



Sun Fire™ V215 and V245 Servers Product Notes

Sun Microsystems, Inc.
www.sun.com

Part No. 819-3040-18
July 2008, Revision A

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

Copyright 2008 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, U.S.A. All rights reserved.

Sun Microsystems, Inc. has intellectual property rights relating to technology that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more of the U.S. patents listed at <http://www.sun.com/patents> and one or more additional patents or pending patent applications in the U.S. and in other countries.

This document and the product to which it pertains are distributed under licenses restricting their use, copying, distribution, and decompilation. No part of the product or of this document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any.

Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

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, docs.sun.com, Sun Fire, SunSolve, SunVTS, OpenBoot, Sun BluePrints, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc., or its subsidiaries, in the U.S. and in other countries.

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

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.

U.S. Government Rights—Commercial use. Government users are subject to the Sun Microsystems, Inc. standard license agreement and applicable provisions of the FAR and its supplements.

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.

Copyright 2008 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, Californie 95054, Etats-Unis. Tous droits réservés.

Sun Microsystems, Inc. a les droits de propriété intellectuels relatants à la technologie qui est décrit dans ce document. En particulier, et sans la limitation, ces droits de propriété intellectuels peuvent inclure un ou plus des brevets américains énumérés à <http://www.sun.com/patents> et un ou les brevets plus supplémentaires ou les applications de brevet en attente dans les Etats-Unis et dans les autres pays.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a.

Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

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

Sun, Sun Microsystems, le logo Sun, docs.sun.com, Sun Fire, SunSolve, SunVTS, OpenBoot, Sun BluePrints, et Solaris sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc., et ses filiales, aux Etats-Unis et dans d'autres pays.

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

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 licenciées de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

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



Adobe PostScript

Contents

Preface vii

1. Sun Fire V215 and V245 Servers Product Notes 1

Shipping Kit Contents 1

Operating System and Required Patches 2

Software Updates 3

▼ To Update Your Sun Fire Server Using the Command-Line Interface 3

ALOM Updates 5

RJ-45 to DB-25 Pinouts 5

Hardware and Software Issues 6

Disabled DIMMs Cause Extra DIMM Entries in `prtdiag` Output (CR 6422484) 6

PCI Components Show Extra Entries in the `prtdiag` Output (CR 6422512) 6

NOTICE: `bge0: ddi_intr_get_supported_types()` in System Messages (CR 6428769) 7

PCI1 Slot `node-id` in FRU Maps to Two FRU Parents in Platform Tree (CR 6431886) 7

BGE Link Breaks IPMP Test Suite (CR 6436499) 7

Banner Displays Incorrect Amount of Memory (CR 6449210) 7

DVD Drive Warning (CR 6440674) 8

WARNING: `px1: spurious interrupt from ino 0x8` (CR 6470835) 8

USB Does Not Respond (CR 6465319)	8
Astec PSU 300-1945-02 Causes Spurious Errors During ALOM Boot and System POST (CR 6557962)	8
FIOB.T_AMB Status Erroneously Reported Critical Prior to poweron After an AC Power Cycle	9
Documentation Errata	9
<i>Sun Fire V215 and V245 Servers Getting Started Guide</i>	9
Maximum Ambient Operating Temperature	9
Power Calculator	10
General Power Figures	10
PCI Card Features	10
Hard Drive Status Indicators	12
<i>Sun Fire V215 and V245 Servers Administration Guide</i>	12
Firmware Update	12
Terminal Server Connection	12
input-device and output-device Values	13
2. Managing Disk Volumes	15
RAID Information	15
RAID Technology	16
RAID 0 – Disk Striping	16
RAID 1 – Disk Mirroring	17
RAID Volume and Drive Status	18
Physical Disk Slot Numbers, Physical Device Names, and Logical Device Names	18
RAID Operations With the Solaris OS Commands	19
Creating a Disk Stripe With the Solaris OS Commands	19
▼ To Create a Disk Stripe With the Solaris OS Commands	19
Creating a Disk Mirror With the Solaris OS Commands	20
▼ To Create a Disk Mirror With the Solaris OS Commands	20

Deleting a Disk Mirror With the Solaris OS Commands	21
▼ To Delete a Disk Mirror With the Solaris OS Commands	22
Updating the RAID Controller Firmware	22
RAID Operations With the OpenBoot PROM Commands	23
▼ To Recover a RAID Volume	23
Labeling a Volume	25
▼ To Label a Volume	25

Preface

This document includes late-breaking information about the Sun Fire™ V215 and V245 servers.

Using UNIX Commands

This document does not contain information about basic UNIX® commands and procedures such as shutting down the system, booting the system, and configuring devices. Refer to the following for this information:

- Software documentation that you received with your system
- Solaris™ Operating System documentation, which is at:

<http://docs.sun.com>

Shell Prompts

Shell	Prompt
C shell	<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. % You have mail.
AaBbCc123	What you type, when contrasted with on-screen computer output	% su 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

Application	Title	Part Number	Format	Location
General information	<i>Sun Fire V215 and V245 Servers Getting Started Guide</i>	819-3041	Printed, HTML, and PDF	Shipping kit and online
Installation	<i>Sun Fire V215 and V245 Servers Installation Guide</i>	819-3037	HTML and PDF	Online
Administration	<i>Sun Fire V215 and V245 Servers Administration Guide</i>	819-3036	HTML and PDF	Online
Service	<i>Sun Fire V215 and V245 Servers Service Manual</i>	819-3038	HTML and PDF	Online
Compliance and safety	<i>Sun Fire V215 and V245 Servers Safety and Compliance Manual</i>	819-3039	HTML and PDF	Online
Lights out management	<i>Sun Advanced Lights Out Manager (ALOM) 1.6 Administration Guide</i>	819-2445	PDF and HTML	Online

You can obtain copies of these documents at the following site:

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

Documentation, Support, and Training

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

Third-Party Web Sites

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

Sun Welcomes Your Comments

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

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

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

Sun Fire V215 and V245 Servers Product Notes, part number 819-3040-18

Sun Fire V215 and V245 Servers Product Notes

This document contains late-breaking news about the Sun Fire V215 and V245 servers.

Topics include:

- [“Shipping Kit Contents” on page 1](#)
- [“Operating System and Required Patches” on page 2](#)
- [“Software Updates” on page 3](#)
- [“ALOM Updates” on page 5](#)
- [“RJ-45 to DB-25 Pinouts” on page 5](#)
- [“Hardware and Software Issues” on page 6](#)
- [“Documentation Errata” on page 9](#)

Shipping Kit Contents

The shipping kit might contain different items than those described on the packing list. For example, to promote eco-responsibility, the kit might no longer contain the RJ-45 Ethernet cable, the antistatic wriststrap, or other ancillary items. Alternatively, serial adapters, fasteners, or other items not listed on the packing list might be included to enhance the customer experience. Contact Sun Microsystems, Inc., to purchase the items you need. These ancillary items also might be available at computer supply stores.

Operating System and Required Patches

If you need to reinstall the operating system for the Sun Fire V215 and V245 servers, you must use the minimum level of the Solaris Operating System (Solaris OS) and apply the minimum version of required patches in the following list:

- Solaris 9 9/05 HW OS

You must use the OS distribution with the HW designation. Required patches with minimum revision levels:

- 113886-40
- 113887-40
- 120812-13
- 118558-37
- 116561-15
- 123761-04
- 123763-03

- Solaris 10 6/06 OS

Required patches with minimum revision levels:

- 118833-22
- 118557-06
- 118777-06
- 119850-20
- 122027-06
- 122517-02
- 122523-02
- 122525-02
- 123334-02
- 123360-01
- 122363-01
- 123526-01
- 120068-03
- 116577-42
- 122032-04

Software Updates

After you have configured the preinstalled software on your server, you should register your system and obtain the latest Solaris OS updates to optimize the performance of your Sun Fire server.

▼ To Update Your Sun Fire Server Using the Command-Line Interface

1. Register your system with the Sun Update Connection:

- a. **Copy the** `/usr/lib/breg/data/RegistrationProfile.properties` **file to** `/tmp/myreg.profile`.
- b. **Edit the** `/tmp/myreg.profile` **file to add your user name.**

This file must be owned by `root` and have permissions set to `400`.

Note – To obtain patches, you must have a registered user name and password. To obtain a user name and password, go to <http://sun.com/mysun> to register.

c. Type:

```
# sconadm proxy -r /tmp/myreg.profile
```

d. Type:

```
# sconadm register -a -r /tmp/myreg.profile
```

You can disregard the warning messages about unknown host exception.

2. Update your system:

- a. **Set the** `patchpro.patchset` **variable:**

```
# smpatch set patchpro.patchset=sfv2x5
```

- b. **Display the list of patches:**

```
# smpatch analyze
```

c. Install the patches:

```
# smpatch update
```

Note – You might see a Problem installing patches message. Disregard this message. The patches will be installed properly after the shut down and boot of the system.

Alternatively, you can download individual patches from the SunSolveSM site:

<http://sunsolve.sun.com/>

ALOM Updates

The first time you log in to the Sun Advanced Lights Out Manager (ALOM) software you will be logged in as the `admin` user, and you will be prompted to set the `admin` password.

A new ALOM NVRAM variable, `if_connection`, has been added to the ALOM software. This variable controls the network connection type (Telnet or SSH). SSH is the default setting.

For more information, refer to the *Sun Advanced Lights Out Manager (ALOM) 1.6 Administration Guide*.

RJ-45 to DB-25 Pinouts

To connect an external Hayes-compatible modem (one that responds to *AT* commands) on either Port A or B, you must configure the modem cable in accordance to [TABLE 1-1](#).

TABLE 1-1 RJ-45 to DB-25 Pinout Connections

RJ-45 on System	Signal Direction	DB-25 on Modem
1 (RTS)	—>	4 (CTS)
2 (DTR)	—>	20 (DTR)
3 (TXD)	—>	2 (RXD)
4 (GND)	<—>	1 (GND)
5 (GND)	<—>	7 (SIG GND)
6 (RXD)	<—	3 (TXD)
7 (DSR)	<—	8 (DCD)
8 (CTS)	<—	5 (RTS)

Note – Port A is the console port and has no dial-in capabilities, but it can dial out. Port B can dial out or dial in. Port B also provides the required DCD signal for dial-in. The DCD signal informs the `ttymon` process that the modem is requesting a login session.

TABLE 1-1 gives the full end-to-end description of the pins on the modem DB-25 port that are connected to the pins on the server's RJ-45 TTYB port. You can use a straight RJ-45 cable, a twisted RJ-45 cable, or converters. The end result should be that pin X on the server's RJ-45 port must be connected to pin Y on the modem's DB-25 port.

For more information about the commands used to configure the modem, refer to Sun InfoDoc 47436 at:

<http://sunsolve.sun.com/>

Hardware and Software Issues

This section contains the title, number, summary, and possible workarounds for hardware and software issues that were found before the release of the product.

Disabled DIMMs Cause Extra DIMM Entries in `prtdiag` Output (CR 6422484)

When you execute the `asr-disable` command on the DIMMs, the `prtdiag` command prints out extra entries for disabled (empty) slots. Also, the PICL tree has extra memory-module class nodes for the disabled slots, which are not in the tree when the slots are enabled.

Workaround: Do not use the `asr-disable` command on empty DIMM slots, or ignore the extra entries.

PCI Components Show Extra Entries in the `prtdiag` Output (CR 6422512)

When you execute the `asr-disable` command on the PCI slots, the `prtdiag` command prints out extra entries in the I/O devices section.

Workaround: Ignore the extra entries.

NOTICE: bge0: ddi_intr_get_supported_types() in System Messages (CR 6428769)

The BGE driver displays debug messages to the system message log.

Workaround: Ignore these messages.

PCI1 Slot node-id in FRU Maps to Two FRU Parents in Platform Tree (CR 6431886)

The node-id in the FRU tree maps to two FRU parents in the platform tree instead of one. One of the node-id values is the correct PCI node for the card, and the other value is the PCI node for slot PCI3. All the other cards map to the correct FRU parents.

Workaround: None.

BGE Link Breaks IPMP Test Suite (CR 6436499)

When the BGE interface is plumbed, the RUNNING flag is not set immediately. The RUNNING flag is set after 3 to 4 seconds.

Workaround: Wait a few seconds after system startup before initializing the IPMP subsystem (in.mpathd).

Banner Displays Incorrect Amount of Memory (CR 6449210)

With eight 1-gigabyte DIMMs installed in the system, the banner displays 8192 Mbytes after the last memory bank is disabled. The memory value should be either 0 Mbytes or 2048 Mbytes. If you use asr-disable to disable a DIMM, the reported amount of memory will be incorrect.

Workaround: Ignore the reported amount of memory.

DVD Drive Warning (CR 6440674)

Timeout warnings might be displayed if the DVD media is damaged or is of poor quality.

Workaround: Ignore the timeout warnings, or remove the media from the DVD drive.

WARNING: px1: spurious interrupt from ino 0x8 (CR 6470835)

You might see this message if applications are running while you run the SunVTS™ serial test that writes output to the serial port. There is no data loss or data corruption.

Workaround: Ignore this message.

USB Does Not Respond (CR 6465319)

If you connect a USB keyboard to the front USB ports on the Sun Fire V245 server, OpenBoot™ PROM will display the following message during the USB probing sequence:

```
ERROR: USB device does not respond
```

The Solaris OS will reprobe USB devices at boot time, so the functionality of the keyboard is not affected at the Solaris OS prompt. However, the keyboard will not work at the ok prompt after the system is halted.

Workaround: Ignore this message, or use the rear USB ports to connect the keyboard.

Astec PSU 300-1945-02 Causes Spurious Errors During ALOM Boot and System POST (CR 6557962)

Some revisions of the Astec power supplies (300-1945-02) might cause spurious errors during the ALOM boot process and the POST. However, all errors and service LEDs return to normal after the Solaris OS boots.

Workaround: You can safely ignore the errors.

FIOB.T_AMB Status Erroneously Reported Critical Prior to poweron After an AC Power Cycle

At cold power on, the system controller boots, initializes the ALOM software, and displays the ALOM prompt. If the system does not automatically boot to the OpenBoot PROM prompt or the Solaris OS prompt (for instance, if the cover is off or if the `autoboot` variable is set to `false`) the EPIC remains at its initial power on state. In this state, the value for `FIOB.T_AMB` within the EPIC register set is zero. The value should be set to `0x80` to represent an invalid temperature. Because a valid value and valid status are returned to the ALOM software, the ALOM software reports the results as a critical environmental issue.

Workaround: This issue has been fixed in the 1.6.4 release of the ALOM software. To obtain this release, go to:

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

Documentation Errata

This section contains documentation errors that were found after the documents were released.

Sun Fire V215 and V245 Servers Getting Started Guide

This section contains the errors found in the *Sun Fire V215 and V245 Servers Getting Started Guide*.

Maximum Ambient Operating Temperature

The printed version of this guide incorrectly states that the maximum ambient operating temperature is 40°C. The maximum ambient operating temperature is 35°C.

Power Calculator

The printed version of this guide includes a URL to a power calculator. The power calculator is not available at this time.

General Power Figures

The printed version and the online version of this guide include general power figures. More specific power figures were obtained by using the functional mode in the SunVTS software. [TABLE 1-2](#) and [TABLE 1-3](#) show these figures.

TABLE 1-2 Sun Fire V215 Server

Parameter	Minimum Configuration	Maximum Configuration
Power in Watts	110V: 188.9 220V: 186.1	110V: 320.1 220V: 309.2
AC voltage	119	119
Current in AMPS	1.39 to 1.41	2.86 to 2.93

TABLE 1-3 Sun Fire V245 Server

Parameter	Minimum Configuration	Maximum Configuration
Power in Watts	110V: 205.8 220V: 202.7	110V: 362.6 220V: 354
AC voltage	119.4	119.8
Current in AMPS	1.44	3.43

PCI Card Features

The getting started guide includes a table that shows the Server Differentiating Features. The PCI numbers are incorrect. With the release of new optional riser cards in October 2007, the guidelines have been updated.

The following guidelines now apply for the Sun Fire V215 server:

TABLE 1-4 PCI Card Guidelines for the Sun Fire V215 Server

Riser	Type	Dimensions	Comments
1UXXL (left side)	PCI-X (slot 1)	Full height, half length [\]	133 MHz
1UEEL (left side)	PCI-E (slot 1 [*])	Full height, half length [†] Low profile, half length	Lane x8 Lane x8, slot x16
1UEER (right side)	PCI-E (slot 0)	Low profile, half length ^d	Lane x8, slot x16

* If you add a 1U or 2U PCI-E riser card, you must remove the PCI-X riser card. The removal results in the loss of available slots for PCI-X option cards.

\ Low-profile cards fit with appropriate full-height faceplate.

d 1UEER is factory installed in all systems.

The following guidelines now apply for the Sun Fire V245 server:

TABLE 1-5 PCI Card Guidelines for the Sun Fire V245 Server

Riser	Type	Dimensions	Comments
2UEXL (left side)	PCI-X (upper slot 3)	Full height, full length ^d	133 MHz
	PCI-X (lower slot 1)	Full height, half length [†]	133 MHz
2UEEL [*] (left side)	PCI-E (upper slot 3)	Full height, full length [‡]	Lane x8, slot x16
	PCI-E (lower slot 1)	Full height, half length [‡]	Lane x8, slot x16
2UEER [\] (right side)	PCI-E (upper slot 2)	Low profile, half length	Lane x4, slot x16
	PCI-E (lower slot 0)	Low profile, half length	Lane x4, slot x16

* If you add a 1U or 2U PCI-E riser card, you must remove the PCI-X riser card. The removal results in the loss of available slots for PCI-X option cards.

\ 2UEER is factory installed in all systems.

d Low-profile cards fit with appropriate full-height faceplate.

You can add only one riser card to the server. If you add the XVR-2500 graphics adapter, you must use the PCI-E riser card assembly that converts all of the PCI-X I/O slots to PCI-E slots. After installation, the I/O subsystem will be entirely PCI-E based (that is, no PCI-X slots will be available). The XVR-2500 graphics adapter and any PCI-X adapter, including the serial, asynchronous adapter (SAI/P), are mutually exclusive within the same system.

Hard Drive Status Indicators

In the *Sun Fire V215 and V245 Servers Getting Started Guide*, the behavior of the hard drive indicators is described as follows:

TABLE 1-6 Hard Drive Indicators

Indicator	LED Color	LED State	Component Status
Removal	Blue	On	Ready to remove
		Off	Not ready to remove
Usage	Amber	On	Being used
		Off	Not being used
Activity	Green	Flashing	Active SCSI transactions
		Off	No activity

The LED state for the Activity indicator should be On for no activity.

Sun Fire V215 and V245 Servers Administration Guide

This section contains the errors found in the *Sun Fire V215 and V245 Servers Administration Guide*.

Firmware Update


The correct command for Step 2 in “To Update the Firmware” on page 22 should be:

```
% telnet 129.xxx.xx.xx
```

Terminal Server Connection

In Figure 1-3 and Figure 1-4, the arrow pointing from the terminal server to the Sun Fire V215 and V245 server points to the TTYB connector. The arrow should point to the middle connector, which is the serial management port. [FIGURE 1-1](#) shows the proper routing of the cable in red.

FIGURE 1-1 Proper Terminal Server Cable Routing



input-device and output-device Values

Throughout the administration guide, the content states that `input-device` and `output-device` should be set to `virtual-console`. This value is not correct. The `input-device` and `output-device` variables should be set to `ttys`.

Managing Disk Volumes

This chapter contains brief descriptions of redundant array of independent disks (RAID) concepts, how to manage disk volumes, and how to configure hardware mirroring using the SAS controller.

This chapter contains the following sections:

- “RAID Technology” on page 16
- “Physical Disk Slot Numbers, Physical Device Names, and Logical Device Names” on page 18
- “RAID Operations With the Solaris OS Commands” on page 19
- “RAID Operations With the OpenBoot PROM Commands” on page 23
- “Labeling a Volume” on page 25

RAID Information

The Sun Fire V215 and V245 servers support RAID 0 and RAID 1 configurations. You should use the Solaris OS `raidctl(1M)` command to configure the volumes.

With Solaris 10 8/07 OS release, the `raidctl(1M)` command has been updated. The usage of the command has changed dramatically. Two potential areas of concern are:

- The previous version of the command required entering only `raidctl` to get volume and disk status. The command now requires the `-l` option and the *volume-name* argument. The change was prompted by the concern here a disk or volume in a nonoptimal or failed state might go unnoticed.
- There are options and features listed that are not supported on many platforms (for example, hot-spares) that might cause you to think they are available.

Before you use the command, read the updated `raidctl(1M)` man page.

For more information about RAID, the use of the `raidctl(1M)` command, and managing volumes, refer to the following resources:

- BigAdmin site: <http://www.sun.com/bigadmin/home/>
There are numerous articles on RAID that can be found by searching on RAID.
- Sun BluePrints™ site: <http://www.sun.com/blueprints/>
- *Solaris Volume Manager Administration Guide (819-2789)*, which can be found at the following site:
<http://docs.sun.com/>
- `raidctl(1M)` man page, which can be found in the Solaris 10 Reference Manual Collection at:
<http://docs.sun.com/>
- *VERITAS Volume Manager Administration Guide, Solaris (875-3529-10)*, which can be found at:
<http://www.sun.com/products-n-solutions/hardware/docs/>



Caution – If you remove the system board, you must run service-only OpenBoot PROM commands to recover the volumes. Contact your service representative for more information about these commands.

Note – The `raidctl(1M)` command does not support non-hot-pluggable devices. To service non-hot-pluggable components, you must shut down the server.

RAID Technology

There are several levels of RAID configurations that provide varying degrees of data availability with corresponding trade-offs in performance and cost. This section describes the RAID configurations that are supported on the Sun Fire V215 and V245 servers.

RAID 0 – Disk Striping

Disk striping (RAID 0) is a technique for increasing system throughput by using several disk drives in parallel. In nonstriped disks the operating system writes a single block to a single disk. In a striped arrangement, each block is divided and portions of the data are written to different disks simultaneously. No data is shared.

Data is written across all disks (concatenated) in exactly one location. Write operations are committed in a round-robin fashion across all member disks, so performance is increased in some workload scenarios. Read operations are satisfied at one location (performs as a single disk).

When you create a RAID volume, all data contained on the drive is lost. If a single disk fails, data is lost. If a RAID 0 volume is disabled (broken), data is lost.

System performance using RAID 0 will be better than using RAID 1, but the possibility of data loss is greater because there is no way to retrieve or reconstruct data stored on a failed disk drive.

RAID 1 – Disk Mirroring

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 (primary and secondary configuration). You can configure hardware disk mirroring using the SAS controller. This configuration provides higher performance than conventional software mirroring using volume management software. If you mirror the boot drive, you must create the RAID volume before you install the operating system.

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, the operating system reads from whichever disk is more readily accessible at the moment, which can result in enhanced performance for read operations.

If a single disk fails, the data is still available, but the volume will be in the degraded mode. You can replace and resynchronize the failed disk. If the volume is disabled (that is, the mirror is broken), you must relabel the drives, which causes all of the data to be lost.

On the Sun Fire V215 and V245 servers, the SAS controller supports internal hardware disk mirroring using the Solaris OS `raidctl` utility. A hardware disk mirror created under the `raidctl` utility behaves slightly differently than a disk mirror created using volume management software. Under a software mirror, each device has its own entry in the virtual device tree, and read-write operations are performed to both virtual devices. Under hardware disk mirroring, only one device (the *master*) appears in the device tree. The mirrored device (the *slave*) is invisible to the OS and is accessed only by the SAS controller.



Caution – Creating or restoring a disk mirror destroys all data previously stored on the disk drive.

RAID Volume and Drive Status

After you create a RAID volume or during the normal operation of the system, the volume can have one of the following status conditions:

- OK – The volume is optimal.
- DEGRADED – The volume has a failed drive (mirror) or is otherwise functionally minimized.
- RESYNCING – The volume is synchronizing (mirror).
- FAILED – The volume is offline.

At any time during the normal operation of the system, any one or more of the drives can have one of the following status conditions:

- OK – Disk is operating normally.
- MISSING – Disk is removed or otherwise offline.
- FAILED – Disk is in a failed state and needs service.

Physical Disk Slot Numbers, Physical Device Names, and Logical Device Names

To perform a RAID procedure, you must know the physical or logical device name for the drive that you want to use. You can use [TABLE 2-1](#) to associate internal disk slot numbers with the logical and physical device names for each hard disk drive.

TABLE 2-1 Disk Slot Numbers, Logical Device Names, and Physical Device Names

Disk Slot Number	Logical Device Name*	Physical Device Name
Slot 0	c0t0d0	/pci@1e,600000/pci@0/pci@a/pci@0/pci@8/scsi@1/sd@0,0
Slot 1	c0t1d0	/pci@1e,600000/pci@0/pci@a/pci@0/pci@8/scsi@1/sd@1,0
Slot 2	c0t2d0	/pci@1e,600000/pci@0/pci@a/pci@0/pci@8/scsi@1/sd@2,0
Slot 3	c0t3d0	/pci@1e,600000/pci@0/pci@a/pci@0/pci@8/scsi@1/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.

RAID Operations With the Solaris OS Commands

This section contains instructions on how to create RAID volumes using the Solaris OS commands.

Creating a Disk Stripe With the Solaris OS Commands

When you create a disk stripe on two or more drives, the data on all of the drives will be lost. In addition, no file systems can be mounted on any of these drives.

The RAID volume has the ability to use the capacity of all the disks that are combined. However, the smallest common size is used.

▼ To Create a Disk Stripe With the Solaris OS Commands

1. **Verify that a hardware disk stripe does not already exist by typing:**

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

The example indicates that no RAID volume exists. 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 [disk3]
```

For example:

```
# raidctl -c -r 0 c0t1d0 c0t2d0 c0t3d0
```

The RAID volume `c0t1d0` is created.

Before you can use the volume, you must label it in the Solaris OS by using the `format(1M)` command. For instructions on labeling the volume, see [“Labeling a Volume” on page 25](#).

Creating a Disk Mirror With the Solaris OS Commands

When you create a disk mirror, all of the data on all of the drives is lost. The RAID volume will have the capacity of one drive. No file systems can be mounted on either drive. The secondary drive must be as large as the primary drive.

▼ To Create a Disk Mirror With the Solaris OS Commands

1. Verify that a hardware disk mirror does not already exist by typing:

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

The example indicates that no RAID volume exists. 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 master slave
```

For example:

```
# raidctl -c c0t0d0 c0t1d0
```

RAID volume `c0t0d0` is created. When you create a RAID mirror, the slave drive (in this case, `c0t1d0`) disappears from the Solaris OS device tree.

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

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

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

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

```
# raidctl -l
  RAID      RAID      RAID      Disk
  Volume    Status    Disk      Status
  -----
  c0t0d0    OK        c0t0d0    OK
                               c0t1d0    OK
```

Under RAID 1 (disk mirroring), all data is duplicated on both drives. If a drive fails, replace that drive with a working drive and restore the mirror.

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

Before you can use the volume, you must label it in the Solaris OS by using the `format(1M)` command. For instructions on labeling the volume, see [“Labeling a Volume” on page 25](#).

Deleting a Disk Mirror With the Solaris OS Commands

Before you perform the procedure in this section, verify which disk drive corresponds with which logical device name and physical device name. For more information, see [“Physical Disk Slot Numbers, Physical Device Names, and Logical Device Names” on page 18](#).

▼ To Delete a Disk Mirror With the Solaris OS Commands

1. Determine the name of the mirrored volume:

```
# raidctl -l
  RAID      RAID      RAID      Disk
  Volume    Status    Disk      Status
  -----
  c0t0d0    OK        c0t0d0    OK
                    c0t1d0    OK
```

In this example, the mirrored volume is c0t0d0.

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

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

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

For example:

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

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

```
# raidctl -l
No RAID volumes found
```

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

After you delete the volume, you must relabel the drive in the Solaris OS by using the `format(1M)` command.

Updating the RAID Controller Firmware

The RAID controller firmware is part of the OpenBoot PROM software patch. If you need to update the firmware for the RAID controller, download and install the latest OpenBoot PROM software patch.

RAID Operations With the OpenBoot PROM Commands

If you have existing RAID volumes and have to replace a component, such as the system board, you must perform the procedure in this section so that the Solaris OS can see the volume.



Caution – The OpenBoot command, `activate-volumes`, should be run only by a qualified service representative.

In the following procedure, the `show-disks` command is used to determine the controller path. You can also use the `probe-scsi-all` command to find the path.

▼ To Recover a RAID Volume

1. Set the `auto-boot?` variable to `false`.

```
ok setenv auto-boot? false
auto-boot? =                false
```

2. Set the `fcode-debug?` variable to `true`.

```
ok setenv fcode-debug? true
fcode-debug? =              true
```

3. Reset the system.

```
ok reset-all
```

4. Type the `show-disks` command to determine the path of the controller.

```
ok show-disks
a) /pci@1e,600000/pci@0/pci@a/pci@0/pci@8/scsi@1/disk
b) /pci@1e,600000/pci@0/pci@1/pci@0/ide@1f/cdrom
c) /pci@1e,600000/pci@0/pci@1/pci@0/ide@1f/disk
q) NO SELECTION
```

5. Enter the appropriate selection, type Control-Y to insert the path into the command line, then press Enter.

```
Enter Selection, q to quit: a
/pci@1e,600000/pci@0/pci@a/pci@0/pci@8/scsi@1/disk has been
selected.
Type ^Y ( Control-Y ) to insert it in the command line.
e.g. ok nvalias mydev ^Y
      for creating devalias mydev for
/pci@1e,600000/pci@0/pci@a/pci@0/pci@8/scsi@1/disk

ok select /pci@1e,600000/pci@0/pci@a/pci@0/pci@8/scsi@1
```

Note – The /disk portion of the path was left off the path because it is the controller you want to select.

6. Type the `show-volumes` command to determine the inactive volume:

```
ok show-volumes
```

7. Activate the volume using the volume number:

```
ok volume-number activate-volume
```

8. Unselect the device:

```
ok unselect-dev
```

9. Reset the OpenBoot PROM variables, and reset the system after you have finished managing the RAID volumes.

```
ok setenv auto-boot? true
auto-boot? = true
ok setenv fcode-debug? false
fcode-debug? = false
ok reset-all
```

If the `auto-boot?` variable was set to `true`, you might see the following message:

```
Base SAS World Wide ID is 0!
This must be fixed immediately using set-sas-wwid
```

If you see this message, set the `auto-boot?` variable to `false`, and issue the `reset-all` command.

Labeling a Volume

After you create a RAID volume by using the `raidctl(1M)` command or by using the OpenBoot PROM commands, you must label the volume correctly before the volume can be used.

▼ To Label a Volume

The Solaris OS must be running to perform this procedure. If you have created the volume by using the OpenBoot PROM commands, start the Solaris OS and log in as the superuser.

1. **Run the `format(1M)` command, and choose the newly created volume from the list of drives.**

You might see messages about the volume having a corrupt label or the block count being incorrect when booting or running the `format` command. This command needs a correct label.

2. **Chose `type` from the Format menu.**

This action displays a list of available drive types.

3. **Chose `Auto configure` from the list.**

This action defines the volume with the correct type.

4. **Write the new label to the volume using the `label` command from the `format>` prompt.**

The system asks you if you are ready to label the disk.

5. **Type `yes`.**

After you perform these steps, the drive is ready for use. You can now partition the drive as needed and use the volume within the Solaris OS.

Note – After you delete a RAID volume, the drives that were previously part of the RAID volume should be relabeled for proper use.
