyboards with the newest systems has made it necessary to change some of the OpenBoot emergency procedures. Specifically, the `Stop-N`, `Stop-D`, and `Stop-F` commands that were available on systems with non-USB keyboards are not supported on systems that use USB keyboards, such as the Sun SPARC Enterprise T5120 or T5220 servers. If you are familiar with the earlier (non-USB) keyboard functionality, this section describes the analogous OpenBoot emergency procedures available in newer systems that use USB keyboards.

OpenBoot Emergency Procedures for SPARC Enterprise T5120 and T5220 Systems

The following sections describe how to perform the functions of the `Stop` commands on systems that use USB keyboards. These same functions are available through Integrated Lights Out Manager (ILOM) system controller software.

Stop-N Functionality

Stop-N functionality is not available. However, you can closely emulate the Stop-N functionality by completing the following steps, provided the system console is configured to be accessible using either the serial management port or the network management port.

▼ To Restore OpenBoot Configuration Defaults

1. Log in to the ILOM service processor.
2. Type the following commands:

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

Note – If you do not issue the `stop /SYS` and `start /SYS` commands or the `reset /SYS` command within 10 minutes, the host server ignores the `set/HOST/bootmode` commands.

You can issue the `show /HOST/bootmode` command without arguments to display the current setting

```
-> show /HOST/bootmode

/HOST/bootmode
  Targets:

  Properties:
    config = (none)
    expires = Tue Jan 19 03:14:07 2038
    script = (none)
    state = normal
```

3. To reset the system, type the following commands:

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

4. To view console output as the system boots with default OpenBoot configuration variables, switch to console mode.

```
-> set /SP/network pendingipdiscovery=dhcp
Set 'pendingipdiscovery' to 'dhcp'

-> set /SP/network commitpending=true
Set 'commitpending' to 'true'

->
```

5. To discard any customized IDPROM values and restore the default settings for all OpenBoot configuration variables, type:

```
-> set /SP reset_to_defaults=all
-> reset /SP
```

Stop-F Functionality

The Stop-F functionality is not available on systems with USB keyboards.

Stop-D Functionality

The Stop-D (Diags) key sequence is not supported on systems with USB keyboards. However, you can closely emulate the Stop-D functionality by setting the virtual keyswitch to `diag`, using the ILOM `set /SYS keyswitch_state=diag` command. For more information, refer to the ILOM guide and the ILOM supplement for your server.

Automatic System Recovery

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

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

Note – ASR is not activated until you enable it. See [“Enabling and Disabling Automatic System Recovery” on page 41.](#)

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

Auto-Boot Options

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

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

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

Note – The default setting for `auto-boot-on-error?` is `false`. The system will not attempt a degraded boot unless you change this setting to `true`. In addition, the system will not attempt a degraded boot in response to any fatal nonrecoverable error, even if degraded booting is enabled. For examples of fatal nonrecoverable errors, see [“Error Handling Summary” on page 38.](#)

Error Handling Summary

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

- If no errors are detected by POST or OpenBoot firmware, the system attempts to boot if `auto-boot?` is `true`.
- If only nonfatal errors are detected by POST or OpenBoot firmware, the system attempts to boot if `auto-boot?` is `true` and `auto-boot-on-error?` is `true`. Nonfatal errors include the following:

- SAS subsystem failure. In this case, a working alternate path to the boot disk is required. For more information, see [“Multipathing Software” on page 46](#).
- Ethernet interface failure.
- USB interface failure.
- Serial interface failure.
- PCI card failure.
- Memory failure. Given a failed DIMM, the firmware will unconfigure the entire logical bank associated with the failed module. Another nonfailing logical bank must be present in the system for the system to attempt a degraded boot.

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

- If a fatal error is detected by POST or OpenBoot firmware, the system does not boot regardless of the settings of `auto-boot?` or `auto-boot-on-error?`. Fatal nonrecoverable errors include the following:
 - Any CPU failed
 - All logical memory banks failed
 - Flash RAM cyclical redundancy check (CRC) failure
 - Critical field-replaceable unit (FRU) PROM configuration data failure
 - Critical system configuration card (SCC) read failure
 - Critical application-specific integrated circuit (ASIC) failure

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

Reset Scenarios

Three ILOM `/HOST/diag` configuration properties, `mode`, `level`, and `trigger`, control whether the system runs firmware diagnostics in response to system reset events.

The standard system reset protocol bypasses POST completely unless the virtual keyswitch or ILOM properties are set as follows:

TABLE 2-3 Virtual Keyswitch Setting for Reset Scenario

Keyswitch	Value
/SYS keyswitch_state	diag

If `keyswitch_state` is set to `diag`, the system can power itself on using preset values of diagnostic properties (`/HOST/diag level=max`, `/HOST/diag mode=max`, `/HOST/diag verbosity=max`) to provide thorough fault coverage. This option overrides the values of diagnostic properties that you might have set elsewhere.

TABLE 2-4 ILOM Property Settings for Reset Scenario

Property	Value
mode	normal or service
level	min or max
trigger	power-on-reset error-reset

The default settings for these properties are:

- `mode = normal`
- `level = min`
- `trigger = power-on-reset error-reset`

For instructions on automatic system recovery (ASR), see [“Enabling and Disabling Automatic System Recovery” on page 41](#).

Automatic System Recovery User Commands

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

- [“Unconfiguring and Reconfiguring Devices” on page 43](#)
- [“To Reconfigure a Device Manually” on page 44](#)
- [“Obtaining Automatic System Recovery Information” on page 42](#)

Enabling and Disabling Automatic System Recovery

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

▼ To Enable Automatic System Recovery

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

```
-> set /HOST/diag mode=normal  
-> set /HOST/diag level=max  
-> set /HOST/diag trigger=power-on-reset
```

2. At the `ok` prompt, type:

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

Note – For more information about OpenBoot configuration variables, refer to the service manual for your server.

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

```
ok reset-all
```

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

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

▼ To Disable Automatic System Recovery

1. At the `ok` prompt, type:

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

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

```
ok reset-all
```

The system permanently stores the parameter change.

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

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

Obtaining Automatic System Recovery Information

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

● At the `->` prompt, type:

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

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

For more information, see:

- [“Automatic System Recovery” on page 37](#)
- [“Enabling and Disabling Automatic System Recovery” on page 41](#)
- [“To Disable Automatic System Recovery” on page 41](#)
- [“Unconfiguring and Reconfiguring Devices” on page 43](#)
- [“To Reconfigure a Device Manually” on page 44](#)

Unconfiguring and Reconfiguring Devices

To support a degraded boot capability, the ILOM firmware provides the `set Device_Identifier component_state=disabled` command, which enables you to unconfigure system devices manually. This command “marks” the specified device as *disabled* by creating an entry in the ASR database. Any device marked `disabled`, whether manually or by the system’s firmware diagnostics, is removed from the system’s machine description prior to the hand-off to other layers of system firmware, such as OpenBoot PROM.

▼ To Unconfigure a Device Manually

- At the `->` prompt, type:

```
-> set Device_Identifier component_state=disabled
```

Where the *Device_Identifier* is one of the device identifiers from [TABLE 2-5](#)

Note – The device identifiers are case sensitive.

TABLE 2-5 Device Identifiers and Devices

Device Identifiers	Devices
/SYS/MB/CMPcpu_number/Pstrand_number	CPU Strand (Number: 0-63)
/SYS/MB/RISERRiser_number/PCIESlot_number	PCIe Slot (Number: 0-5)
/SYS/MB/RISERRiser_number/XAUIcard_number	XAUI card (Number: 0-1)
/SYS/MB/NETport_number	Ethernet ports (Number: 0-3)
/SYS/MB/PCIE	PCIe root complex
/SYS/MB/USBnumber	USB ports (Number: 0-1, located on rear of chassis)
/SYS/MB/CMP0/L2_BANKnumber	(Number: 0-3)
/SYS/DVD	DVD

TABLE 2-5 Device Identifiers and Devices (Continued)

Device Identifiers (Continued)	Devices (Continued)
/SYS/USBBD/USBnumber	USB ports (Number: 2-3, located on front of chassis)
/SYS/TTYA	DB9 Serial Port
/SYS/MB/CMP0/BRbranch_number/CHchannel_number/Ddimm_number	DIMMS

▼ To Reconfigure a Device Manually

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

```
-> set Device_Identifier component_state=enabled
```

where the *Device_Identifier* is any device identifier from [TABLE 2-5](#)

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

You can use the ILOM `set Device_Identifier component_state=enabled` command to reconfigure any device that you previously unconfigured with the `set Device_Identifier component_state=disabled` command.

Displaying System Fault Information

ILOM software lets you display current valid system faults.

▼ To Display Current Valid System Faults

- Type:

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

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

For example:

```
-> show /SP/faultmgmt

/SP/faultmgmt
Targets:
    0 (/SYS/PS1)

Properties:

Commands:
    cd
    show

->
```

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

▼ To Clear a Fault

- 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.

Storing FRU Information

▼ To Store Information in Available FRU PROMs

- At the `->` prompt type:

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

Multipathing Software

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

For the Sun SPARC Enterprise T5120 or T5220 Servers, three different types of multipathing software are available:

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

For More Information

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

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

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

Managing Disk Volumes

This chapter describes redundant array of independent disks (RAID) concepts, and how to configure and manage RAID disk volumes using the SPARC Enterprise T5120 or T5220 server's onboard serial attached SCSI (SAS) disk controller.

This chapter contains the following sections:

- [“Requirements” on page 47](#)
- [“Disk Volumes” on page 47](#)
- [“RAID Technology” on page 48](#)
- [“Hardware Raid Operations” on page 50](#)

Requirements

To configure and use RAID disk volumes on the SPARC Enterprise T5120 or T5220 server, you must install the appropriate patches. For the latest information on patches for the SPARC Enterprise T5120 or T5220 server, see the latest product notes for your system.

Installation procedures for patches are included in text README files that accompany the patches.

Disk Volumes

From the perspective of the on-board disk controller on the SPARC Enterprise T5120 and T5220 servers, *disk volumes* are logical disk devices comprising one or more complete physical disks.

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

The onboard disk controller of the SPARC Enterprise T5120 or T5220 server provides for the creation of as many as two hardware RAID volumes. The controller supports either two-disk RAID 1 (integrated mirror, or IM) volumes, or up to eight-disk RAID 0 (integrated stripe, or IS) volumes.

Note – Due to the volume initialization that occurs on the disk controller when a new volume is created, properties of the volume such as geometry and size are unknown. RAID volumes created using the hardware controller must be configured and labeled using `format(1M)` prior to use with the Solaris Operating System. See [“To Configure and Label a Hardware RAID Volume for Use in the Solaris Operating System” on page 57](#), or the `format(1M)` man page for further details.

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

RAID Technology

RAID technology enables the construction of a logical volume, made up of several physical disks, in order to provide data redundancy, increased performance, or both. The SPARC Enterprise T5120 or T5220 server’s on-board disk controller supports both RAID 0 and RAID 1 volumes.

This section describes the RAID configurations supported by the onboard disk controller:


- Integrated Stripe, or IS volumes (RAID 0)
- Integrated Mirror, or IM volumes (RAID 1)

Integrated Stripe Volumes (RAID 0)

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

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

FIGURE 3-1 Graphical Representation of Disk Striping




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

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

Integrated Mirror Volumes (RAID 1)

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

FIGURE 3-2 Graphical Representation of Disk Mirroring



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



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

Hardware Raid Operations

On the SPARC Enterprise T5120 or T5220 server, the SAS controller supports mirroring and striping using the Solaris OS `raidctl` utility.

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

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

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

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

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

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

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

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

▼ To Create a Hardware Mirrored Volume

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

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

See “Physical Disk Slot Numbers, Physical Device Names, and Logical Device Names for Non-RAID Disks” on page 51.

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

# raidctl				
RAID Volume	Volume Type	RAID Status	RAID Disk	Disk Status

c0t0d0	IM	OK	c0t0d0	OK
			c0t1d0	OK

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

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

The RAID status might be:

- OK – Indicating that the RAID volume is online and fully synchronized.
- RESYNCING – Indicating that the data between the primary and secondary member disks in an IM are still synchronizing.
- DEGRADED – Indicating that a member disk is failed or otherwise offline.
- FAILED – Indicating that volume should be deleted and reinitialized. This failure can occur when any member disk in an IS volume is lost, or when both disks are lost in an IM volume.

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

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

# raidctl				
RAID Volume	Volume Type	RAID Status	RAID Disk	Disk Status

c0t0d0	IM	DEGRADED	c0t0d0	OK
			c0t1d0	MISSING

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

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

2. Type the following command:

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

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

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

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

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

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

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

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

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

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

# raidctl				
RAID	Volume	RAID	RAID	Disk
Volume	Type	Status	Disk	Status

c0t0d0	IM	OK	c0t0d0	OK
			c0t1d0	OK

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

Under RAID 1 (disk mirroring), all data is duplicated on both drives. If a disk fails, replace it with a working drive and restore the mirror. For instructions, see [“To Perform a Mirrored Disk Hot-Plug Operation” on page 62](#).

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

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

Due to the volume initialization that occurs on the disk controller when a new volume is created, the volume must be configured and labeled using the `format(1M)` utility prior to use with the Solaris Operating System (see [“To Configure and Label a Hardware RAID Volume for Use in the Solaris Operating System” on page 57](#)). Because of this limitation, `raidctl(1M)` blocks the creation of a hardware RAID volume if any of the member disks currently have a file system mounted.

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

1. Determine which disk is the default boot device

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

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

ok devalias disk
disk                  /pci@0/pci@0/pci@2/scsi@0/disk@0,0
```

2. Type the `boot net -s` command

```
ok boot net -s
```

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

See [“To Create a Hardware Mirrored Volume” on page 51](#). For example:

```
# raidctl -c -r 1 c0t0d0 c0t1d0
Creating RAID volume c0t0d0 will destroy all data on member disks,
proceed
(yes/no)? yes
Volume  c0t0d0  created
#
```

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

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

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

▼ To Create a Hardware Striped Volume

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

See [“Disk Slot Numbers, Logical Device Names, and Physical Device Names” on page 51](#).

To verify the current RAID configuration, type:

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

The preceding example indicates that no RAID volume exists.

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

2. Type the following command:

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

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

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

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

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

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

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

```
# raidctl
```

RAID Volume	Volume Type	RAID Status	RAID Disk	Disk Status
c0t1d0	IS	OK	c0t1d0	OK
			c0t2d0	OK
			c0t3d0	OK

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

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

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

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

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

1. Start the `format` utility:

```
# format
```

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

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

In this example, c0t2d0 is the logical name of the volume.

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

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

For example:

```
format> type

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

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

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

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

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

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

```
format> disk

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

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

7. Exit the `format` utility.

The volume can now be used in the Solaris OS.

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

▼ To Delete a Hardware RAID Volume

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

See [“Disk Slot Numbers, Logical Device Names, and Physical Device Names”](#) on page 51.

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

# <code>raidctl</code>				
RAID	Volume	RAID	RAID	Disk
Volume	Type	Status	Disk	Status

c0t0d0	IM	OK	c0t0d0	OK
			c0t1d0	OK

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

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

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

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

For example:

```
# raidctl -d c0t0d0  
RAID Volume 'c0t0d0' deleted
```

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

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

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

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

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

```
# raidctl
```

For example:

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

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

▼ To Perform a Mirrored Disk Hot-Plug Operation

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

See [“Disk Slot Numbers, Logical Device Names, and Physical Device Names”](#) on page 51.

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

```
# raidctl
```

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

For example:

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

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

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

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

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

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

The RAID utility automatically restores the data to the disk.

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

```
# raidctl
```

For example:

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

This example indicates that RAID volume c0t1d0 is resynchronizing.

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

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

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

▼ To Perform a Nonmirrored Disk Hot-Plug Operation

1. Verify which hard drive corresponds with which logical device name and physical device name.
See [“Disk Slot Numbers, Logical Device Names, and Physical Device Names”](#) on page 51.

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

2. Type the following command:

```
# cfgadm -al
```

For example:

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

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

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

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

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

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

3. Remove the hard drive from the device tree.

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

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

For example:

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

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

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

Type the following command:

```
# cfgadm -al
```

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

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

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

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

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

7. Configure the new hard drive.

Type the following command:

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

For example:

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

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

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

Type the following command:

```
# cfgadm -al
```

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

Note that c0t3d0 is now listed as configured.

Logical Domains Software

SPARC Enterprise servers support the Logical Domains (LDoms) 1.0 software that is used to create and manage logical domains. The software comprises LDoms-enabling code in the Solaris 10 11/06 OS, LDoms-enabling in System Firmware 6.4, and the Logical Domains Manager, which is the command-line interface.

About Logical Domains Software

Logical Domains software enables you to allocate the system resources of your server (such as a boot environment, CPUs, memory, and I/O devices) into logical domains. By using a logical domains environment, you can increase resource usage, improve scaling, and gain greater control of security and isolation.

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

A logical domain is a discrete logical grouping with its own operating systems, resources, and identity within a single computer system. Applications software can run in logical domains. Each logical domain can be created, destroyed, reconfigured, and rebooted independently, without requiring a power cycle of the server. There are several roles that logical domains can perform as shown in the following table.

TABLE 4-1 Logical Domain Roles

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

Logical Domain Configurations

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

Logical Domains Software Requirements

The following software is required or recommended for Logical Domains on SPARC Enterprise T5120 or T5220 servers:

- (Required) Solaris 10 11/06 Operating System – Refer to the Solaris 10 Collection for more information.
- (Required) Solaris 10 11/06 OS patches:
 - 124921-02, which contains updates to the Logical Domains 1.0 drivers and utilities. Logical Domains networking will be broken without this patch.
 - 125043-01, which contains updates to the console (qcn) drivers. This patch depends on KU 118833-36, so if this is not already updated on your system, you will need to do so.

For further information about patches, see the product notes for your server.

- (Required) System firmware version 6.4.x – Refer to your server's documentation for more information.
- (Required) Logical Domains Manager 1.0.x software.
- (Recommended) Solaris Security Toolkit 4.2 software – Refer to the Solaris Security Toolkit 4.2 Administration Guide and Solaris Security Toolkit 4.2 Reference Manual for more information.

Note – During the boot process, domains that use virtual boot devices must wait for their service domains to come online first. This can prolong the boot process.

OpenBoot Configuration Variables

TABLE A-1 describes the OpenBoot firmware configuration variables stored in non-volatile memory on the system. The OpenBoot configuration variables are printed here in the order in which they appear when you issue the following command:
-> **show -o table -level all /SYS**

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

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

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

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

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

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

Index

Symbols

/etc/remote file, 14
 modifying, 15

-> commands

 reset /SYS, 36
 set /HOST/bootmode, 36
 set /SYS/LOCATE, 34, 35
 show /SYS/LOCATE, 35

-> prompt

 about, 21, 29
 accessing from network management port, 23
 accessing from serial management port, 23
 multiple sessions, 22
 system console escape sequence (#.), 22
 system console, switching between, 19
 ways to access, 22

A

Activity (disk drive LED), 66

alphanumeric terminal

 accessing system console from, 16
 setting baud rate, 17

auto-boot (OpenBoot configuration variable), 23, 38

automatic system recovery (ASR)

 about, 37
 commands, 40
 disabling, 41
 enabling, 41
 obtaining recovery information, 42

B

Break key (alphanumeric terminal), 27

C

cables, keyboard and mouse, 18

cfgadm (Solaris command), 64

cfgadm install_device (Solaris command),
 cautions against using, 64

cfgadm remove_device (Solaris command),
 cautions against using, 64

Cisco AS2511-RJ Terminal Server, connecting, 11

command prompts, explained, 21

communicating with the system

 about, 1
 options, table, 2

console configuration, connection alternatives
 explained, 7

D

default system console configuration, 5

device identifiers, listed, 43

device reconfiguration, manual, 44

device unconfiguration, manual, 43

device, disabling, 43

device, enabling, 44

disabling a device, 43

disk configuration

 RAID 0, 48
 RAID 1, 49

disk drives

 LEDs

- Activity, 66
- OK-to-Remove, 65
- logical device names, table, 51
- disk hot-plug
 - mirrored disk, 62
 - nonmirrored disk, 63
- disk slot number, reference, 51
- disk volumes
 - about, 47
 - deleting, 61
- dtterm (Solaris utility), 15
- Dynamic Host Configuration Protocol (DHCP)
 - client on network management port, 11

E

- enabling a device, 44
- environmental information, viewing, 31
- error handling, summary, 38
- escape sequence (#.), service processor, 22

F

- fsck (Solaris command), 26

G

- go (OpenBoot command), 26
- graceful system halt, 24, 27
- graphics monitor
 - accessing system console from, 17
 - connecting to PCI graphics card, 18
 - restrictions against using for initial setup, 17
 - restrictions against using to view POST output, 17

H

- halt, gracefully, advantages of, 24, 27
- hardware disk mirror
 - about, 50
 - hot-plug operation, 62
- hardware disk mirrored volume
 - checking the status of, 53
- hardware disk stripe
 - about, 48
- hardware disk striped volume
 - checking the status of, 56
- hardware disk striping, about, 48
- hot-plug operation

- nonmirrored disk drive, 63
- on hardware disk mirror, 62

I

- ILOM commands
 - reset /SYS, 25
 - show, 11
- ILOM, *See* Integrated Lights Out Manager (ILOM)
- init (Solaris command), 24, 27
- input-device (OpenBoot configuration variable), 18, 28
- Integrated Lights Out Manager (ILOM)
 - > prompt, *See* -> prompt
 - commands, *See* -> prompt
 - escape sequence (#.), 22
 - logging in, 30
 - multiple connections to, 22

K

- keyboard, attaching, 18

L

- LEDs
 - Activity (disk drive LED), 66
 - OK-to-Remove (disk drive LED), 65
- LEDs, Locator (system status LED), 33
- Locator (system status LED)
 - controlling from -> prompt, 34, 35
- Locator (system status LED), controlling, 33
- logging in to the Integrated Lights Out Manager (ILOM), 30
- logical device name (disk drive), reference, 51

M

- manual device reconfiguration, 44
- manual device unconfiguration, 43
- manual system reset, 25, 27
- monitor, attaching, 17
- multiple ILOM sessions, 22

N

- network management port (NET MGT)
 - activating, 9
- nonmirrored disk hot-plug operation, 63

O

- ok prompt
 - about, 23
 - accessing with Break key, 25
 - accessing with graceful system shutdown, 24
 - accessing with ILOM set /HOST
 - break_action=break command, 25
 - accessing with ILOM set /HOST
 - send_break_action=break command, 24
 - accessing with manual system reset, 24, 25
 - risks in using, 26
 - suspension of Solaris operating system, 26
 - ways to access, 24, 26
- OK-to-Remove (disk drive LED), 65
- OpenBoot commands
 - go, 26
 - probe-ide, 25
 - probe-scsi-all, 25
 - reset-all, 18
 - setenv, 18
- OpenBoot configuration variables
 - auto-boot, 23, 38
 - described, table, 71
 - input-device, 18, 28
 - output-device, 18, 28
 - system console settings, 28
- OpenBoot emergency procedures
 - performing, 35
 - USB keyboard commands, 35
- OpenBoot firmware
 - scenarios for control, 23
- operating system software, suspending, 26
- output-device (OpenBoot configuration variable), 18, 28

P

- parity, 17
- patch panel, terminal server connection, 12
- PCI graphics card
 - configuring to access system console, 17
 - connecting graphics monitor to, 18
 - frame buffers, 17
- physical device name (disk drive), 51
- probe-ide (OpenBoot command), 25
- probe-scsi-all (OpenBoot command), 25

R

- RAID (redundant array of independent disks), xiii, 47
- RAID 0 (striping), 48
- RAID 1 (mirroring), 49
- raidctl (Solaris command), 51 to 63
- reset
 - manual system, 25, 27
 - scenarios, 39
- reset /SYS (ILOM command), 25
- reset-all (OpenBoot command), 18
- run levels
 - explained, 23
 - ok prompt and, 23

S

- SER MGT, *See* serial management port
- serial management port (SER MGT)
 - acceptable console device connections, 5
 - as default communication port on initial startup, 2
 - configuration parameters, 8
 - default system console configuration, 5
 - using, 8
- set /HOST/bootmode (-> command), 36
- set /SYS/LOCATE (-> command), 35
- setenv (OpenBoot command), 18
- show (ILOM CMT command), 11
- shutdown (Solaris command), 24, 27
- Solaris commands
 - cfgadm, 64
 - cfgadm install_device, cautions against using, 64
 - cfgadm remove_device, cautions against using, 64
 - fsck, 26
 - init, 24, 27
 - raidctl, 51 to 63
 - shutdown, 24, 27
 - tip, 13, 15
 - uadmin, 24
 - uname, 16
 - uname -r, 15
- Stop-D (USB keyboard functionality), 37
- Stop-F (USB keyboard functionality), 37
- Stop-N (USB keyboard functionality), 36

- suspending the operating system software, 26
- system console
 - > prompt, switching between, 19
 - accessing with alphanumeric terminal, 16
 - accessing with graphics monitor, 17
 - accessing with terminal server, 2, 11
 - accessing with tip connection, 13
 - alphanumeric terminal connection, 2, 16
 - alternate configurations, 7
 - configuring local graphics monitor to access, 17
 - connection using graphics monitor, 8
 - default configuration explained, 2, 5
 - default connections, 4
 - defined, 1
 - Ethernet attachment through network
 - management port, 2
 - graphics monitor connection, 2, 8
 - multiple view sessions, 22
 - setting OpenBoot configuration variables for, 28
- system reset scenarios, 39
- system status LEDs
 - Locator, 34
- system status LEDs, Locator, 33

T

- terminal server
 - accessing system console from, 5, 11
 - connection through patch panel, 12
 - pinouts for crossover cable, 12
- tip (Solaris command), 15
- tip connection
 - accessing system console, 13
 - accessing terminal server, 13
- tip connection
 - accessing system console, 13

U

- uadmin (Solaris command), 24
- uname (Solaris command), 16
- uname -r (Solaris command), 15