

Sun Blade 6000 Disk Module

Administration Guide



Part No. 820-4922-16
June 2010, Revision A

Copyright © 2009, 2010 Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related software documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle America, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications which may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. UNIX is a registered trademark licensed through X/Open Company, Ltd.

This software or hardware and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

Copyright © 2009, 2010, Oracle et/ou ses affiliés. Tous droits réservés.

Ce logiciel et la documentation qui l'accompagne sont protégés par les lois sur la propriété intellectuelle. Ils sont concédés sous licence et soumis à des restrictions d'utilisation et de divulgation. Sauf disposition de votre contrat de licence ou de la loi, vous ne pouvez pas copier, reproduire, traduire, diffuser, modifier, breveter, transmettre, distribuer, exposer, exécuter, publier ou afficher le logiciel, même partiellement, sous quelque forme et par quelque procédé que ce soit. Par ailleurs, il est interdit de procéder à toute ingénierie inverse du logiciel, de le désassembler ou de le décompiler, excepté à des fins d'interopérabilité avec des logiciels tiers ou tel que prescrit par la loi.

Les informations fournies dans ce document sont susceptibles de modification sans préavis. Par ailleurs, Oracle Corporation ne garantit pas qu'elles soient exemptes d'erreurs et vous invite, le cas échéant, à lui en faire part par écrit.

Si ce logiciel, ou la documentation qui l'accompagne, est concédé sous licence au Gouvernement des Etats-Unis, ou à toute entité qui délivre la licence de ce logiciel ou l'utilise pour le compte du Gouvernement des Etats-Unis, la notice suivante s'applique :

U.S. GOVERNMENT RIGHTS. Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle America, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

Ce logiciel ou matériel a été développé pour un usage général dans le cadre d'applications de gestion des informations. Ce logiciel ou matériel n'est pas conçu ni n'est destiné à être utilisé dans des applications à risque, notamment dans des applications pouvant causer des dommages corporels. Si vous utilisez ce logiciel ou matériel dans le cadre d'applications dangereuses, il est de votre responsabilité de prendre toutes les mesures de secours, de sauvegarde, de redondance et autres mesures nécessaires à son utilisation dans des conditions optimales de sécurité. Oracle Corporation et ses affiliés déclinent toute responsabilité quant aux dommages causés par l'utilisation de ce logiciel ou matériel pour ce type d'applications.

Oracle et Java sont des marques déposées d'Oracle Corporation et/ou de ses affiliés. Tout autre nom mentionné peut correspondre à des marques appartenant à d'autres propriétaires qu'Oracle.

AMD, Opteron, le logo AMD et le logo AMD Opteron sont des marques ou des marques déposées d'Advanced Micro Devices. Intel et Intel Xeon sont des marques ou des marques déposées d'Intel Corporation. Toutes les marques SPARC sont utilisées sous licence et sont des marques ou des marques déposées de SPARC International, Inc. UNIX est une marque déposée concédée sous licence par X/Open Company, Ltd.

Ce logiciel ou matériel et la documentation qui l'accompagne peuvent fournir des informations ou des liens donnant accès à des contenus, des produits et des services émanant de tiers. Oracle Corporation et ses affiliés déclinent toute responsabilité ou garantie expresse quant aux contenus, produits ou services émanant de tiers. En aucun cas, Oracle Corporation et ses affiliés ne sauraient être tenus pour responsables des pertes subies, des coûts occasionnés ou des dommages causés par l'accès à des contenus, produits ou services tiers, ou à leur utilisation.



Adobe PostScript

Contents

Using This Documentation	xi
1. Overview of Using Sun Blade 6000 Disk Modules With Server Modules	1
Terminology	1
Overview of the Sun Blade 6000 Disk Module	2
Sun Blade 6000 Disk Module Front Panel	4
The Server Module SAS Host Bus Adapter	6
Server Modules and Disk Modules in a Chassis	7
2. Inserting, Moving and Replacing the Sun Blade 6000 Disk Module	9
Inserting and Controlling a Disk Module	9
Moving Disk Blades	10
Moving Pairs of Server Blades and Disk Blades in the Same or a Different Chassis	10
Replacing Disk Blades	10
3. System Requirements	13
Supported Servers, Firmware and Operating Systems	14
Meeting the Minimum Requirements for Disk Blade Support	15
Adding Patches to Solaris	16
Enclosure Driver Required for Windows	16

Important OS and Server Considerations	16
Supported Disk Drives	17
Supported RAID Controllers and Host Bus Adapters (HBA)	18
Supported Network Express Modules	21
Rules for Mixing NEMs	21
Supported Blade Chassis	22
Drivers, Firmware and Software	22
Minimum Operational Requirements	22
Oracle Common Array Manager (CAM) Software	23
Oracle's Sun Installation Assistant (SIA) Software	23
lsiutil Software	23
Where to Get Drivers, Firmware and Software	24
Obtaining the Latest Information	24
4. Upgrading Systems	25
Replacing Older X6220 Server Blades	26
▼ To Verify the Correct 1068E Version	26
Upgrading the Chassis ILOM	27
Overview of Upgrading Firmware on x86 Server Blades	27
Downloading SIA for Each x86 Server Blade	28
Checking SP Firmware and System BIOS Versions on x86 Server Blades	28
Upgrading SP Firmware and System BIOS on x86 Server Blades	31
Checking Host Bus Adapter Firmware Versions on x86 Server Blades	31
▼ To Check LSI Adapter Firmware Version by Rebooting	31
▼ To Check Adaptec Adapter Firmware Version Without Rebooting	33
▼ To Check Adaptec Adapter Firmware Version by Rebooting	33
Upgrading Host Bus Adapter Firmware on x86 Server Blades	34
Updating Operating System Drivers on x86 Server Blades	34
Linux Drivers	36

Solaris Drivers	36
Windows Drivers	37
▼ To Upgrade Your Blade Server Windows Driver	37
Updating the Sun Blade 6000 Enclosure Device Driver in Windows	40
▼ To Install the Driver	40
Upgrading Firmware on SPARC Server Blades	41
▼ To Upgrade Your SPARC Blade Server System Firmware	41
▼ To Upgrade Your SPARC Blade Server Disk Controller Firmware	43
Upgrading Disk Blade and SAS-NEM Firmware	44
Updating Solaris Network Installation Servers	45
For SPARC Servers	45
▼ To Have the Patches Added Automatically	45
For x86 Servers	45
▼ To Update x86 Solaris Network Installation Servers	46
5. Installing, Removing or Replacing Disks	47
Moving or Replacing Disks	47
Disk Drive Compatibility Rules	48
Multipathing	48
Replacing a Hard Disk Drive	49
▼ To Remove a Hard Disk Drive.	50
Installing a Disk Drive	52
6. Multipathing and RAID	53
Multipathing Overview	53
Multipathing Using LSI Host Bus Adapters	54
One SAS-NEM (Single Path)	54
Two SAS-NEMs (Dual Path)	55
Multipathing Using Adaptec Host Bus Adapters	57

Multipathing at the OS Level	57
Hardware RAID Hides Dual Paths From the OS	58
Multipathing With RHEL 5.x Linux and LSI Host Bus Adapters	58
Installing the OS on Multipath Disks	59
▼ To Configure Multipathing After Installation	59
Multipathing With the Solaris OS and LSI Host Bus Adapters	59
▼ To Install Multipathing for the Solaris OS	60
Multipathing With Windows Server 2008 and LSI Host Bus Adapters	62
▼ To Install Multipathing for Windows Server 2008	62
▼ To Select Disks for Multipathing	62
▼ To Set Load Balancing Property to Failover Only	63
▼ To View a Single Instance of MPIO Disk Drives	64
Failover	65
Multipathing with VMware ESX 3.5 and LSI Host Bus Adapters	65
▼ To View Status and Configure Multipath Storage Devices	66
7. How LSI Host Bus Adapters Assign Target IDs to Expanders, Disks, and RAID Volumes	69
About Target IDs With LSI Controllers	69
Allowed Range of Target IDs for LSI Controllers	70
Target ID Mapping With LSI Controllers	70
Assigning Volume IDs With LSI Controllers	71
Collision With Rebuilt Volumes	72
8. How Adaptec Host Bus Adapters Assign Target IDs to Expanders, Disks, and RAID Volumes	73
About Adaptec Host Bus Adapters	73
Target ID Mapping With Adaptec Host Bus Adapters	74
Assigning Volume IDs With Adaptec Host Bus Adapters	77

9.	Setting Up Hardware RAID Volumes With LSI and Adaptec Host Bus Adapters	81
	RAID Options	81
	Setting Up Hardware RAID With LSI Host Bus Adapters	83
	Deciding Whether to Configure RAID Before or After OS Installation	84
	Configuring RAID for the Windows and Linux OS After the OS Is Installed	84
	Configuring RAID for the Solaris OS After the OS Is Installed	84
	▼ To Configure RAID for SPARC Systems Before Installing the OS	85
	▼ To Configure LSI RAID for x86 Servers With Any OS From the BIOS Before the OS Is Installed	85
	Setting Up Hardware RAID With the Adaptec Host Bus Adapters	91
	Configuring Adaptec RAID After the OS Is Installed	91
	Configuring Adaptec RAID for x86 Servers With Any OS from the BIOS	91
	▼ To Create a RAID Array	92
10.	ILOM For the Sun Blade 6000 Disk Module	97
	ILOM on the CMM	97
	ILOM 2.0 Proxy CLI Program For SAS-NEMs and Disk Modules	98
	▼ To Start the ILOM 2.0 Proxy CLI Program	99
	Navigation Using the CMM ILOM CLI	99
	CMM ILOM Management and Monitoring for SAS-NEMs and Disk Modules	101
	Upgrading CMM ILOM Firmware	102
	▼ To upgrade ILOM Firmware Using the CLI	102
	▼ To upgrade ILOM Firmware Using the ILOM Web Interface	102
11.	Oracle's Sun Storage Common Array Manager (CAM)	105
	About Enclosure Management	105
	Assigning Storage Resources	106
	CAM	106

	CAM Agent	106
	Problem Resolution	107
	Obtaining CAM Software	107
	Using CAM With Disk Blades and SAS-NEMs	108
	Monitoring Component Health	108
	Upgrading Expander Firmware Using CAM	109
	▼ To Upgrade Expander Firmware	110
A.	Using the SIA Application to Upgrade x86 Server Blade and Host Bus Adapter Firmware	117
	Using SIA to Upgrade x86 Systems and HBAs	117
	Local and Remote Media Options	117
	▼ To Perform a Firmware upgrade Task With SIA Using Local or Remote Media	118
B.	Using the lsiutil Software	121
	Where to Obtain the lsiutil Software	121
	Installing lsiutil	122
	▼ To Use lsiutil From the Command Line	122
	▼ To Use the Interactive lsiutil Menus	122
	Why Save LSI Host Bus Adapter Persistent Mappings?	126
	When to Save LSI Host Bus Adapter Persistent Mappings	126
	Saving and Restoring a Persistence Map	127
	▼ To Save a Snapshot of Your Host Bus Adapter Persistent Mappings	127
	▼ To Restore a Snapshot of Your Host Bus Adapter Persistent Mappings	131
C.	Upgrading the Disk Module to Allow a SAS-2 Compatible Server Blade in the Chassis	135
	Server with a SAS-2 REM Might Hang When Inserted into a Chassis with SAS-1 Devices	135

Using This Documentation

This document contains procedures for connecting, managing and maintaining disk drives in Oracle's Sun Blade 6000 Disk Module.

Product Information

For product information and downloads for the Sun Blade 6000 Disk Module, go to the following web site:

(<http://www.oracle.com/goto/blades>)

Related Documentation

For a description of the document set for the Sun Blade 6000 Disk Module, see the *Where To Find Documentation* sheet that is packed with your system and also posted at the product's documentation site:

(<http://docs.sun.com/app/docs/coll/blade6000dskmod>)

Translated versions of some of these documents are available at this web site in French, Simplified Chinese, and Japanese. English documentation is revised more frequently and might be more up-to-date than the translated documentation.

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.

Documentation Feedback

Submit comments about this document by clicking the Feedback[+] link at (<http://docs.sun.com>). Include the title and part number of your document with your feedback:

Sun Blade 6000 Disk Module Administration Guide, part number 820-4922-16.

Overview of Using Sun Blade 6000 Disk Modules With Server Modules

This chapter contains the following topics:

- “Terminology” on page 1
- “Overview of the Sun Blade 6000 Disk Module” on page 2
 - “Sun Blade 6000 Disk Module Front Panel” on page 4
- “The Server Module SAS Host Bus Adapter” on page 6
- “Server Modules and Disk Modules in a Chassis” on page 7

Terminology

The following terminology is used in this document.

disk module or disk blade	The Sun Blade 6000 Disk Module (blade), the subject of this document The terms “disk module” and “disk blade” are used interchangeably.
server module or server blade	Any server module (blade) that will interoperate with the disk module. The terms “server module” and “server blade” are used interchangeably.
chassis	The Sun Blade 6000 Modular System blade enclosure. A chassis in which server and disk blades are installed
CMM	Chassis Monitoring Module. An Integrated Lights Out Manager component of the Sun Blade 6000 Modular System used to access and manage blades in a chassis.
NEM	Any Network Express Module that plugs into a Sun Blade 6000 chassis.

NEM 0, NEM 1	The slots for Network Express Modules on the rear of the chassis. NEM 0 is the lower slot and NEM 1 is the upper slot.
SAS-NEM	A Network Express Module that supports SAS interconnectivity. For example, the Sun Blade Multi-Fabric Network Express Module (abbreviated Multi-Fabric NEM) or the Sun Blade 10GbE Multi-Fabric Network Express Module (abbreviated 10GbE Multi-Fabric NEM).
REM	A RAID Expansion Module. Also referred to as a Host Bus Adapter (HBA). Supports the creation of RAID volumes on disks in the server and disk blades.

Note – From time to time, this document refers to SAS-NEMs. SAS-NEM is a general category: it includes Sun Blade 6000 Multi-Fabric NEMs, as well as other NEMs with SAS connectivity such as the Sun Blade 6000 10GbE Multi-Fabric NEM. From the standpoint of connecting server blades with disk blades, all SAS-NEMs provide the same functionality.

Overview of the Sun Blade 6000 Disk Module

The Sun Blade 6000 Disk Module is a SAS-1 (up to 3 Gb/s data transfer speed) disk module for the Sun Blade 6000 Modular System.

As a disk module, the Sun Blade 6000 Disk Module does not contain a CPU or service processor. It does, however, contain SAS expanders, which are switching circuits that can connect disks in complex patterns.

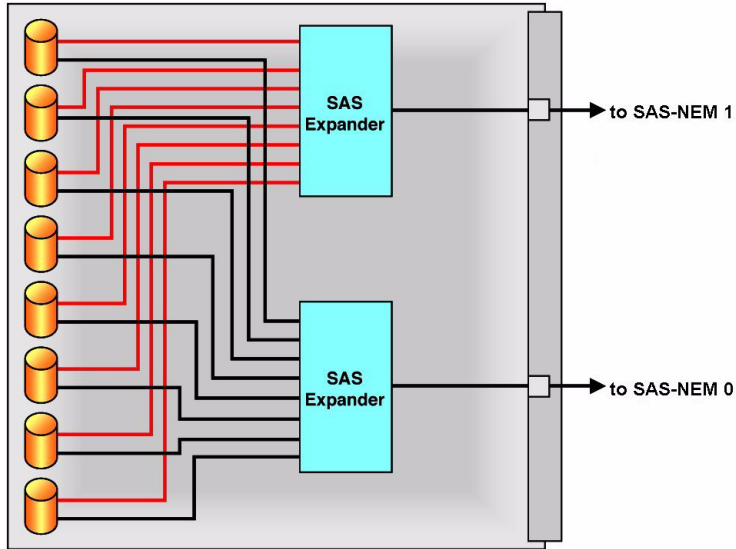
A fully loaded Sun Blade 6000 Disk Module consists of eight SAS disks, with each disk connected to SAS expanders. The SAS expanders, in turn, connect to a Sun Blade 6000 Modular System chassis' SAS-NEMs.

The disk module works only with a SAS-NEM. It does not work with a NEM that lacks SAS connectivity. For the disk module to function, there must be a SAS-NEM in at least slot NEM 0. Dual paths to SAS disks require SAS-NEMs in both NEM slots.

Note – You can use a plain NEM (one without SAS connectivity) in slot NEM 1, but not in slot NEM 0. If you do this, you lose the redundant second path to SAS disks.

FIGURE 1-1 shows a schematic view of the disk module. Each SAS port connects to a single SAS-NEM module.

FIGURE 1-1 Main Components of the Sun Blade 6000 Disk Module



Note – SATA drives are not supported for the disk module at this time.

FIGURE 1-2 shows the interior of the disk module, fully loaded with eight SAS disks.

FIGURE 1-2 Interior of the Sun Blade 6000 Disk Module





Figure Legend

-
- 1 Disk drives
 - 2 SAS Expanders
 - 3 Connectors to chassis midplane
-

Sun Blade 6000 Disk Module Front Panel

The front panel of the disk module is shown in [FIGURE 1-3](#).

FIGURE 1-3 Sun Blade 6000 Disk Module Front Panel



Refer to [TABLE 1-1](#) for descriptions of the LED behavior.

TABLE 1-1 Front Panel LED Functions

LED Name	Description
1 Locate button/LED (white)	<p>This LED helps you to identify which system in the rack you are working on in a rack full of servers. It can be set remotely. Locally you can:</p> <ul style="list-style-type: none"> • Push and release this button to make the Locate LED blink for 30 minutes. • If the LED is blinking, push and release this button to make the LED stop blinking. • Hold down the button for 5 seconds to initiate a “push-to-test” mode that illuminates all LEDs on the disk blade and the disks in it for 15 seconds.
2 Ready-to-Remove LED (blue)	This LED is not used.
3 Module Fault LED (amber)	<p>This LED has two states:</p> <ul style="list-style-type: none"> • Off: No fault. • On: A hardware malfunction has been detected, and service action is required.

TABLE 1-1 Front Panel LED Functions (*Continued*)

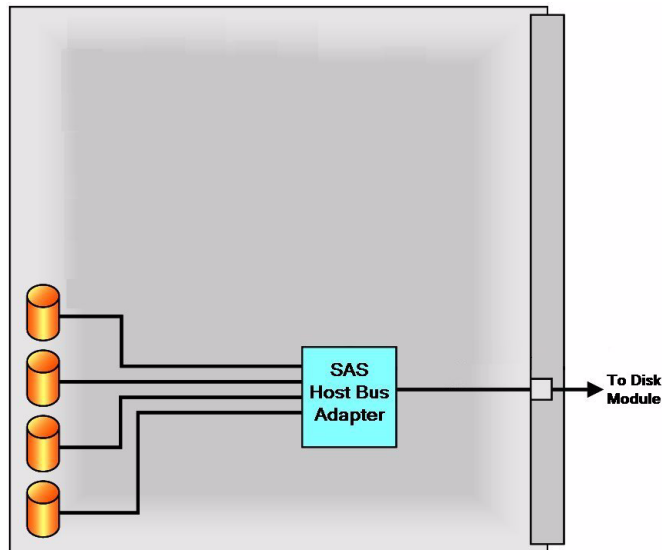
	LED Name	Description
4	Module OK LED (green)	This LED has three states: <ul style="list-style-type: none">• Off: Module is offline.• On: Module is online.• Blinking: Module is configuring, or a flash upgrade is in progress.
5	Disk Drive OK LED (green)	This LED has three states: <ul style="list-style-type: none">• Off: Disk drive is offline.• On: Power is on.• Blinking: Random blink means normal disk activity; slow blink means that a RAID volume is being rebuilt on that disk.
6	Disk Drive Fault LED (amber)	This LED has four states: <ul style="list-style-type: none">• Off: Normal operation.• On: Disk drive fault. Service action required.• Slow blink: Disk drive failure predicted.• Fast blink: Indicates locate function.
7	Disk Drive LED (blue)	This LED is not used.

The Server Module SAS Host Bus Adapter

The disks on a Sun Blade 6000 Disk Module are controlled by a SAS host bus adapter on the server module. This adapter can be either an on-board chip or a RAID Expansion Module (REM) card. [FIGURE 1-4](#) shows the key components of a server that are used to control disks. All of the server modules except the X6450 (no disks) and the T6340 (two disks maximum) can have up to four disks of their own. The SAS host bus adapter on the server module controls its own disks as well as disks on the disk module. The disks on the disk module can be connected to the SAS host bus adapter through two SAS-NEM modules, providing a dual path to each SAS disk.

Note – Server modules with a SAS-2 (up to 6 Gb/s data transfer speed) REM are not supported to work with the Sun Blade 6000 Disk Module. However, SAS-1 NEMs and Sun Blade 6000 Disk Modules might need to be upgraded to allow a server module with a SAS-2 REM to coexist in the chassis with a SAS-1 devices. For more information, refer to [Appendix C](#).

FIGURE 1-4 Connection of SAS Host Bus Adapter for Sun Blade Servers



Server Modules and Disk Modules in a Chassis


The Sun Blade 6000 Disk Module can interoperate with any one of the supported server blades (see [Chapter 3](#)).

The server blades and disk blades must be placed in the chassis in pairs. The pairs must be in slots 0+1, 2+3, 4+5, 6+7, or 8+9. No other combination of slots can be used. A server blade in slot n (n is an even number) sees its own on-blade disks and all the disks in a disk blade in slot $n+1$, but no other disks.

Note – You should place the disk module in an odd-numbered slot. The server module that controls its disks should then be in the even-numbered slot to the left of the disk module. The 10 chassis slots are numbered 0 through 9 from the left to the right, viewed facing the chassis.

[FIGURE 1-5](#) shows a Sun Blade 6000 Modular System chassis, fully loaded with server modules and disk modules.

FIGURE 1-5 Front View of Chassis Filled With Server Modules and Disk Modules in Pairs



Inserting, Moving and Replacing the Sun Blade 6000 Disk Module

This chapter contains the following topics:

- [“Inserting and Controlling a Disk Module”](#) on page 9
- [“Moving Disk Blades”](#) on page 10
- [“Replacing Disk Blades”](#) on page 10

Inserting and Controlling a Disk Module



Caution – You need to save a file that contains the LSI host bus adapter persistence mappings for every server module that: 1) runs the Solaris OS, 2) uses an LSI host bus adapter, and 3) is paired with a disk module. Use the procedure described in [“To Save a Snapshot of Your Host Bus Adapter Persistent Mappings”](#) on page 127.

If the server blade uses an Adaptec host bus adapter, this procedure is not necessary.

The disk module should be placed in an odd-numbered slot. The server module that controls its disks should then be in the even-numbered slot to the left of the disk module. The 10 chassis slots are numbered 0 through 9 from the left to the right, viewed facing the chassis.

The disk module receives its power directly from the chassis. It powers on automatically when inserted.

After insertion, you can power on and off the disk module using the Sun Blade 6000 Oracle Integrated Lights Out Manager (ILOM) (see [Chapter 10](#)).

Moving Disk Blades

Disk blades can be moved from one location to another, with these caveats:

- If a disk blade does not contain a boot disk or RAID volume, it can be moved to another slot in any chassis and paired with a different server blade.
- If a disk blade contains a RAID volume and is moved to another slot in any chassis, the RAID volume becomes “inactive.” If you have a reason to use the RAID volume, you need to reactivate it when the disk blade is in its new slot. Instructions for activating RAID volumes can be found in the appendix to the *Sun Blade 6000 Disk Module Service Manual* (part number 820-1703).

Note – The statements above apply only if the disk blade is not *replacing* another disk blade. If you are moving a disk blade to replace another one, you must follow the procedure described in the *Sun Blade 6000 Disk Module Service Manual* (part number 820-1703).

Moving Pairs of Server Blades and Disk Blades in the Same or a Different Chassis

When powered down, a server blade and disk blade pair can be moved to different slots in the same or a different chassis. The pair can be placed in slots 0+1, 2+3, 4+5, 6+7, or 8+9, but not in any other combination of slots.

Replacing Disk Blades

If your Sun Blade 6000 Disk Module fails, you must replace it.



Caution – You cannot simply remove a disk blade and plug in another one. A service procedure is required.

The procedures for replacing the disk module vary, depending on the following:

- What type of SAS host bus adapter is used.
- Which OS the server blade paired with the disk blade is running.

These procedures are discussed in detail in the *Sun Blade 6000 Disk Module Service Manual* (part number 820-1703).

System Requirements

This chapter describes the requirements for using the Sun Blade 6000 Disk Blade with a Sun blade server module. This chapter contains the following topics:

- “Supported Servers, Firmware and Operating Systems” on page 14
 - “Meeting the Minimum Requirements for Disk Blade Support” on page 15
 - “Adding Patches to Solaris” on page 16
 - “Enclosure Driver Required for Windows” on page 16
 - “Important OS and Server Considerations” on page 16
- “Supported Disk Drives” on page 17
- “Supported RAID Controllers and Host Bus Adapters (HBA)” on page 18
- “Supported Network Express Modules” on page 21
- “Supported Blade Chassis” on page 22
- “Drivers, Firmware and Software” on page 22
 - “Minimum Operational Requirements” on page 22
 - “Oracle Common Array Manager (CAM) Software” on page 23
 - “lsiutil Software” on page 23
 - “Where to Get Drivers, Firmware and Software” on page 24
- “Obtaining the Latest Information” on page 24

Supported Servers, Firmware and Operating Systems

At the time of initial release, the Sun Blade 6000 Disk Module works with the combinations of server blades and operating systems shown in TABLE 3-1. For later combinations, look online (see “[Obtaining the Latest Information](#)” on page 24).

TABLE 3-1 Supported Servers, Firmware and Operating Systems

Server	SPARC/ x86	Minimum System Firmware and Patches	Linux	VMware	Windows	Solaris OS and Patches (in order of installation)
T6300	SPARC	SysFW: 6.6.5 + 136931-04 Disk Controller: 1.24.93.00, + 139418-01	N/A	N/A	N/A	S10 05/08: + 119254-64, + 125555-02, + 138866-02
T6320 (with RAID 0/1 REM)	SPARC	SysFW: 6.6.5 + 136931-04 Disk Controller: 1.24.93.00, + 138445-02	N/A	N/A	N/A	S10 10/08 S10 05/08: + 119254-64, + 125555-02, + 138866-02
T6320 (with RAID 0/1 G2 REM)	SPARC	SysFW: 7.1.6 + 136931-04 Disk Controller: 1.24.94.00, + 139419-01	N/A	N/A	N/A	S10 10/08 S10 05/08: + 119254-64, + 125555-02, + 138866-02
T6340 (with RAID 0/1 G2 REM)	SPARC	SysFW: 7.1.6.d Disk Controller: 1.24.94.00, + 139419-01	N/A	N/A	N/A	S10 10/08 S10 05/08: + 119254-64, + 125555-02, + 138866-02
X6220	x86	BIOS: 110 ILOM: 2.0.3.10, build 36968	RHEL 4.6 (32/64), RHEL 4.7 (32/64), RHEL 5.1 (64), SLES 10 SP1 (64)	ESX 3.0.3, ESX 3.5 U3	2003 R2 SP2 (32/64), 2008 Data Cntr. (32/64)	S10 10/08 S10 05/08: + 138880-01
X6240	x86	BIOS: 1.08 ILOM: 2.0.3.9, build 36997	RHEL 4.6 (32/64), RHEL 5.1 (64), SLES 9 SP4 (64), SLES 10 SP2 (64)	ESX 3.5 U2, ESXi 3.5 U2	2003 R2 SP2 (32/64), 2008 Data Cntr. (32/64)	S10 05/08: + 138880-01

TABLE 3-1 Supported Servers, Firmware and Operating Systems (*Continued*)

Server	SPARC/ x86	Minimum System Firmware and Patches				Solaris OS and Patches (in order of installation)
			Linux	VMware	Windows	
X6250	x86	BIOS: v11 ILOM: 2.0.3.6	RHEL 4.6 (32/64), RHEL 4.7 (32/64), RHEL 5.0 (64), RHEL 5.1 (64), RHEL 5.2 (64), RHEL 5.3 (64), SLES 9 SP3 (64), SLES 9 SP4 (64), SLES 10 SP1 (64), SLES 10 SP2 (64)	ESX 3.0.2 U1	2003 R2 SP2 (32/64), 2008 Data Cntr. (32/64)	S10 10/08 S10 05/08: + 138880-01
X6270	x86	Initial release supported	RHEL 4.6 (32/64), RHEL 4.7 (32/64), RHEL 5.1 (64), RHEL 5.2 (64), SLES 9 SP4 (64), SLES 10 SP1 (64), SLES 10 SP2 (64)	ESX 3.0.2 U1, ESX 3.5 U1, ESX 3.5 U2	2003 R2 SP2 (32/64), 2008 Data Cntr. (32/64)	S10 10/08 S10 05/08: + 138880-01
X6440	x86	Initial release supported	RHEL 4.6 (32/64), RHEL 5.1 (64), SLES 9 SP4 (64), SLES 10 SP2 (64)	ESX 3.5 U2, ESXi 3.5 U2	2003 R2 SP2 (32/64), 2008 Data Cntr. (32/64)	S10 05/08: + 138880-01
X6450	x86	BIOS: 3B16 ILOM: 2.0.3.6	RHEL 4.6 (32/64), RHEL 4.7 (32/64), RHEL 5 (64), RHEL 5.2 (64), SLES 9 SP4 (64), SLES 10 SP2 (64)	ESX 3.5, ESX 3.5 U2	2003 R2 SP2 (32/64), 2008 Data Cntr. (32/64)	S10 05/08: + 138880-01

Meeting the Minimum Requirements for Disk Blade Support

You must ensure that all server blades and their disk controllers, or host bus adapters, meet the minimum firmware requirements whether or not you intend to pair them with a disk blade. Additional details on upgrading server blade firmware may be found in [Chapter 4](#).

At the release of this document, there is currently no support for using the disk module with a server module that has a SAS-2 REM (SAS-2 supports 6 Gb per second data transfer speeds). The Sun Blade 6000 Disk Module is only supported

with server modules that have SAS-1 disk controllers or REMs. However, if you plan on installing a server module with a SAS-2 REM into the chassis, even if it will not be used with the disk module, a firmware upgrade to allow SAS-1/SAS-2 coexistence in the chassis might be required. Refer to [Appendix C](#).

Adding Patches to Solaris

The patches listed in [TABLE 3-1](#) for the Solaris OS are *in addition to* the patches that are required for the server blade. The server blade has its own patch list. These are additional patches needed for use with disk blades.

Enclosure Driver Required for Windows

The Windows Enclosure Device Driver for Sun Blade 6000 Disk Modules is required for all supported versions of Windows. This Windows-certified driver allows Windows to recognize the enclosure services devices within the storage chassis. See [“Where to Get Drivers, Firmware and Software”](#) on page 24.

Important OS and Server Considerations

- If you are using two SAS-NEMs for redundant paths to SAS disks, are using an LSI controller, and you have disks that are not in hardware RAID volumes, you *must* enable multipathing at the OS level on all active server blades in the chassis. This ensures that the primary path to disks is not lost, which can cause an OS panic. Oracle does not support dual paths to SAS disks if OS multipathing is not enabled.
- For all operating systems, if you put your boot disk on the disk blade, it must be in a hardware RAID volume.
- For SPARC systems, you must either put the OS on a single disk in your server blade or in a RAID 1 volume with both member disks in your server blade.
- For the Solaris OS, you must create RAID volumes before enabling `mpxio`. Once `mpxio` is enabled, you can no longer create RAID volumes or change RAID configurations.
- SUSE Linux can manage dual paths to disk drives at the OS level, *except for boot disks*. You should either create a RAID volume for your OS with the SAS host bus adapter firmware or put the OS on a disk or volume on your server. The disks on the server blades have only one path.

- RHEL 4 Linux can manage dual paths to disk drives at the OS level, *except for boot disks*. You should either create a RAID volume for your OS with the SAS host bus adapter firmware or put the OS on a disk or volume on your server. The disks on the server blades have only one path.
- RHEL 5.0, 5.1, and 5.2 can manage dual paths to disk drives either at the OS level (*including the boot disk*) or through the SAS host bus adapter firmware. Installation may require special boot parameters passed to the kernel for dual-path disks recognized correctly
- Windows Server 2008 includes an MPIO utility and supports dual-pathing at the OS level. It can also handle multipathing through the SAS host bus adapter firmware.
- At the release of this document, Windows Server 2003 does not include an MPIO utility, so dual paths are not supported at the OS level. However, dual-pathing on a server running Windows Server 2003 can be managed through the SAS host bus adapter firmware.

Supported Disk Drives

TABLE 3-2 lists the hard disk drives, both mechanical and Solid State Drives (SSD), that have been tested for use in the Sun Blade 6000 Disk Module.

TABLE 3-2 Supported Disk Drives

Capacity (GB)	Speed (RPM)	Part Number (With Bracket)	Manufacturer	Type	SAS/SATA
32	N/A	540-7841-xx	Intel	SSD	SATA
73	10K	540-7354-xx	Fujitsu	mechanical	SAS
73	10K	540-7354-xx	Hitachi	mechanical	SAS
73	10K	540-7354-xx	Seagate	mechanical	SAS
73	15K	540-7361-xx	Seagate	mechanical	SAS
146	10K	540-7355-xx	Fujitsu	mechanical	SAS
146	10K	540-7355-xx	Hitachi	mechanical	SAS
146	10K	540-7355-xx	Seagate	mechanical	SAS
146	10K	540-7864-xx	Hitachi	mechanical	SAS
146	10K	540-7868-xx	Seagate	mechanical	SAS
300	10K	540-7868-xx	Seagate	mechanical	SAS
300	10K	540-7869-xx	Hitachi	mechanical	SAS

The following rules apply when mixing drives in the disk blade:

- You can mix SAS disk drives with different specifications in a disk blade.
- You can have up to eight SSD drives in the disk blade (for HBA requirements when using SSDs, see [“Supported RAID Controllers and Host Bus Adapters \(HBA\)”](#) on page 18).

Note – Currently, the X4620A Sun Blade RAID 5 Expansion Module (Intel/Adaptec) does not support SSDs. Check the *Sun Blade 6000 Disk Module Product Notes* (820-1709) for the latest information.

- Do not mix SSD and mechanical drives in a single RAID volume.
- Do not mix SATA and SAS drives in a single RAID volume
- Hot spares assigned to a volume must be of the same type as disks in the RAID volume (SAS or SATA; mechanical or SSD).

Supported RAID Controllers and Host Bus Adapters (HBA)

At the time of initial release, the supported SAS RAID controllers and HBAs are shown in [TABLE 3-3](#). As new servers are released, see [“Obtaining the Latest Information”](#) on page 24 for the latest supported combinations.

Note – All of the controllers and HBAs listed here are SAS-1 (up to 3 Gb/s data transfer speed), if your server has a SAS-2 HBA (6 Gb/s), refer to [Appendix C](#).

For information on obtaining the latest firmware, see “Drivers, Firmware and Software” on page 22.

TABLE 3-3 RAID Controllers and Host Bus Adapters

Server	On-board RAID Controller Chip	RAID Expansion Module HBA	Minimum HBA Firmware Required	Minimum Disk Blade and SAS-NEM Firmware Required	Minimum HBA Firmware Required for SSD Support
T6300	LSI 1068E	Not supported	LSI: FW 1.24.93.0	5.02.10	LSI: FW 1.26.90.0
T6320	Not supported	<ul style="list-style-type: none"> • T6320 RAID 0/1 Expansion Module (LSI) • Sun Blade RAID 0/1 G2 Expansion Module (LSI) 	<ul style="list-style-type: none"> • LSI: FW 1.24.93.0 • LSI: FW 1.24.94.0 	5.02.10	<ul style="list-style-type: none"> • LSI: FW 1.26.90.0 • LSI: FW 1.26.90.0
T6340	Not supported	Sun Blade RAID 0/1 G2 Expansion Module (LSI)	LSI: FW 1.24.94.0	5.02.10	LSI: FW 1.26.90.0
X6220	LSI 1068E	Not supported	LSI: FW 1.24.93.00, BIOS 6.20.03	5.02.10	LSI: FW 1.26.90.00, BIOS 6.24.01
X6240	Not supported	<ul style="list-style-type: none"> • Sun Blade RAID 0/1 G2 Expansion Module (LSI) • Sun Blade RAID 5 Expansion Module (Intel/Adaptec) 	<ul style="list-style-type: none"> • LSI: FW 1.24.94.00, BIOS 6.20.03, Fcode 1.00.49 • Adaptec: FW 5.2.0-16795, BIOS 5.2.0-15825, Fcode 5.2.0-15825 	5.02.10	<ul style="list-style-type: none"> • LSI: FW 1.26.90.0, BIOS 6.24.01, Fcode 1.00.49 • Adaptec: No SSD support
X6250	Not supported	Sun Blade RAID 5 Expansion Module (Intel/Adaptec)	Adaptec: FW 5.2.0-16795, BIOS 5.2.0-15825, Fcode 5.2.0-15825	5.02.10	Adaptec: No SSD support

TABLE 3-3 RAID Controllers and Host Bus Adapters (Continued)

Server	On-board RAID Controller Chip	RAID Expansion Module	HBA Minimum HBA Firmware Required	Minimum Disk Blade and SAS-NEM Firmware Required	Minimum HBA Firmware Required for SSD Support
X6270	Not supported	<ul style="list-style-type: none"> Sun Blade RAID 0/1 G2 Expansion Module (LSI) Sun Blade RAID 5 Expansion Module (Intel/Adaptec) 	<ul style="list-style-type: none"> LSI: FW 1.24.94.00, BIOS 6.20.03, Fcode 1.00.49 Adaptec: FW 5.2.0-16795, BIOS 5.2.0-15825, Fcode 5.2.0-15825 	5.02.10	<ul style="list-style-type: none"> LSI: FW 1.26.90.0, BIOS 6.24.01, Fcode 1.00.49 Adaptec: No SSD support
X6440	Not supported	<ul style="list-style-type: none"> Sun Blade RAID 0/1 G2 Expansion Module (LSI) Sun Blade RAID 5 Expansion Module (Intel/Adaptec) 	<ul style="list-style-type: none"> LSI: FW 1.24.94.00, BIOS 6.20.03, Fcode 1.00.49 Adaptec: FW 5.2.0-16795, BIOS 5.2.0-15825, Fcode 5.2.0-15825 	5.02.10	<ul style="list-style-type: none"> LSI: FW 1.26.90.0, BIOS 6.24.01, Fcode 1.00.49 Adaptec: No SSD support
X6450	Not supported	<ul style="list-style-type: none"> Sun Blade RAID 0/1 G2 Expansion Module (LSI) Sun Blade RAID 5 Expansion Module (Intel/Adaptec) 	<ul style="list-style-type: none"> LSI: FW 1.24.94.00, BIOS 6.20.03, Fcode 1.00.49 Adaptec: FW 5.2.0-16795, BIOS 5.2.0-15825, Fcode 5.2.0-15825 	5.02.10	<ul style="list-style-type: none"> LSI: FW 1.26.90.0, BIOS 6.24.01, Fcode 1.00.49 Adaptec: No SSD support

For supported RAID controllers and HBAs, note the following

- The on-board LSI 1068E chip, the T6320 RAID 0/1 Expansion Module, and the Sun Blade RAID 0/1 G2 Expansion Module are functionally equivalent. The latter two are REMs that use the LSI 1068E chip. Each of the three can create two RAID volumes that are RAID 0, 1, or 1E. Documentation for these LSI host bus adapters can be found in the *Sun LSI 106x RAID User's Guide* (part number 820-4933).
- The Sun Blade RAID 5 Expansion Module is a REM with Intel circuitry and Adaptec firmware. It can create 24 RAID volumes that are RAID 0, 1, 1E, 10, 5, 5EE, 50, 6, or 60. See the *Sun Intel Adaptec BIOS RAID Utility User's Manual* (part number 820-4708) and the *Sun StorageTek RAID Manager Software User's Guide* (part number 820-1177).
- For servers using the on-board LSI 1068E chip, firmware upgrades for the chip are done when upgrading the server's system BIOS and ILOM firmware.

- Support for Solid State Drives (SSD) requires compatible firmware on both the HBA and the disk blade. For more on upgrading firmware, see [“Drivers, Firmware and Software”](#) on page 22.

Note – Currently, the X4620A Sun Blade RAID 5 Expansion Module (Intel/Adaptec) does not support SSDs. Check the *Sun Blade 6000 Disk Module Product Notes* (820-1709) for the latest information.

Supported Network Express Modules

The Sun Blade 6000 Disk Module requires, at minimum, a SAS-NEM in NEM slot 0 for communication with server blades. For dual paths to SAS disk drives, a SAS-NEM is also required in NEM slot 1. SAS-NEMs also include firmware, which may be upgraded using the Oracle Common Array Manager (CAM), see [“Drivers, Firmware and Software”](#) on page 22. The following SAS-NEMs are supported for use with the disk blade:

- Sun Blade 6000 Multi-Fabric Network Express Module (abbreviated Multi-Fabric NEM)
- Sun Blade 6000 10GbE Multi-Fabric Network Express Module (abbreviated 10GbE Multi-Fabric NEM)
- Sun Blade 6000 Virtualized Multi-Fabric 10GbE Network Express Module (abbreviated Virtualized 10GbE Multi-Fabric NEM)

Rules for Mixing NEMs

You can mix SAS-NEMs as follows:

- You can use two Multi-Fabric NEMs, two 10GbE Multi-Fabric NEMs, or two Virtualized 10GbE Multi-Fabric NEMs.
- You can mix a Multi-Fabric NEM and a 10GbE Multi-Fabric NEM. It does not matter which type is in slot NEM 0.
- You can not mix a Virtualized 10GbE Multi-Fabric NEM with any other type of NEM.
- You can use a plain NEM (one without SAS connectivity) in slot NEM 1, but not in slot NEM 0. However, **in this configuration you lose the redundant second path to SAS disks.**
- In all cases, a SAS-NEM must be in slot NEM 0.

- The expander firmware version on each of the SAS-NEMs and the disk blade module must be the same.

Supported Blade Chassis

The Sun Blade 6000 Disk Module is supported in the following Sun blade chassis:

- Sun Blade 6000 Modular System chassis
- Sun Blade 6048 Modular System chassis

Each chassis contains a Chassis Monitoring Module (CMM). The minimum version of CMM required on the chassis is ILOM 2.0.3.10, build 36968.

Drivers, Firmware and Software

This section lists the drivers, firmware and software you need for your installation. Refer to the *Sun Blade 6000 Disk Module Product Notes* (part number 820-1709) for the most current versions.

Minimum Operational Requirements

Every installation requires:

- ILOM 2.0.3.10, build number 36968 on the chassis CMM
- The latest expander firmware on the disk blade
- The latest expander firmware on the SAS-NEM (same version as the disk blade)

Every server must have:

- For x86 blades, the latest SP firmware and system BIOS
- For SPARC blades, the latest SysFW
- The latest LSI firmware if there is an on-board 1068E chip (X6220, T6300)
- The latest LSI REM firmware if a RAID 0/1 REM or RAID 0/1 G2 REM is supported
- The latest Adaptec REM firmware if the RAID 5 REM is supported
- The latest LSI drivers for each OS version supported when using LSI 1068E chip or RAID 0/1 REM (see [“Updating Operating System Drivers on x86 Server Blades” on page 34](#))

- The latest Adaptec drivers for each supported OS when using the RAID 5 REM (see “Updating Operating System Drivers on x86 Server Blades” on page 34)
- For servers running Windows, the Sun Blade 6000 Disk Module enclosure driver is required (see “Enclosure Driver Required for Windows” on page 16).
- Installations utilizing multipath, refer to [Chapter 6](#) for adding the multipath drivers to your operating system.

Oracle Common Array Manager (CAM) Software

Support for managing your Sun Blade 6000 Disk Modules and Sun Blade 6000 System chassis supported SAS-NEMs is available beginning with CAM version 6.1.2. Use CAM to upgrade firmware on your disk blade and SAS-NEMs. To obtain the latest CAM software, go to:

(<http://www.oracle.com/us/products/servers-storage/storage/storage-software/031603.htm>).

Note – At the release of this document, there is currently no support for using the Common Array Manager to upgrade disk module and SAS-NEM firmware to a level that allows SAS-1 device coexistence with server modules with SAS-2 REMs. For more information on upgrading disk module and SAS-NEM firmware to a level that allows SAS-1/SAS-2 coexistence in the chassis, see [Appendix C](#).

Oracle’s Sun Installation Assistant (SIA) Software

Oracle’s Sun Installation Assistant (SIA) can perform a variety of tasks including system BIOS, ILOM service processor and HBA firmware upgrades. Each server blade has its own version of the SIA. To obtain an ISO CD-ROM image of the most recent SIA for each of your server blades go to the SIA download site:

(<http://www.sun.com/systemmanagement/sia.jsp>)

lsiutil Software

The `lsiutil` application is used for service-related procedures and version 1.60 or later should be installed on every server. To obtain the latest `lsiutil` software, go to the disk blade site:

(<http://www.oracle.com/goto/blades>)

Where to Get Drivers, Firmware and Software

- For updating the necessary **system firmware and drivers**, go to:
(<http://www.oracle.com/goto/blades>)
- For updating **HBA Expansion Module firmware and drivers**, go to:
 - Adaptec Support
Site (<http://support.intel.com/support/motherboards/server/sunraid/index.htm>)
 - LSI Product Support (<http://www.lsillogic.com/support/sun>)
- For updating **hard disk firmware**, identify the firmware version that you need, obtain the patches and software to install them at:
(<http://www.oracle.com/goto/blades>)
- For updating the **disk blade and SAS-NEM firmware**, use the Common Array Manager.

Note – At the release of this document, there is currently no support for using the Common Array Manager to upgrade disk module and SAS-NEM firmware to a level that allows SAS-1 device coexistence with server modules with SAS-2 REMs. For more information on upgrading disk module and SAS-NEM firmware to a level that allows SAS-1/SAS-2 coexistence in the chassis, see [Appendix C](#).

Obtaining the Latest Information

If you do not see your server blade listed in this chapter, or if you think there might have been recent changes to the system requirements, drivers, or firmware for your server blade, you can find the latest version information in the *Sun Blade 6000 Disk Module Product Notes* (820-1809):

(<http://docs.sun.com/app/docs/coll/blade6000dskmod>)

Alternatively, to find the most recent versions and download them, go to the disk blade site:

(<http://www.oracle.com/goto/blades>)

Upgrading Systems

You need to make sure that your server blades and host bus adapters all have the minimum supported firmware and software. If they do not, they need to be upgraded.

Similarly, you need the most recent drivers so that your OS works correctly with the firmware on the host bus adapter.

This chapter discusses updating software and firmware, updating drivers, and other related subjects. It contains the following topics:

- [“Replacing Older X6220 Server Blades” on page 26](#)
- [“Upgrading the Chassis ILOM” on page 27](#)
- [“Overview of Upgrading Firmware on x86 Server Blades” on page 27](#)
- [“Downloading SIA for Each x86 Server Blade” on page 28](#)
- [“Checking SP Firmware and System BIOS Versions on x86 Server Blades” on page 28](#)
- [“Upgrading SP Firmware and System BIOS on x86 Server Blades” on page 31](#)
- [“Checking Host Bus Adapter Firmware Versions on x86 Server Blades” on page 31](#)
- [“Upgrading Host Bus Adapter Firmware on x86 Server Blades” on page 34](#)
- [“Updating Operating System Drivers on x86 Server Blades” on page 34](#)
- [“Upgrading Firmware on SPARC Server Blades” on page 41](#)
- [“Upgrading Disk Blade and SAS-NEM Firmware” on page 44](#)
- [“Updating Solaris Network Installation Servers” on page 45](#)


Replacing Older X6220 Server Blades

Older Sun Blade X6220 servers might have a version of the LSI 1068E SAS host bus adapter hardware that is incompatible with the Sun Blade 6000 Disk Module. Many older server blades have been upgraded to the newer 1068E controller, however, so you must check each older server that you intend to use with the disk blade.

You can use the LSI BIOS configuration utility to determine if your X6220 server blade has the correct hardware version of the 1068E controller.

▼ To Verify the Correct 1068E Version

1. **Reboot your server.**
2. **Watch the server BIOS screen during reboot for the LSI BIOS entry. When you see it, press Ctrl-C.**



```
LSI Logic Corp. MPT SAS BIOS
MPTBIOS-6.16.00.00 (2007.05.07)
Copyright 2000-2007 LSI Logic Corp.

Initializing..!_
```

3. **When the configuration utility opens, highlight the adapter (SAS1068E) and press Enter.**
The “Adapter Properties -- SAS1068E” screen opens.
4. **Use the arrow keys to highlight Advanced Adapter Properties and press Enter.**
The “Advanced Adapter Properties -- SAS1068E” screen opens.

```
LSI Corp Config Utility      v6.20.01.00 (2008.03.13)
Advanced Adapter Properties -- SAS1068E

IRQ                          0A
NUM                          Yes
IO Port Address              7800
Chip Revision ID             04

Advanced Device Properties
Adapter Timing Properties
PHY Properties
```

5. If the Chip Revision ID is 04, you have the correct version. If the Chip Revision ID is 02, you must have your server blade reworked or replaced by Oracle before you can use it with the Sun Blade 6000 Disk Module.

Upgrading the Chassis ILOM

At a minimum the Chassis Monitoring Module (CMM) firmware needs to be at ILOM version 2.0.3.10, build 36968.

To check the CMM ILOM version, open a browser, browse to the IP address of the CMM, and log in. The opening screen will show the current ILOM firmware version. If it is not the version listed above, upgrade it before continuing. See “[Upgrading CMM ILOM Firmware](#)” on page 102.

Overview of Upgrading Firmware on x86 Server Blades

Checking service processor (SP) firmware and system BIOS versions for x86 server blades can be done using ILOM without the need to shut down the blade.

Checking host bus adapter firmware versions requires a reboot.

Upgrading SP firmware and system BIOS (they are always upgraded together) and host bus adapter firmware all require shutting down the server blade and booting with Oracle’s Sun Installation Assistant (SIA) CD or image.

Downloading SIA for Each x86 Server Blade

To upgrade your x86 server blade or host bus adapter firmware, use SIA. Each server blade has its own version of the SIA. To obtain an ISO CD-ROM image of the most recent SIA for each of your server blades go to:

(<http://www.sun.com/systemmanagement/sia.jsp>)

Scroll to your server and click Download. On the login page, select **Installation** from the Platform drop-down list box.

You can also obtain the SIA image from the disk blade page:

(<http://www.oracle.com/goto/blades>)

Checking SP Firmware and System BIOS Versions on x86 Server Blades

These are the minimum versions required in order to use the Sun Blade 6000 Disk Module:

TABLE 4-1 Supported Server Blade BIOS and ILOM Firmware

Server Blade	Minimum System BIOS	Minimum SP Firmware
X6220	110	2.0.3.10, build 36968
X6240	1.08	2.0.3.9, build 36997
X6250	v11	2.0.3.6
X6270	Initial release supported	Initial release supported
X6440	Initial release supported	Initial release supported
X6450	3B16	2.0.3.6
X6270 M2	Initial release supported	Initial release supported

Note – The system BIOS listed on the download site is often not the full name. For example, the complete name for the X6220 blade is OABJT110.

Open a browser, browse to the IP address of the CMM, and log in. Select each server blade in turn to leave the CMM and go to the server's SP. The opening screen for each server's SP shows the SP firmware version. In the following illustration it is 2.0.3.10, build number 36968.

System Information				System Monitoring				Configurat			
Versions		Session Time-Out		Components		Identi					

Versions

View the version of ILOM firmware currently in use.


Version Information	
Property	Value
SP Firmware Version	2.0.3.10
SP Firmware Build Number	36968
SP Firmware Date	Mon Sep 15 21:06:29 PDT 2008
SP Filesystem Version	0.1.14

To view the BIOS version, select the Components tab. A list of all the components appears.

System Information		System Monitoring		Configuration		User Management	
Versions	Session Time-Out	Components	Identification Information				

Component Management

View component information, or prepare to install or remove a component from this page. Components without radio buttons cannot be modified. Choosing the *Prepare to Remove* action shuts d

Component Management Status		
— Actions —		
	Component Name	Type
<input type="radio"/>	/SYS	Host System
<input type="radio"/>	/SYS/SP	Service Processor
<input type="radio"/>	/SYS/SP/NET0	Network Interface
<input type="radio"/>	/SYS/MB	Motherboard
<input type="radio"/>	/SYS/MB/NET0	Network Interface
<input type="radio"/>	/SYS/MB/NET1	Network Interface
<input type="radio"/>	<u>/SYS/MB/BIOS</u>	BIOS

Click the component name that corresponds to the component type BIOS (here it is /SYS/MB/BIOS). A dialog box opens that shows the BIOS version, which is OABJT110 here.

/SYS/MB/BIOS	
Property	Value
fru_name	SYSTEM BIOS
fru_manufacturer	AMERICAN MEGATRENDS
fru_version	OABJT110
fru_part_number	AMIBIOS8

Upgrading SP Firmware and System BIOS on x86 Server Blades

The SP firmware and system BIOS are upgraded together using SIA. See [Appendix A](#).

Note – The SIA CD or ISO image includes the correct firmware versions.

Tip – If you use SIA to upgrade the SP firmware and system BIOS, you need to boot to the SIA image. Because the SIA *also* upgrades host bus adapter firmware (for both LSI and Adaptec adapters), you should elect to upgrade the adapter firmware before exiting SIA. When you choose SIA's Upgrade HBA Firmware task, SIA tells you the firmware version you have currently and the latest version. This saves you from needing to check your adapter firmware with an additional reboot.

Checking Host Bus Adapter Firmware Versions on x86 Server Blades

At the time of initial release of the disk blade, the minimum firmware levels for on-board chip controllers or HBAs are shown below:

TABLE 4-2 Supported HBA Firmware Versions

Server	On-Board Chip or Host Bus Adapter	Minimum Firmware for Disk Blade Support	Minimum Firmware for SSD Support
X6220	On-board LSI 1068E chip	LSI: FW 1.24.93.00, BIOS 6.20.03	LSI: FW 1.26.90.0, BIOS 6.24.01
X6240, X6270, X6440, X6450	Sun Blade RAID 0/1 G2 Expansion Module (LSI)	LSI: FW 1.24.94.00, BIOS 6.20.03, Fcode 1.00.49	LSI FW 1.26.90.0, BIOS 6.24.01, Fcode 1.00.49
X6240, X6250, X6270, X6440, X6450	Sun Blade RAID 5 Expansion Module (Intel/Adaptec)	Adaptec: 15825	No SSD support

▼ To Check LSI Adapter Firmware Version by Rebooting

1. Power-cycle your server module.

The BIOS screen appears.

2. Watch for the LSI Logic Corp. screen which comes up soon after the opening screen.




```
LSI Logic Corp. MPT SAS BIOS
MPTBIOS-6.04.07.00 (2005.11.03)
Copyright 2000-2005 LSI Logic Corp.

Press Ctrl-C to start LSI Logic Configuration Utility...
```

3. When the BIOS screen shows the LSI Logic Corp. message, press Ctrl-C to start the LSI Logic Configuration Utility.

The first screen of the utility appears after a short delay.



```
LSI Corp Config Utility      v6.20.01.00 (2008.03.13)
Adapter List Global Properties
Adapter  PCI  PCI  PCI  PCI  FW Revision  Status  Boot
         Bus Dev Fnc Slot
SAS1068E 02  00  00  00  0.24.30.00-IR  Enabled  0
```

4. With the LSI 1068E adapter highlighted in the first screen, press Enter.

The main screen of the utility opens.

```

LSI Corp Config Utility   v6.20.01.00 (2008.03.13)
Adapter Properties -- SAS1068E

Adapter                SAS1068E
PCI Slot                00
PCI Address(Bus/Dev/Func) 02:00:00
MPT Firmware Revision   0.24.30.00-IR
SAS Address             50000200:00422470
NVDATA Version         2D.03
Status                 Enabled
Boot Order              0
Boot Support            Enabled BIOS & OS

RAID Properties

SAS Topology

Advanced Adapter Properties

Esc = Exit Menu      F1/Shift+1 = Help
Enter = Select Item  -/+Enter = Change Item

```

The MPT BIOS version number is at the top (LSI Corp Config Utility). Here it is 6.20.01.00. The MPT Firmware Revision shows the firmware version. Here it is 0.24.30.00.

▼ To Check Adaptec Adapter Firmware Version Without Rebooting

If you are running the Sun StorageTek RAID Manager software, you can check the Adaptec host bus adapter firmware version without rebooting.

1. In the main screen of the GUI, select the host bus adapter (labeled controller) in the navigation tree on the left.
2. Click the Properties button on the tool bar.

The firmware version is shown.

▼ To Check Adaptec Adapter Firmware Version by Rebooting

1. Reboot and watch the BIOS screen.
2. When the Adaptec section of the BIOS appears, the firmware version is shown.

The host bus adapter firmware version is shown at the top as the build number. Here it is 15825.

```
Adaptec RAID BIOS V5.2-0 [Build 15825]
(c) 1998-2008 Adaptec, Inc. All Rights Reserved.

<<< Press <Ctrl><A> for Adaptec RAID Configuration Utility!

Booting the Controller Kernel.....Controller started

Controller #00: Sun STK RAID REM at PCI Slot:00, Bus:01, De
Waiting for Controller to Start...Controller started
Controller monitor V5.2-0[15829], Controller kernel V5.2-0[
```

Note – You do not need to enter the Adaptec RAID configuration utility. Just note the number and continue booting.

Upgrading Host Bus Adapter Firmware on x86 Server Blades

Use SIA. See [Appendix A](#).

Note – The SIA CD or ISO image includes the correct firmware versions.

Updating Operating System Drivers on x86 Server Blades

If you have an x86 server blade, such as a X6220, X6240, X6250, X6270, X6440 or X6450 server blade you need to confirm that you have the correct disk controller driver version installed to ensure compatibility with your disk module. The following tables list supported drivers as follows:

- [TABLE 4-3](#) lists the minimum HBA driver versions required to support the Sun Blade 6000 Disk Module
- [TABLE 4-4](#) lists the minimum HBA driver versions required to support SSD drives in the Sun Blade 6000 Disk Module

If your drivers are older versions, newer versions can be found at.

(<http://www.oracle.com/goto/blades>)

TABLE 4-3 Minimum HBA Drivers Required for Disk Module Support

Operating System	LSI 1068E On-Board Chip	LSI RAID 0/1 G2 REM	Adaptec RAID 5 REM
Solaris 10 05/08 OS (x86)	MPT 1.77	MPT 1.77	15817
Solaris 10 10/08 OS (x86)	MPT 1.80	MPT 1.80	15872
Red Hat Enterprise Linux (RHEL) 4.6 (32/64)	MPT 3.12.27.00-2	MPT 3.12.27.00-2	1.1.5.2459
RHEL 4.7 (32/64)	MPT 3.12.27.00-2	MPT 3.12.27.00-2	1.1.5.2459
RHEL 5 (64)	MPT 4.00.38.07-1	MPT 4.00.38.07-1	1.1.5.2459
RHEL 5.1 (64)	MPT 4.00.38.07-1	MPT 4.00.38.07-1	1.1.5.2459
RHEL 5.2 (64)	MPT 4.00.41.00-1	MPT 4.00.41.00-1	1.1.5.2459
SUSE Linux Enterprise Server (SLES) 9 SP 4	MPT 3.12.27.00-2	MPT 3.12.27.00-2	1.1.5.2459
SLES 10 SP1	MPT 4.00.38.07-1	MPT 4.00.38.07-1	1.1.5.2459
SLES 10 SP2	MPT 4.00.41.00-1	MPT 4.00.41.00-1	1.1.5.2459
Windows Server 2003 (32/64)	1.28.03.00	1.25.10.05	5.2.0-15816
Windows Server 2008 (32/64)	1.28.03.00	1.26.05.00	5.2.0-15816

TABLE 4-4 Minimum Drivers Required for SSD Support in a Disk Module

Operating System	LSI 1068E On-Board Chip	LSI RAID 0/1 G2 REM	Adaptec RAID 5 REM
Solaris 10 05/08 OS (x86)	MPT 1.77	MPT 1.77	No SSD support
Solaris 10 10/08 OS (x86)	MPT 1.80	MPT 1.80	No SSD support
Red Hat Enterprise Linux (RHEL) 4.6 (32/64)	MPT 3.16.00.00	MPT 3.16.00.00	No SSD support
RHEL 4.7 (32/64)	MPT 3.16.00.00	MPT 3.16.00.00	No SSD support
RHEL 5 (64)	MPT 4.16.00.00	MPT 4.16.00.00	No SSD support
RHEL 5.1 (64)	MPT 4.16.00.00	MPT 4.16.00.00	No SSD support
RHEL 5.2 (64)	MPT 4.16.00.00	MPT 4.16.00.00	No SSD support
SUSE Linux Enterprise Server (SLES) 9 SP 4	MPT 3.16.00.00	MPT 3.16.00.00	No SSD support
SLES 10 SP1	MPT 4.16.00.00	MPT 4.16.00.00	No SSD support

TABLE 4-4 Minimum Drivers Required for SSD Support in a Disk Module (*Continued*)

Operating System	LSI 1068E On-Board Chip	LSI RAID 0/1 G2 REM	Adaptec RAID 5 REM
SLES 10 SP2	MPT 4.16.00.00	MPT 4.16.00.00	No SSD support
Windows Server 2003 (32/64)	1.28.03.00	1.28.03.00	No SSD support
Windows Server 2008 (32/64)	1.28.03.00	1.28.03.00	No SSD support

Linux Drivers

To determine the current LSI or Adaptec driver on your server.

- For driver information on an LSI chip or HBA, enter the command:

```
% lsmod | grep mpt
```

- For driver information on an Adaptec HBA, enter the command:

```
% lsmod | grep aac
```

If your drivers are older versions than those listed in [TABLE 4-3](#) or [TABLE 4-4](#) (if you are using SSD drives in your disk blade), replace them with the listed ones. The latest drivers for your server are available at:

(<http://www.oracle.com/goto/blades>)

Solaris Drivers

To determine the current LSI or Adaptec driver on your server.

- For driver information on an LSI chip or HBA, enter the command:

```
% modinfo | grep mpt
```

- For driver information on an Adaptec HBA, enter the command:

```
% modinfo | grep aac
```

If your drivers are older versions than those listed in [TABLE 4-3](#) or [TABLE 4-4](#) (if you are using SSD drives in your disk blade), replace them with the listed ones. The latest drivers for your server are available at:

(<http://www.oracle.com/goto/blades>)

Windows Drivers

If you are running Microsoft Windows Server 2003 or Windows Server 2008, you will need to confirm that you have the correct disk controller driver version installed. If your drivers are older versions than those listed in [TABLE 4-3](#) or [TABLE 4-4](#) (if you are using SSD drives in your disk blade), replace them with the listed ones.

The latest drivers for your server are available at:

(<http://www.oracle.com/goto/blades>)

Follow these procedures to determine if your server requires a driver update.

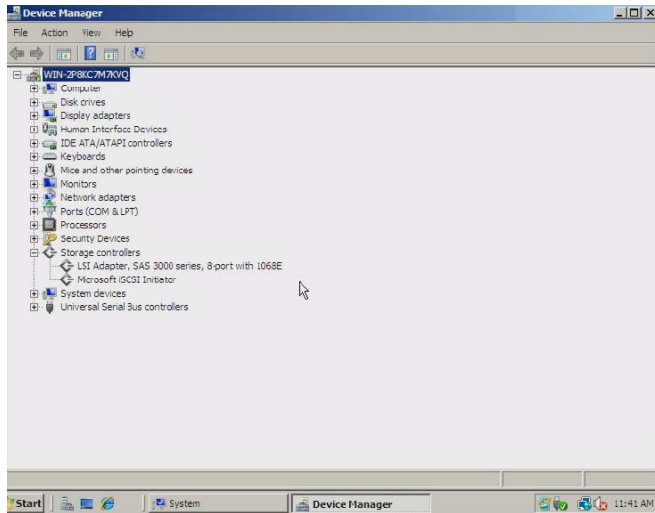
▼ To Upgrade Your Blade Server Windows Driver

Follow the steps below to check your disk controller driver version and, if necessary, perform a driver update.

- 1. Log in to the server blade you want to update.**
- 2. Click Start, and then click Control Panel.**
- 3. Double-click the System icon.**
- 4. If you are running Windows Server 2008, go to Step 5. If you are running Windows Server 2003, select the Hardware tab.**
- 5. Click Device Manager.**
- 6. Scroll down the list of hardware types and do one of the following:**
 - Expand "SCSI and RAID Host Bus Adapter" if your OS is Windows Server 2003.
 - Expand "Storage Controller" if your OS is Windows Server 2008.

An example showing an LSI 1068E on-board chip can be seen in [FIGURE 4-1](#).

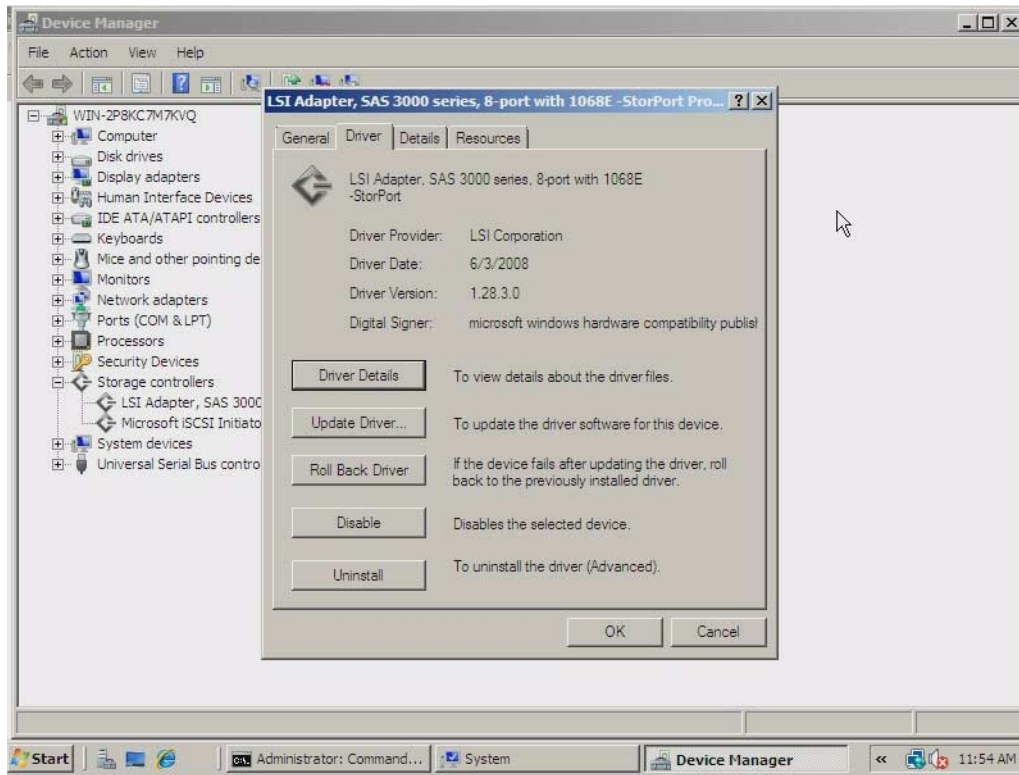
FIGURE 4-1 Example of Device Manager Listing



7. Right-click on the disk controller, and then click Properties.

You will see the Properties dialog box displayed (see [FIGURE 4-2](#)).

FIGURE 4-2 Example Device Properties Dialog Box



8. Compare the driver version listed with the minimum supported driver version required for your disk controller.

- For the minimum version required for disk blade support, see [TABLE 4-3](#).
- For the minimum version required for SSD support in your disk blade, see [TABLE 4-4](#).

9. If your disk controller driver is at or above the minimum version, skip to the next section.

If your disk controller driver is earlier than the minimum version required, do the following:

a. Go to the disk blade download page for your server and obtain the latest system software:

(<http://www.oracle.com/goto/blades>)

You can download the server's latest Tools and Drivers CD ISO image, or the drivers specific to your Windows OS.

- b. Copy (or extract if you downloaded a zip file) the driver files to a folder accessible by the server being updated.
- c. Click the **Update Driver** button from the driver Properties dialog box (see [FIGURE 4-2](#)), and follow the instructions to update the driver with the new driver you downloaded from the Oracle download site.

If you are not at the driver properties dialog box, repeat steps 2-7.

Once the disk controller driver is updated, proceed to the next section.

Updating the Sun Blade 6000 Enclosure Device Driver in Windows

The Sun Blade 6000 Disk Module requires an update to the standard enclosure device driver in Microsoft Windows Server 2003 or 2008 for the enclosure to be seen correctly by Windows Device Manager. You need to download and install an updated driver package from the disk blade site at:

(<http://www.oracle.com/goto/blades>)

Click the View by Category tab. Choose **Systems Administration > Storage Management** and look for the appropriate Sun Blade 6000 Disk Module downloadable package.

▼ To Install the Driver

1. Download the executable file (.exe) to a folder on the server blade to be updated.
2. From the server blade, navigate to the folder where you downloaded the executable file and double-click it.
3. The device driver installation wizard starts.

Follow the on-screen instructions to complete the update.

Upgrading Firmware on SPARC Server Blades

If you have a T6300, T6320 or T6340 server blade, you need to check your server's system firmware version and disk controller version to ensure compatibility with your disk module.

TABLE 4-5 Server Blade Minimum Required Firmware Versions

Server	Minimum Firmware Versions Required	Patches With Minimum Firmware Versions
T6300	<ul style="list-style-type: none">• SysFW: 6.6.5• Disk Controller: 1.24.93.00	<ul style="list-style-type: none">• Patch 136931-04• Patch 139418-01
T6320 (with RAID 0/1 REM - revision B2)	<ul style="list-style-type: none">• SysFW: 7.1.6• Disk Controller: 1.24.93.00	<ul style="list-style-type: none">• Patch 136931-04• Patch 138445-02
T6320 (with RAID 0/1 G2 REM - revision B3)	<ul style="list-style-type: none">• SysFW: 7.1.6• Disk Controller: 1.24.94.00	<ul style="list-style-type: none">• Patch 136931-04• Patch 139419-01
T6340 (with RAID 0/1 G2 REM - revision B3)	<ul style="list-style-type: none">• SysFW: 7.1.6.d• Disk Controller: 1.24.94.00	<ul style="list-style-type: none">• No patch currently required• Patch 139419-01

Follow these procedures to determine if your server requires a firmware upgrade patch.

▼ To Upgrade Your SPARC Blade Server System Firmware

For checking your server's system firmware version, use the Advanced Lights Out Management (ALOM) feature.

If your server is configured to operate under an ILOM shell by default, you can create an ALOM compatibility shell to use the ALOM command specified in this section.

1. **Log on to the service processor in CLI mode as root by entering the command:**

```
ssh root@ipaddress
```

Where *ipaddress* is the IP address of your servers' service processor.

2. After login, you want to be at the ALOM prompt:

```
sc>
```

- If you see the ALOM prompt, skip to Step 3.
- If you see an ILOM prompt:
- ->
- You need to create an ALOM shell. Perform the following steps:

a. Create a user named `admin`, and then set the `admin` account role to `Administrator` and the CLI mode to `alom`.

If a user `admin` with the `Administrator` role has not yet been created, you can combine the `create` and `set` commands on a single line:

```
-> create /SP/users/admin role=Administrator cli_mode=alom
Creating user...
Enter new password: *****
Enter new password again: *****
Created /SP/users/admin
```

Note – The asterisks in the example do not appear when you enter your password.

If the user `admin` with the `Administrator` role already exists, you need only to change the CLI mode to `alom`:

```
-> set /SP/users/admin cli_mode=alom
Set 'cli_mode' to 'alom'
```

b. Log out of the root account after you have finished creating the `admin` account:

```
-> exit
```

You are returned to the ILOM login prompt.

c. Log in to the ALOM CLI shell from the ILOM login prompt:

```
SUNSPxxxxxxxxxx login: admin
Password:
Waiting for daemons to initialize...
Daemons ready
Sun(TM) Integrated Lights Out Manager
Version 2.0.0.0
Copyright 2007 Sun Microsystems, Inc. All rights reserved.
Use is subject to license terms.
sc>
```

The `sc>` prompt indicates that you are in the ALOM CMT compatibility shell.

3. Enter the CLI command to show your server's system firmware (SysFW) version:

```
sc> showhost
```

You will see output similar to the following example:

```
Sun System Firmware 7.1.6.d 2008/09/15 17:10
Hypervisor flash versions:
Hypervisor 1.6.7.a 2008/08/30 05:18
OBP 4.29.0.a 2008/09/15 12:01
```

4. Compare the system firmware version displayed with the required versions listed in TABLE 4-5.

5. If your system firmware is earlier than the minimum version required, go to:

<http://www.oracle.com/goto/blades>

If your system firmware is at or above the minimum version, skip to the next section to confirm your disk controller firmware.

▼ To Upgrade Your SPARC Blade Server Disk Controller Firmware

You can check your server's disk controller firmware version from the Solaris command line:

1. As root, enter the following command:

```
# /usr/sbin/prtconf -vp | egrep  
'(LSI|firmware-version|revision-id|mpt-version)'
```

Your output will look similar to the following example:

```
....  
firmware-version: `1.18.00.00`  
mpt-version: `1.05`  
model: `LSI,1068E`  
name: `scsi`  
revision-id: 00000002  
....
```

2. Compare the disk controller firmware version displayed with the required versions listed in [TABLE 4-5](#).
3. If your disk controller firmware is earlier than the minimum version required, go to:

(<http://www.oracle.com/goto/blades>)

If your disk controller firmware is at or above the minimum version, you need do nothing further at this time.

Upgrading Disk Blade and SAS-NEM Firmware

Sun Blade 6000 Disk Modules and Sun Blade 6000 SAS-NEMs both contain SAS expanders. It is critical that these expanders have the correct firmware. The firmware versions must match for both devices.

The Common Array Manager (CAM) checks that these devices have the correct firmware and generates alarms if they do not. CAM also has a simple procedure for updating the firmware (see “[Upgrading Expander Firmware Using CAM](#)” on [page 109](#)).

Note – At the release of this document, there is currently no support for using the Common Array Manager to upgrade disk module and SAS-NEM firmware to a level that supports SAS-2. For more information on upgrading disk module and SAS-NEM firmware to a level for SAS-2 host support, see [Appendix C](#).

Updating Solaris Network Installation Servers

If you use a Solaris network installation server for server blades that are paired with disk blades, and you are using Solaris 10 05/08, you need to add patches to the Solaris 10 05/08 OS image. Patches are not required for later versions of Solaris. For a list of required patches for Solaris 10 05/08, see [TABLE 3-1](#).

For SPARC Servers

The following is the command you use to apply a patch to a Solaris network installation image's miniroot:

```
# patchadd -C <netinstall image>/Solaris_10/Tools/Boot <patch>
```

If you do not use a Jumpstart server, you need to install Solaris 10 5/08 and then use patchadd to apply the patches.

▼ To Have the Patches Added Automatically

If you use a Jumpstart server, you can have the patches automatically added by following these steps:

1. **Install Solaris 10 05/08.**
2. **Download a stock image of Solaris 10 05/08 and copy the patches into the Solaris_10/Patches directory.**

Note – The order of patch installation is based on the timestamp of the patches, oldest installed first.

For x86 Servers

This procedure must be done on an x86 system running the Solaris 10 05/08 with the latest available packaging and patching utilities already installed. You must then add required patches (see “[Supported Servers, Firmware and Operating Systems](#)” on [page 14](#)).

▼ To Update x86 Solaris Network Installation Servers

1. Change to your Solaris_10/Tools directory.

2. Run `setup_install_server` to a local directory:

```
# ./setup_install_server -b /export/home/s10u5_patch
```

3. Unpack the miniroot.

```
# /boot/solaris/bin/root_archive unpackmedia  
/export/home/s10u5_patch /export/home/s10u5_patch_mr
```

4. Install the patch.

```
# patchadd -C /export/home/s10u5_patch_mr <patch directory>
```

5. Pack up the new miniroot.

```
# /boot/solaris/bin/root_archive packmedia  
/export/home/s10u5_patch /export/home/s10u5_patch_mr
```

On your installation server, use `setup_install_server` and then copy the newly generated `x64.miniroot` over:

6. Change to your Solaris_10/Tools directory.

7. Run `setup_install_server` to a local directory:

```
# ./setup_install_server /export/home/s10u5_patch
```

8. Save the old `x86.miniroot` file:

```
# cd /export/home/s10u5_patch/boot  
# cp -p x86.miniroot x86.miniroot.orig
```

9. Copy the new `x86.miniroot` file from the machine on which you built it.

For example:

```
# cp -p  
/net/<machine_name>/export/home/s10u5_patch/boot/x86.miniroot
```

Installing, Removing or Replacing Disks

This chapter contains these topics:

- “Moving or Replacing Disks” on page 47
- “Disk Drive Compatibility Rules” on page 48
- “Multipathing” on page 48
- “Replacing a Hard Disk Drive” on page 49

Moving or Replacing Disks



Caution – If you remove a disk, wait at least one minute before you insert another in order for the HBA to rescan and recover from the action. This applies to disks in the server blade as well as disks in the disk blade.

Disks can be moved from one location to another, with several caveats:

- A disk that is not in a hardware RAID volume can be moved to a different slot. You do not need to power down the disk blade to remove or insert a disk.

Note – Depending on your OS, you might have trouble locating the disk after it has been moved. The target IDs assigned to a disk are those of the slot it is in. When you move a disk, the target ID assigned by the SAS host bus adapter changes.

- If a disk is a member of a mirrored RAID volume and is removed, the RAID volume will need to rebuild.
 - If there is a hot spare, the rebuild will occur automatically.

- If there is no hot spare, you must replace the disk you removed with a new one in the same slot, and the mirror will use that disk to rebuild.
- If the disk you removed did not fail and you replace it in a different slot, the mirror will rebuild automatically, using all the same physical disks that were present before, including the one you moved.
- If a disk is a member of a striped RAID volume and is removed, various possibilities exist:
 - If the disk has failed, you have lost your striped volume. There is nothing you can do to recover it.
 - If the disk has not failed, you must reinsert it in the same slot or a different slot. The RAID volume will be restored.
 - If the disk has not failed, and you do not reinsert it, you have lost your striped volume. Reinserting the disk you removed will restore the RAID volume.

Disk Drive Compatibility Rules

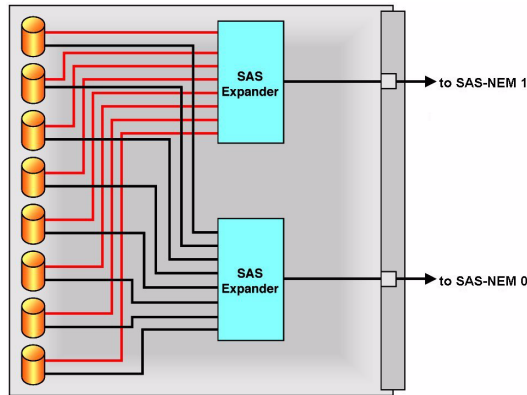
The following rules apply when mixing drives in the disk blade:

- You can mix SAS disk drives with different specifications in a disk blade
- You can have up to eight SSD drives in the disk blade
- Do not mix SSD and mechanical drives in a single RAID volume
- Do not mix SATA and SAS drives in a single RAID volume
- Hot spares assigned to a volume must be of the same type as disks in the RAID volume (SAS or SATA; mechanical or SSD)
- Use only supported disk drives in the disk blade (see [“Supported Disk Drives”](#) on [page 17](#) for a complete list).

Multipathing

Multipathing to the disks on the Sun Blade 6000 Disk Module is accomplished by using SAS drives with two installed SAS-NEMs. [FIGURE 5-1](#) shows the configuration that is used.

FIGURE 5-1 Dual Paths to the Disk Drives on the Sun Blade 6000 Disk Module



Procedures for setting up multipathing are described in [Chapter 6](#).

Replacing a Hard Disk Drive



Caution – Before handling components, attach an ESD wrist strap to the grounding post that is built into the rear of the chassis. The system’s hard disk drives contain components that are extremely sensitive to static electricity.

On occasion, a disk drive might fail. The status of the drive is indicated by its LEDs:

- Green LED - Random or fast blink means normal disk activity, slow blink means that data is being rebuilt on that drive, and off means power is off. On means normal operation.
- Amber LED - This LED is controlled by the expander firmware on the disk blade. Off indicates normal operation. Slow blink predicts imminent disk failure. On indicates that the system has detected a fault and service action is required.
- Blue LED - Not used.



Caution – **Possible data loss:** You must not insert a disk that has been previously configured with a RAID volume into a new RAID volume. Before permanently removing a disk that is part of an active RAID volume, use the appropriate RAID tool to delete the RAID volume from the disk.

A single disk drive failure does not cause a data failure if the disk drives are configured as a mirrored RAID volume. When there is no hot spare assigned to the mirror, the failed disk drive can be hot-swapped; when the new disk drive is inserted, the contents are automatically rebuilt from the rest of the array with no need to reconfigure the RAID parameters.

If the mirror was configured with a hot spare, the mirror is automatically rebuilt with the hot spare.



Caution – Possible data loss: You can remove the failed disk while the mirror is rebuilt to the hot spare, but you must not insert a new disk in its place until the rebuilding of the mirror is completed. While data is being rebuilt, the green LED on the rebuilding drive will blink slowly. The rebuild process can take a number of hours for large mirrors.

The Sun Blade 6000 Disk Module disk drives can be replaced by users.

▼ To Remove a Hard Disk Drive.

1. Observe the LEDs on the faces of the disk drives and identify the defective disk drive.
2. Execute the software commands appropriate to the software that you are using to prepare the hard drive for removal.



Caution – Do not operate the system with empty slots. Always insert a filler into an empty slot to reduce the possibility of module shut down.

SAS disk drives in the Sun Blade 6000 Disk Module are hot-pluggable, so you can remove a SAS disk drive when the disk module is installed in the chassis, as well as when it is out of the chassis. See [FIGURE 5-2](#) and [FIGURE 5-3](#).

Note – When you return the disk drives to their slots, you should return each disk drive to the *same* slot from which it was removed. Use an adhesive note or another method to temporarily label the disk drives when you remove them.

FIGURE 5-2 Removing a Hard Disk Drive From a Disk Module After Removing the Disk Module From the Chassis





FIGURE 5-3 Removing a Hard Disk Drive From a Disk Module Without Removing the Disk Module From the Chassis



3. Press the button on the face of the disk drive to release the spring-loaded securing latch. See [FIGURE 5-2](#) and [FIGURE 5-3](#).

4. Grasp the securing latch and remove the disk drive from the drive slot.

Installing a Disk Drive

Installation is the reverse of the preceding procedure. When you install a disk drive, open its securing latch before you push the drive into the slot. Push the disk drive into the slot until it stops, and then close the securing latch to fully engage the connector on the disk drive backplane.

If the disk drives were previously configured as a mirrored RAID 1 array, an automatic resynchronization is invoked and the contents are automatically rebuilt from the rest of the array with no need to reconfigure the RAID parameters. If the bad disk drive was configured as a hot spare, the new disk drive is automatically configured as a new hot spare.

Multipathing and RAID

This chapter covers the behavior of the supported LSI and Adaptec SAS host bus adapters, with both single and dual paths to SAS disks.

Here are the topics:

- [“Multipathing Overview” on page 53](#)
- [“Multipathing Using LSI Host Bus Adapters” on page 54](#)
- [“Multipathing Using Adaptec Host Bus Adapters” on page 57](#)
- [“Multipathing at the OS Level” on page 57](#)
 - [“Hardware RAID Hides Dual Paths From the OS” on page 58](#)
 - [“Multipathing With RHEL 5.x Linux and LSI Host Bus Adapters” on page 58](#)
 - [“Multipathing With the Solaris OS and LSI Host Bus Adapters” on page 59](#)
 - [“Multipathing With Windows Server 2008 and LSI Host Bus Adapters” on page 62](#)
 - [“Multipathing with VMware ESX 3.5 and LSI Host Bus Adapters” on page 65](#)

Multipathing Overview

With a server blade in a Sun Blade 6000 chassis, there can be two different paths to each SAS disk drive in a disk blade. One path goes through one SAS-NEM and the other goes through a second SAS-NEM.

If there is only one SAS-NEM in the chassis, it must be in slot NEM 0. In such a case there is only one path to each SAS disk.

Multipathing Using LSI Host Bus Adapters

The information in this section is equally valid for an on-board LSI chip or an LSI RAID Expansion Module (REM). Much of the discussion is also relevant to Adaptec controllers.

One SAS-NEM (Single Path)

Configuration: A server blade with four disk drives, one SAS-NEM in slot NEM 0 (providing a single path), and a disk blade with eight disk drives.

Initially, before any attempt to create volumes with the LSI SAS host bus adapter, all twelve disk drives seen by the SAS host bus adapter are passed to the server blade BIOS and the OS. The BIOS sees twelve disk drives at the outset and will continue to do so. The OS, which gets its disk drive information from the BIOS, can make RAID volumes and see a combination of volumes and disk drives.

Note – By definition, software RAID refers to RAID volumes created by an OS, while hardware RAID refers to RAID volumes created by firmware on the SAS HBA (typically configured through an HBA BIOS-based utility). In addition, there are software applications provided with the HBA that run on various operating systems that can utilize the HBA firmware to create hardware RAID (for example, LSI-based RAID controllers use MegaRAID Storage Manager which is available for Linux and Windows). Most of the discussion in this chapter involves hardware RAID.

The one SAS-NEM configuration is straightforward—the server blade BIOS and OS see all disk drives on the server blade and the disk blade as unique disk drives.

You can change this situation by using the LSI BIOS configuration utility or the Solaris utility `raidctl` to create SAS host bus adapter hardware RAID volumes:

- For x86 blades with LSI host bus adapters, you can create hardware RAID volumes by entering the LSI HBA BIOS configuration utility available during server boot (by pressing Ctrl-C when prompted).
- For SPARC blades with LSI host bus adapters, you can create hardware RAID volumes with the Solaris utility `raidctl`.

The 1068E SAS host bus adapter can create only two volumes. The two volumes can contain up to twelve disk drives plus one or two hot spares (no more than fourteen disk drives total).

RAID 0 can have from two to ten striped disk drives.

RAID 1 can have two and only two mirrored disk drives, plus one or two hot spares.

RAID 1E (RAID 1 Enhanced) can have up to twelve mirrored disk drives plus one or two hot spares (if all disk drives are in one volume and there is no other volume).

RAID 1E requires a minimum of three disk drives.

Note – Hot spares created by the LSI HBA BIOS configuration utility are hidden from the server blade BIOS and OS.

Example 1:

Volume V1 contains three server blade disk drives in a RAID 1E volume. The fourth server blade disk drive is a hot spare.

Volume V2 contains seven blade disk drives in a RAID 1E volume. The eighth blade disk drive is a hot spare.

Now the BIOS and OS recognize just two “disk drives” (V1 and V2). The hot spares, which are known to the SAS host bus adapter, are unknown to the BIOS and OS.

Example 2:

Volume V1 contains two server blade disk drives in a RAID 1 volume. One server blade disk drive is a hot spare. One is not included in a volume.

Volume V2 contains three disk blade disk drives in a RAID 1E volume. The fourth disk blade disk drive is a hot spare. The other four disk blade disk drives are not included in a volume.

The server blade BIOS and OS recognize seven disk drives: V1, V2, and five other disk drives not including the hot spares.

Two SAS-NEMs (Dual Path)

Configuration: A server blade with four disk drives, two SAS-NEMs (providing a dual path), and a disk blade with eight disk drives. The second SAS-NEM allows the SAS host bus adapter to connect to each disk drive on the disk blade by two paths—one through each SAS-NEM.

The server blade disk drives have only one path (hard-wired).

Initially, the SAS host bus adapter passes twenty disk drives to the server blade BIOS and OS—four on the server blade and sixteen on the disk blade. The sixteen from the disk blade are actually just eight disk drives, each with two unique SAS addresses, one for each port on the disk.

The server blade BIOS and OS see two instances of each dual-pathed disk drive. OS multipathing, if enabled, hides the second path to dual-pathed disks (see [“Multipathing at the OS Level”](#) on page 57).

Note – The LSI BIOS configuration utility displays a single path in the RAID creation menu so that you cannot create a RAID volume with a disk drive and itself.

After disk drives are configured as RAID volumes, only one instance will be displayed by the server blade BIOS and OS.

One path (the primary) is used by the SAS host bus adapter to read and write from the volume. The other path (secondary) is held in passive reserve in case the primary path is lost.

Note – If failover of the active (primary) path occurs, the SAS host bus adapter re-establishes contact through the passive (secondary) path without user intervention.

Example 1:

Volume V1 contains three server blade disk drives in a 1E RAID volume. The fourth server blade disk drive is a hot spare.

Volume V2 contains seven blade disk drives in a RAID 1E volume. The eighth blade disk drive is a hot spare.

The BIOS and OS report just two “disk drives” (V1 and V2) and not the hot spares.

Example 2:

Volume V1 contains two server blade disk drives in a RAID 1 volume. One server blade disk drive is hot spare. One is not included in a volume.

Volume V2 contains three disk blade disk drives in a RAID 1E volume. A fourth disk blade disk drive is a hot spare. The other four disk blade disk drives are not included in a volume.

The server blade BIOS and OS report eleven disk drives: V1, V2, and nine disk drives not in RAID (but not the hot spares). There are five disk drives not in volumes or hot spares and four of these are recognized twice. If multipathing is enabled, the OS hides the second path.

Multipathing Using Adaptec Host Bus Adapters

There are many similarities between the LSI and Adaptec SAS host bus adapters, and much of the discussion in [“Multipathing Using LSI Host Bus Adapters” on page 54](#) is relevant to Adaptec controllers as well. There are, however, important differences between the LSI and Adaptec functionality.

When you have an Adaptec REM and you plug in a disk blade with eight dual-pathed disk drives, the REM recognizes the disk drives as eight disks with dual paths, but it does not present *any* disk drives to the server blade BIOS or the OS until volumes are created. Volume creation using the Adaptec REM utilities is required. The volumes can be either one-disk drive volumes or multi-disk drive RAID volumes. The Adaptec REM can create up to 24 volumes.

Note – In a situation of one SAS-NEM and single paths, the same volume creation requirement applies.

When volumes are created, the Adaptec REM presents each volume to the server blade BIOS and OS as *one* disk drive. The dual paths are hidden.

This means that OS software for multipathing cannot be used because it cannot see the two paths. With the Adaptec REM, the OS will never be presented with a disk drive that is not in a volume. Using the Adaptec RAID manager software, you can determine that two paths exist, but the OS still cannot recognize more than one.

Multipathing at the OS Level

The Solaris, Red Hat Linux 5.x, Windows Server 2008, and VMware ESX 3.5. operating systems all have utilities for managing dual paths at the OS level. This section contains the following subsections:

- [“Hardware RAID Hides Dual Paths From the OS” on page 58](#)
- [“Multipathing With RHEL 5.x Linux and LSI Host Bus Adapters” on page 58](#)
- [“Multipathing With the Solaris OS and LSI Host Bus Adapters” on page 59](#)
- [“Multipathing With Windows Server 2008 and LSI Host Bus Adapters” on page 62](#)
- [“Multipathing with VMware ESX 3.5 and LSI Host Bus Adapters” on page 65](#)

Hardware RAID Hides Dual Paths From the OS

If you use a RAID configuration utility at the BIOS level to create RAID volumes, these volumes are managed by the SAS host bus adapter (as a hardware RAID) and they are seen by your OS as a single disk. All multipath behavior of disks in the volumes is hidden from the OS by the SAS adapter.

Note – A RAID configuration utility at the BIOS level is not available for SPARC systems. You can still create hardware RAID volumes that are managed by the SAS host bus adapter, however, by using the Solaris OS utility `raidctl`. RAID volumes created with `raidctl` are seen by the Solaris OS as a single disk, with multipath behavior hidden from OS.

Your options for using multipathing with your host bus adapter are as follows:

- With an LSI host bus adapter, using its RAID configuration utility at the BIOS level (or `raidctl` at the OS level for Solaris) to create hardware RAID volumes is *optional*. Multipath behavior of disks in these volumes is hidden from the OS by the SAS adapter.

If you decide to use an OS-based software RAID, you can use multipathing software at the OS level for your disk blade drives.

- With an Adaptec host bus adapter (such as the Sun Blade RAID 5 Expansion Module), using its RAID configuration utility at the BIOS level to create array volumes (for RAID or non-RAID installations) is *mandatory*. Consequently, you cannot use multipathing software at the OS level.

Multipathing With RHEL 5.x Linux and LSI Host Bus Adapters

The RHEL 5.x linux utilities for multipathing are the `multipath` utility and the `multipathd` daemon. Ensure that you have the minimum supported driver version for your LSI HBA listed in [“Updating Operating System Drivers on x86 Server Blades” on page 34](#).

Note – Due to limitations of the OS, you cannot install your OS onto multipathed disks with RHEL 4.x or SLES 9 or 10. With these Linux OS versions, you have two alternatives: 1) Install your OS on a disk or RAID volume on your server’s on-board disks (which are single path). OR, 2) Create a RAID volume with the SAS host bus adapter BIOS configuration utility and install your OS on it.

Installing the OS on Multipath Disks

With RHEL 5.x you can install your OS on dual-pathed SAS disks (requires two installed SAS-NEMs). By passing the `mpath` parameter during boot, you enable the multipath daemon (`multipathd`) during installation. You see any disks on the server as `sda`, `sdb`, and so forth. Disks on the server blade have only one path. You see the disks on the disk blade, which have two paths, as `mpath0` through `mpath7`.

You can choose to install the OS onto any of the multipathed disks.

▼ To Configure Multipathing After Installation

For all Linux operating systems, you can configure multipathing after your OS is installed. To do this, you run the `multipathd` daemon:

1. Find the `multipath.conf` file and edit it to remove the “blacklisting” of the disks you want to see.

You should blacklist the disk drives on your server blade because they do not have dual paths. Add devnode "`sd[a-d]`" to the blacklist and remove devnode `sd*` if it is present.

2. Run the `multipathd` daemon as root.

```
# service multipathd start
```

3. The multipathed disks are seen as `/dev/mapper/mpath*`. Partition them with `fdisk`.

For example, to partition `/dev/mapper/mpath0`, run

```
# fdisk /dev/mapper/mpath0
```

After the multipathed disks are partitioned, they can be used as regular disks.

Note – Henceforth, the multipath daemon must start at bootup. To start the daemon for run levels 1, 2, 3, 4, and 5 enter:

```
# chkconfig --level 12345 multipathd on
```

Multipathing With the Solaris OS and LSI Host Bus Adapters

The Sun Blade 6000 Disk Module supports multipathing with Solaris 10 05/08 and later Solaris 10 updates.

If you are using Solaris 10 05/08, you cannot manage disks with both the `mpxio` and `raidctl` utilities at the same time on Solaris systems. You must choose one or the other for managing your disks.

- You can run `raidctl` first to make volumes and then run `mpxio`, but `mpxio` does not interact with the volume you created with `raidctl`. The OS sees the RAID volumes (for example, as `c0t1d0`, where 1 is the target ID of the RAID volume), but the RAID volumes are not enumerated under `mpxio`.
- When `mpxio` is running, `raidctl` doesn't see any of the disks that `mpxio` is managing.

▼ To Install Multipathing for the Solaris OS

1. Before enabling multipathing, use the `format` command to identify the SAS drives by target number, as in the following example. For systems with two SAS-NEMs installed, there are two entries for each SAS drive.

```
# format
```

Your output will be similar to the following.

```
Searching for disks...done
AVAILABLE DISK SELECTIONS:
  0. c2t0d0 <DEFAULT cyl 8921 alt 2 hd 255 sec 63>
    /pci@0,0/pci10de,5d@c/pci1000,1000@0/sd@0,0
  1. c2t1d0 <DEFAULT cyl 8921 alt 2 hd 255 sec 63>
    /pci@0,0/pci10de,5d@c/pci1000,1000@0/sd@1,0
  2. c2t2d0 <DEFAULT cyl 8921 alt 2 hd 255 sec 63>
    /pci@0,0/pci10de,5d@c/pci1000,1000@0/sd@2,0
  3. c2t3d0 <DEFAULT cyl 8921 alt 2 hd 255 sec 63>
    /pci@0,0/pci10de,5d@c/pci1000,1000@0/sd@3,0
```

2. To enable multipathing, enter the command:

```
# stmsboot -e
```

Your output will be similar to the following. Answer the prompts as shown.

```
WARNING: stmsboot operates on each supported multipath-capable
controller detected in a host. In your system, these controllers are
```

```
/devices/pci@0,0/pci10de,5d@c/pci1000,1000@0
```

```
If you do NOT wish to operate on these controllers, please quit
stmsboot and re-invoke with -D { fp | mpt } to specify which
controllers you wish to modify your multipathing configuration for.
```

```
Do you wish to continue? [y/n] (default: y) y
```



```
Checking mpxio status for driver fp
Checking mpxio status for driver mpt
WARNING: This operation will require a reboot.
Do you want to continue ? [y/n] (default: y) y
updating /platform/i86pc/boot_archive...this may take a minute
The changes will come into effect after rebooting the system.
Reboot the system now ? [y/n] (default: y) y

The system reboots twice in succession. The first reboot automatically configures
the multipath devices and reboots again.
```

3. **At the Solaris prompt, use the `format` command to identify the drives by world-wide identifier, as in the following example. For systems with two SAS-NEMs installed, there is only a single entry for each SAS drive.**

```
# format
```

Your output will be similar to the following.

```
Searching for disks...done
AVAILABLE DISK SELECTIONS:
    0. c3t5000C50000C29A63d0 <DEFAULT cyl 8921 alt 2 hd 255...>
       /scsi_vhci/disk@g5000c50000c29a63
    1. c3t5000C50000C2981Bd0 <DEFAULT cyl 8921 alt 2 hd 255...>
       /scsi_vhci/disk@g5000c50000c2981b
    2. c3t5000C50000C2993Fd0 <DEFAULT cyl 8921 alt 2 hd 255...>
       /scsi_vhci/disk@g5000c50000c2993f
    3. c3t500000E016AA4FF0d0 <DEFAULT cyl 17845 alt 2 hd 255...>
       /scsi_vhci/disk@g500000e016aa4ff0
```

Note – By default when the Solaris OS is installed, multipath IO (MPxIO) to disks on a disk blade is disabled. When this feature is enabled, the `load-balance` variable in the file `/kernel/drv/scsi_vhci.conf` defaults to `round-robin`. It should be reset to `none`.

Setting `load-balance="none"` will cause only one path to be used for active IO, with the other path used for failover.

A serious performance degradation will result if the `load-balance` variable is left set to `round-robin`, because that would result in IO being attempted on the passive path.

4. **To disable multipathing, run the command `stmsboot -d`, and follow the prompts.**

Multipathing With Windows Server 2008 and LSI Host Bus Adapters

The Windows uses the MPIO utility for multipathing. Support for multipathing is currently only available for Windows Server 2008.

▼ To Install Multipathing for Windows Server 2008

An MPIO utility is part of Windows Server 2008, but you must install it manually:

1. **Click Start --> Administrative Tools and select Server Manager.**
2. **Select Features in the left-hand tree.**
3. **Click Add Features.**
4. **Check the Multipath I/O check box and click Next.**
5. **Click Install.**
6. **When installation is finished, click Close and reboot.**
7. **Click Start and select Control Panel. Verify that MPIO properties is present.**

▼ To Select Disks for Multipathing

After MPIO is installed, you must inform Windows of which types of disks should be multipathed:

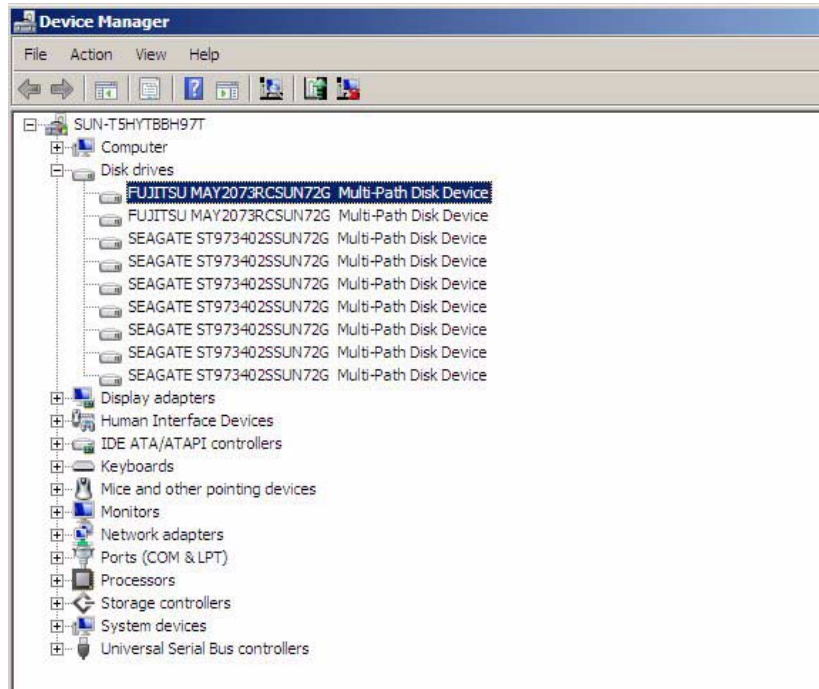
1. **Click Start and select Control Panel. Double-click MPIO properties.**
2. **Select the Discover Multi-Paths tab.**
You see a list of all the disk drives by manufacturer and type.
3. **Select the disk drives to multipath and click Add.**
4. **Click OK and reboot.**

At this point, MPIO is installed and configured for dual paths. The second path is hidden by Windows. By default, MPIO is configured for active passive path failover.

You can view the disks seen by the OS in several ways.

To configure MPIO properties for each disk, click Start and select Control Panel. Double-click the Device Manager. You see the view of the disks in [FIGURE 6-1](#).

FIGURE 6-1 Dual-Path Disks as Seen in the Device Manager

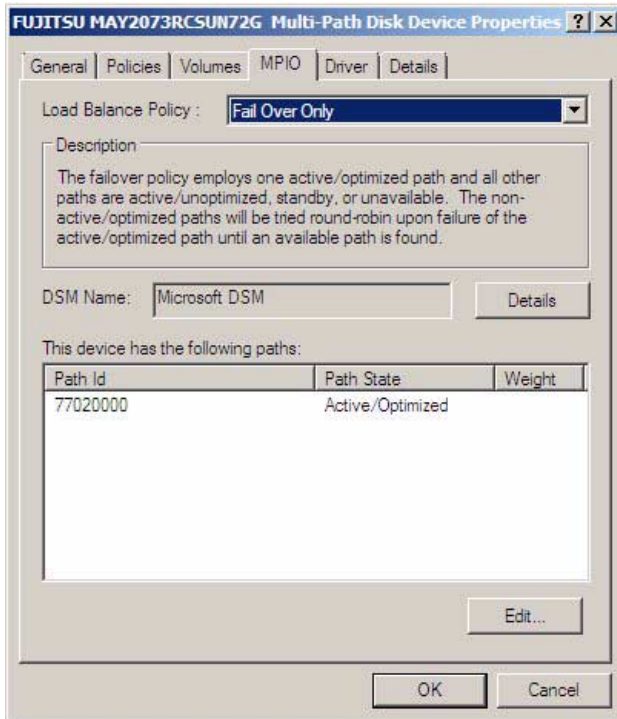


▼ To Set Load Balancing Property to Failover Only

Next, you must make sure that the Load Balancing property for each disk is Failover Only.

- 1. Right-click a disk drive and select Properties.**
- 2. Select the MPIO tab.**
- 3. Choose Failover Only from the Load Balance Policy drop-down list.**
- 4. Repeat steps 1 through 3 for each disk drive in the Device Manager window.**

FIGURE 6-2 Multi-Path Disk Device Properties Dialog Box



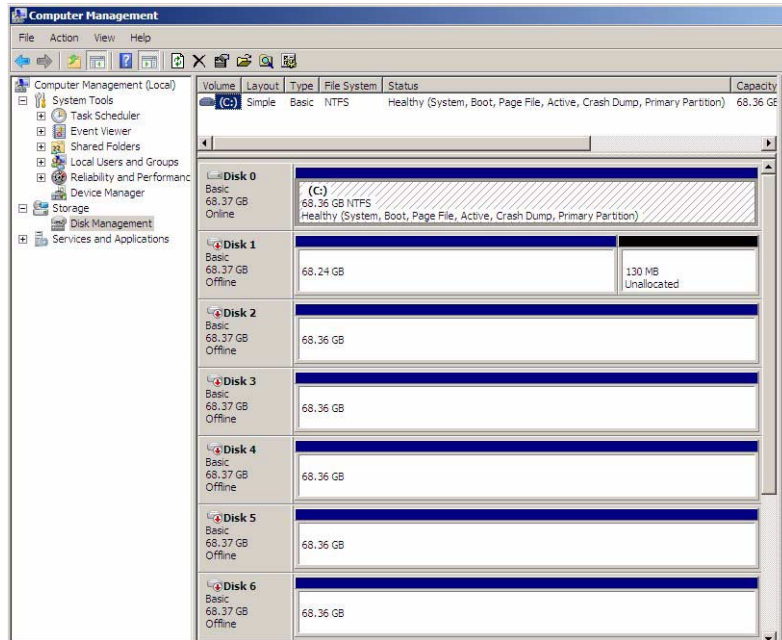
▼ To View a Single Instance of MPIO Disk Drives

To view a single instance of MPIO disk drives in the Windows Disk Manager:

1. **Click Start and select Run.**
2. **Type `diskmgmt.msc` in the Open dialog box.**

Windows Disk Manager starts and displays a single instance of each MPIO disk drive.

FIGURE 6-3 Dual-Path Disks as Seen in the Disk Management View



As before, Windows hides the second path to each SAS disk drive.

Failover

If you remove one path to a SAS disk by removing one of the two SAS-NEMs in your chassis, failover is handled automatically by MPIO. You continue to see the exact same list of disks in the Windows Device Manager and Windows Disk management views.

Multipathing with VMware ESX 3.5 and LSI Host Bus Adapters

To support multipathing, ESX Server 3.5 does not require specific failover drivers. After ESX is installed, it scans the dual SAS-NEM paths to the drives. The default failover policy is *active/active* and works as follows:

- NEM 0 is the default active path to disk blade drives (the path used by the host to communicate with storage devices)
- NEM 1 is the failover path which is unused until needed

If failover of the active (primary) path occurs, the SAS host bus adapter re-establishes contact through the secondary path without user intervention.

Note – For active/active failover on the disk blade-connected ESX Server host, the **Fixed** multipathing policy is required (**Round Robin** is not supported).

▼ To View Status and Configure Multipath Storage Devices

Use the VMware Virtual Infrastructure (VI) Client to show status for the storage devices.

1. **Login to the VI Client and select a server from the inventory panel.**
2. **Click the Configuration tab, and then Storage Adapters from the left pane.**

FIGURE 6-4 shows paths to SCSI targets in the disk blade. Each disk contains two SCSI targets because each drive has two SAS ports. In this example, SCSI Targets 5 and 15 belong to the same disk. SCSI Target 5 (vmhba0:5:0) is the primary, or canonical, path to the disk. SCSI Target 15 (vmhba0:15:0) provides a redundant path in the event that the primary path fails.

The canonical name for a disk is the first path (and the lowest device number) ESX finds to the disk after a scan.

FIGURE 6-4 Storage Adapters Page

The screenshot displays the 'Storage Adapters' page. On the left, there are navigation tabs for 'Getting Started', 'Summary', 'Virtual Machines', 'Resource Allocation', 'Performance', 'Configuration', 'Users & Groups', 'Events', and 'Permissions'. Below these are two main sections: 'Hardware' and 'Software'. The 'Storage Adapters' section is active, showing a table with one entry: 'vmhba0' of type 'Block SCSI'. Below this, the 'Details' section for 'vmhba0' is shown, including the model 'LSI1068E', 7 targets, and a list of SCSI targets with their paths and capacities.

Device	Type	SAN Identifier
LSI1068E		
vmhba0	Block SCSI	

Details

vmhba0
 Model: LSI1068E
 Targets: 7

SCSI Target 2

Path	Canonical Path	Type	Capacity	LUN ID
vmhba0:2:0	vmhba0:2:0	disk	68.37 GB	0

SCSI Target 5

Path	Canonical Path	Type	Capacity	LUN ID
vmhba0:5:0	vmhba0:5:0	disk	68.37 GB	0

SCSI Target 6

Path	Canonical Path	Type	Capacity	LUN ID
vmhba0:6:0	vmhba0:6:0	disk	68.37 GB	0

SCSI Target 8

Path	Canonical Path	Type	Capacity	LUN ID
vmhba0:8:0	vmhba0:8:0	disk	68.37 GB	0

SCSI Target 15

Path	Canonical Path	Type	Capacity	LUN ID
vmhba0:15:0	vmhba0:5:0	disk	68.37 GB	0

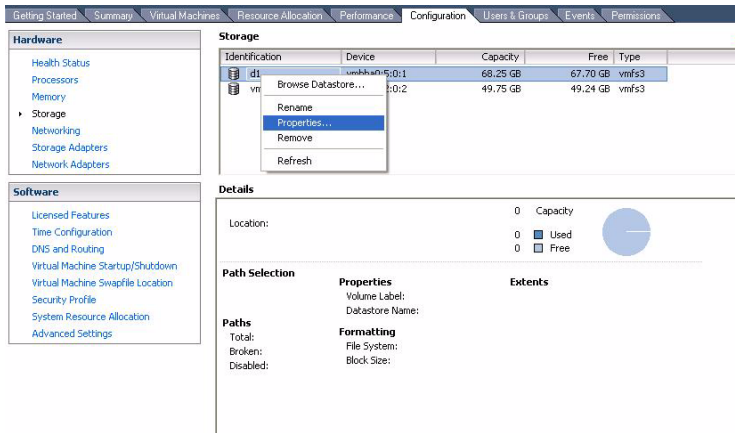
SCSI Target 16

Path	Canonical Path	Type	Capacity	LUN ID
vmhba0:16:0	vmhba0:6:0	disk	68.37 GB	0

Note – When creating a VMFS datastore using the Add Storage Wizard, only the canonical paths are displayed.

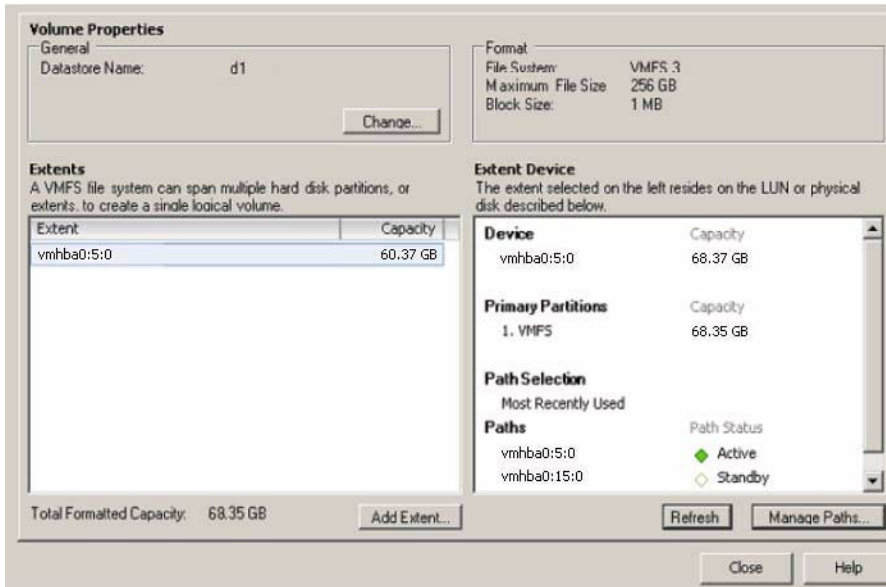
3. To manage multipath settings for a storage device, click on Configuration tab, and then click on Storage from the left pane.
4. From the Storage page, right-click on a storage device, and then click Properties from the drop-down menu (see FIGURE 6-5).

FIGURE 6-5 Storage Device Page



5. From the properties dialog box you can enable/disable paths, change failover policy, and view the failover status of any path.

FIGURE 6-6 Volume Properties Dialog Box



Follow the instructions described in ESX Server 3 Configuration Guide at:
http://www.vmware.com/support/pubs/vi_pages/vi_pubs_35.html

How LSI Host Bus Adapters Assign Target IDs to Expanders, Disks, and RAID Volumes

This chapter contains these topics:

- [“About Target IDs With LSI Controllers” on page 69](#)
- [“Target ID Mapping With LSI Controllers” on page 70](#)
- [“Assigning Volume IDs With LSI Controllers” on page 71](#)

Terms:

- SAS-NEMs and blades are in “slots” in the Sun Blade 6000 chassis.
- Disks are in “bays” in server blades or disk blades, numbered from 0 (lower left) to 7 (upper right).

About Target IDs With LSI Controllers

When you are using expanders (for example, the expanders on SAS-NEMs and disk blades), each disk has a “target ID.” These are the IDs that are used by the server blades’ SAS host bus adapters to identify the disks and that the adapter presents to the BIOS and OS.

When you have an LSI SAS host bus adapter, target IDs are determined by “enclosure slot mapping.” This occurs automatically and means that the target ID of a disk is determined by the disk bay (0-7) it is in. If a disk moves to another bay, its target ID changes: it assumes the ID of the bay it moves into.

Expanders are also considered “targets” and have a target ID. Each SAS-NEM has one expander. Each disk blade has two. One of the expanders on each disk blade connects to the SAS controller through one SAS-NEM expander. If there is only one SAS-NEM in the chassis, only one of the expanders on a disk blade is a target.

Allowed Range of Target IDs for LSI Controllers

In the Sun Blade Modular System chassis, any number from 0 through 113 can be used as an ID for any target in the chassis.

Target ID Mapping With LSI Controllers

The mapping of target IDs might seem random and confusing if you do not know how target IDs are assigned.

When a server blade is booted up in a chassis for the first time, its SAS host bus adapter performs SAS discovery, following this algorithm:

1. It finds the bays for its own on-blade disks (*always* target IDs 0 through 3).
2. Next it discovers all the SAS-NEM expanders that are currently in the chassis.
3. Next, the controller discovers all the disk blades connected through SAS-NEM 0. The discovery algorithm starts with disk blade 9, then 7, then 5, 3, and 1.

When any disk blade is found, whether or not it is paired with the controller’s host server blade, the controller reserves nine IDs: eight for the disks and one for the disk blade expander connected to SAS-NEM 0. These IDs are reserved for path 1, the first of the dual paths that goes through SAS-NEM 0.

Note – The nine IDs are reserved whether or not there are disks in the bays.

4. After the controller has discovered and reserved IDs for all the disk blades through path 1 (NEM 0), it goes on to discover disk blades through path 2 (NEM 1).
Once again, it reserves nine IDs for each disk blade discovered through this path, starting with disk blade 9, then 7, then 5, then 3, then 1.
5. The controller stores a persistence map of the target IDs assigned in its nonvolatile memory.

The persistence map is retained by the controller until a new discovery is requested. This map can also be stored as a file so that in the event of an HBA failure, the configurations can be restored to the new HBA (see [Appendix B](#) for details).

Note – The controller reserves a set of nine IDs for each path to each disk blade, including the one that is paired with its server, but it does not show them all in a SAS discovery list. The list shows all the expanders, but it only shows disks that are *physically present* in bays in its *own* paired disk blade. It does not show disks that are present in any other disk blade, and it does not show disks in empty bays in its own disk blade. When disks are added to such empty bays later, however, it will show them. The IDs of all the bays, whether they are occupied or not, are maintained in the persistence map.

In a paired zoning configuration, the server’s SAS controller can see at most:

- Four on-board disks
- Two SAS-NEM expanders
- Ten disk blade expanders (two per blade, for up to five blades)
- Sixteen disks in its paired disk blade: eight disks through SAS-NEM 0 and the *same* eight disks through SAS-NEM 1, providing a dual path to each disk (for more on multipathing, see [Chapter 6](#)).

Note – All LSI controllers use the same discovery algorithm, but the controller on each server blade has its own view of the targets. There is no communication between the controllers.

Assigning Volume IDs With LSI Controllers

When a RAID volume is created, it is given the same ID as the lowest target ID of its member disks. For example, if disks 3 and 7 are used to create a RAID 1 mirror, the volume has an ID of 3.

Note – Recall that the disk IDs are actually the IDs of the bay the disk is in.

The RAID volume ID is stored in metadata on the disks in the volume and does not change, even if the member disks change through failure and rebuilding to a hot-swap disk.



Caution – Never insert or remove a disk in a RAID volume when it is resynching.

Collision With Rebuilt Volumes

If the lowest ID disk in a mirrored RAID volume fails and is replaced with a hot spare, you can have a situation where the volume ID is not the same as the lowest ID of the RAID members. For example, consider a mirror that is composed of disks 3 and 7 with disk 4 as a hot spare. The volume ID is 3. If disk 3 fails and the mirror is rebuilt with disk 4, volume 3 (the volume ID that is written in metadata persists) now is composed of disks 4 and 7. If a new disk is inserted in bay 3, the SAS controller does not recognize it because it does not accept a volume and a disk with the same ID unless the disk is in the volume.



Caution – Do not insert another disk in an empty slot belonging to a RAID volume until the volume has been rebuilt (resynched). Rebuilding can take hours, depending on the size of the disks.

If (after resynching is completed) you insert a new disk in the bay that has the same target ID as the volume, the new disk is assigned the lowest target ID that is presently unused (including that of an empty bay) and will be added to the mirror as a hot spare.

In such a case, you have a disk that does *not* have the same target ID of the bay it is in, but because it is a hot spare, it is hidden by the SAS host bus adapter.

How Adaptec Host Bus Adapters Assign Target IDs to Expanders, Disks, and RAID Volumes

The Sun Blade RAID 5 Expansion Module is a SAS host bus adapter that is based on an Intel chip set with Adaptec firmware. This chapter refers to it as the Adaptec host bus adapter, Adaptec disk controller, or just the Adaptec controller.

This chapter contains these topics:

- [“About Adaptec Host Bus Adapters” on page 73](#)
- [“Target ID Mapping With Adaptec Host Bus Adapters” on page 74](#)
- [“Assigning Volume IDs With Adaptec Host Bus Adapters” on page 77](#)

About Adaptec Host Bus Adapters

With the Adaptec host bus adapter disks may not be used by the BIOS or OS unless they have been made members of an array. With the Sun Blade RAID 5 Expansion Module, an Adaptec host bus adapter, you *must* include every disk in an array volume if you want the OS to see it.

You can host bus adapter’s BIOS utility called the Adaptec RAID Configuration Utility (ARCU) to create arrays. This utility is entered by pressing Ctrl-A when prompted during the server boot process. Refer to the *Sun Intel Adaptec BIOS RAID Utility User’s Manual* (820-4708) for details.

FIGURE 8-1 shows the opening screen of the ARCU. There are three options:

- **Array Configuration Utility:** Use this for creating and managing arrays (including single-disk arrays), initializing disk drives, and re-scanning disk drives.

- **SerialSelect Utility:** Use this for modifying adapter settings.
- **Disk Utilities:** Use this option for viewing all drives, verifying their existence, and formatting them.

FIGURE 8-1 Opening Screen of the Adaptec RAID Configuration Utility (ARCU)



The ARCU can create up to 24 volumes (arrays). If a volume contains more than one disk, it is created as a RAID volume. You can also create a volume that contains only one disk. This is the mechanism you use to present a single disk to the OS.

Note – The volumes created by the ARCU take into account dual paths to SAS disks and hide the dual paths from the OS. See [“Multipathing Using Adaptec Host Bus Adapters” on page 57](#). Because there are no disks presented to the OS that are not in volumes, you *cannot* manipulate dual paths at the OS level.

Target ID Mapping With Adaptec Host Bus Adapters

The Adaptec host bus adapter uses slot-based mapping. This means that a disk’s target ID is the ID of the bay it is in. If a disk is moved to a different bay, the adapter changes the disk’s ID to that of the bay it moves into.

The adapter assigns target IDs to disks using two parameters:

- The enclosure number, which is labeled the “Box” number.
- The bay that the disk is in, which is labeled the “Slot” number.

For example, a disk in bay 3 of a disk blade might have a target ID of Box04:Slot03, where Box04 refers to the disk blade as an enclosure.

The bays in a disk blade are numbered like this:

TABLE 8-1 Disk Blade Bay Numbering (as seen from front of chassis)

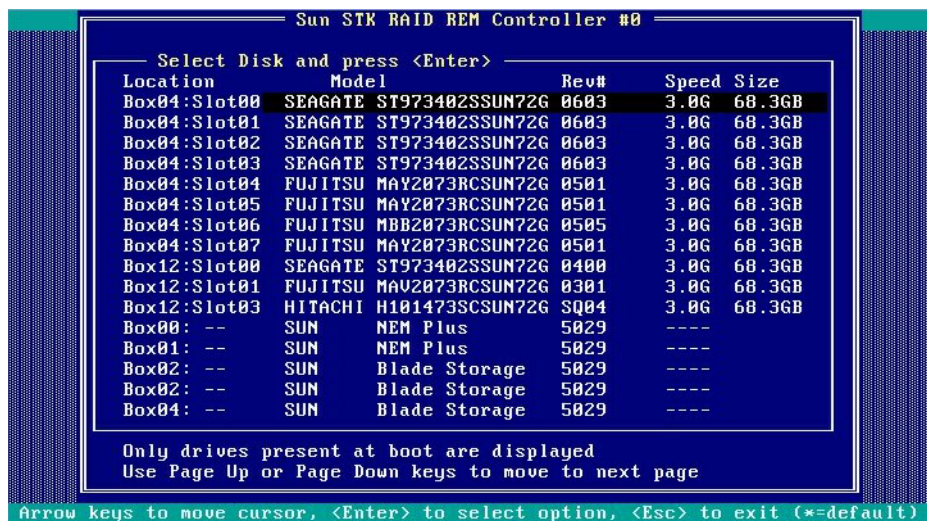
6	7
4	5
2	3
0	1

You can view the enclosure IDs and disk IDs of all the SAS targets that the Adaptec adapter sees by using the following tools:

- The Adaptec RAID Configuration Utility (ARCU). See the *Sun Intel Adaptec BIOS RAID Utility User’s Manual* (820-4708) for a description.
- The Sun StorageTek RAID Manager software. See the *Sun StorageTek RAID Manager Software User’s Guide* (820-1177) for a description.

FIGURE 8-2 shows a list of disks and enclosures seen by the Adaptec controller (a Sun Blade RAID 5 Expansion Module) on an X6250 server blade, using the ARCU’s Disk Properties option.

FIGURE 8-2 Disks and Enclosures Seen by an Adaptec Controller on an X6250 Blade



The listed targets are:

- `Box04:Slot00` through `Box04:Slot07` are the disks in bays 0 through 7 of the disk blade that is paired with the X6250 server blade. The disk blade is an enclosure with the label `Box04`.
- `Box12:Slot00` through `Box12:Slot03` are the disks in bays 0 through 3 of the X6250 server blade. (The server blade, as an enclosure, is labeled `Box12`.)

Note – If a disk bay is empty (for example, `Box12:Slot02`) no disk is reported by the ARCU. If a disk is inserted in this bay at a later time, it becomes `Box12:Slot02`.

- `Box00` and `Box01` are the two SAS-NEMs in the Sun Blade 6000 chassis. Note that the Multi-Fabric NEMs are identified as “NEM Plus” in the above screen example.
- `Box02` is another disk blade (“Blade Storage”) somewhere else in the Sun Blade 6000 chassis.

We can infer from this list that the chassis holds two SAS-NEMs (`Box00` and `Box01`) and two disk blades (`Box02` and `Box04`), one of which (`Box04`) is paired with the X6250 server blade.

The enclosure numbers (`Box00`, etc.) are assigned when the X6250 server blade is booted up for the first time and the Adaptec controller scans for SAS devices.

The Adaptec host bus adapter sees:

- All of the enclosures in the chassis. This includes SAS-NEMs, disk blades, and server blades if they can have on-board disks.
- Some, but not all, of the disks in the chassis. The disks on its own server blade and those on a disk blade that is paired with the server blade are seen. The disk numbers used are always `SLOT00`, `SLOT01`, `SLOT02`, . . . Disks that are not in either the server blade or its paired disk blade are not seen.

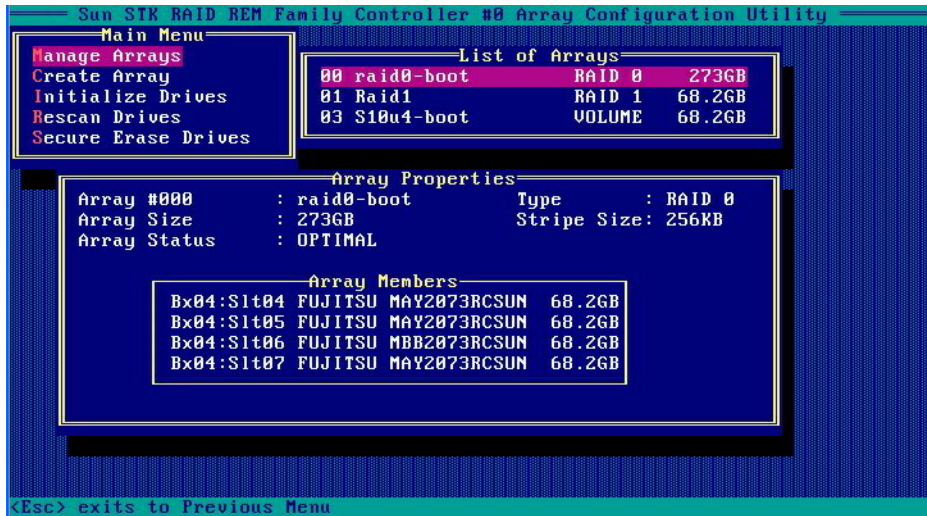
Assigning Volume IDs With Adaptec Host Bus Adapters

Volume IDs (Array IDs) are assigned in the order in which they are created with the ARCU. [FIGURE 8-3](#) shows the opening screen of the ARCU’s Array Configuration Utility option. You can see that four arrays have been created, 00 through 03. Subsequently, `Array 02` was deleted or has become inactive and is no longer listed.

If Array 02 was deleted, the next array to be created will be labeled 02 (the lowest available number). However, if the array still exists but is inactive, it will retain the label 02 and a new array will become Array 04.

Note – In naming new arrays, the lowest available number is always used.

FIGURE 8-3 Opening Screen of the ARCU's Array Configuration Utility Option Showing the Contents of Array 00



Array 00 is highlighted and you can see that it comprises four disks in a RAID 0 array (volume). The disks are Box04:Slot04 (abbreviated Bx04:Slot04) through Box04:Slot07.

FIGURE 8-4 shows the contents of Array 01, a RAID 1 mirror

FIGURE 8-4 Contents of Array 01

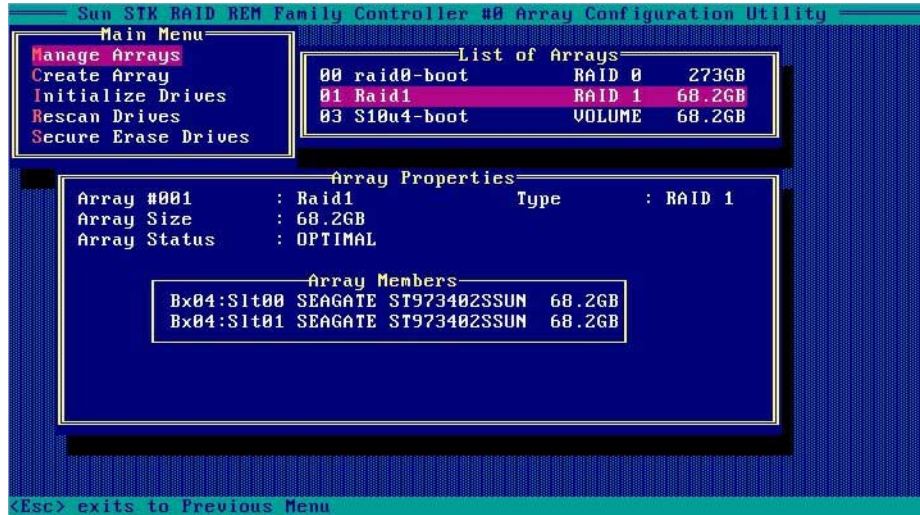
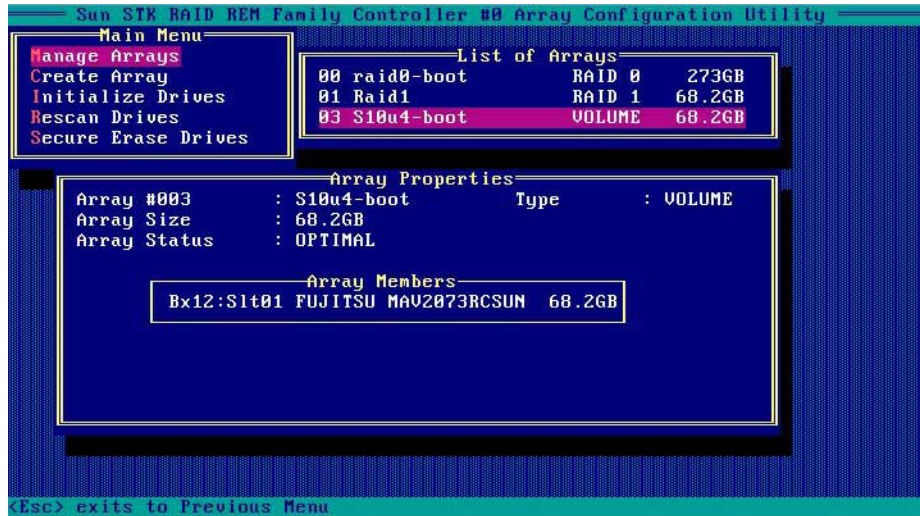


FIGURE 8-5 shows the contents of Array 03, a single disk volume.

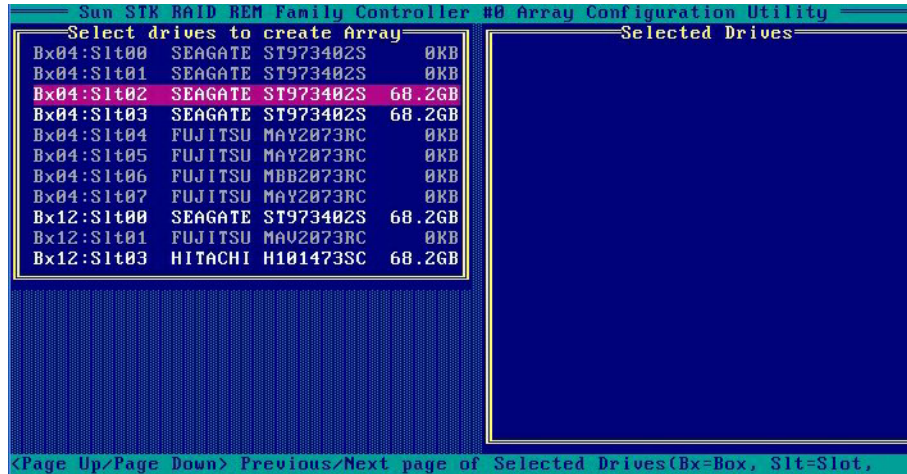
FIGURE 8-5 Contents of Array 03



Looking at the preceding figures, you can see that the three existing arrays include only seven of the eleven disks listed in FIGURE 8-2. Because all disks must be in arrays, you need to put the other four in arrays (volumes) if you want your OS to see them.

Highlight the **Create Array** menu item in the Array Configuration Utility screen and press **Enter**. The next screen that opens shows you the entire list of drives.

FIGURE 8-6 Select Drives to Create an Array



The drives that are dimmed are already in arrays. You can select one or more of the others to create additional arrays. The procedure is outlined in “Setting Up Hardware RAID With the Adaptec Host Bus Adapters” on page 91.

Setting Up Hardware RAID Volumes With LSI and Adaptec Host Bus Adapters

This chapter contains these topics:

- [“RAID Options” on page 81](#)
- [“Setting Up Hardware RAID With LSI Host Bus Adapters” on page 83](#)
- [“Setting Up Hardware RAID With the Adaptec Host Bus Adapters” on page 91](#)

RAID Options



Caution – Possible data loss: You must not insert a disk that has been previously configured with a RAID volume into a new RAID volume. Before permanently removing a disk that is part of an active RAID volume, use the appropriate RAID tool to delete the RAID volume from the disk.

The server modules that interoperate with the disk modules all have SAS host bus adapters that are capable of creating hardware RAID volumes.

The X6220 and T6300 server blades have an LSI 1068E SAS host bus adapter chip (with LSI firmware) on the blade’s motherboard.

Some server blades (see [“Supported RAID Controllers and Host Bus Adapters \(HBA\)” on page 18](#)) can use a RAID Expansion Module (REM) that uses the same LSI 1068E chip. This REM card is named the Sun Blade RAID 0/1 G2 Expansion Module.

Note – REMs are daughter cards that mount on the server blades.

Some server blades (see “Supported RAID Controllers and Host Bus Adapters (HBA)” on page 18) can use a REM card that uses Intel chips and Adaptec firmware. This REM card is named the Sun Blade RAID 5 Expansion Module.

All of these SAS host bus adapter options can control disks on the server blades (all the blades except X6450 can have on-board disks) and on the disk blade as well.

Both the LSI and the Adaptec controllers have two types of software for creating and managing RAID volumes. One type is a configuration utility that can be run from the BIOS before (or after) installation of an OS. The other type is a software application that runs on the server after an OS is installed.

TABLE 9-1 lists the various RAID possibilities.

TABLE 9-1 RAID Software

SAS Host Bus Adapter	Firmware	HBA BIOS Utility	RAID Manager Application
On-board LSI 1068E chip (SPARC server)	LSI	Not available for SPARC systems	<code>raidctl</code> for the Solaris OS
On-board LSI 1068E chip (x86 server)	LSI	LSI Logic Configuration Utility	MegaRAID Storage Manager (MSM). Runs on Linux and Windows OS only. <code>raidctl</code> for the Solaris OS
T6320 RAID 0/1 Expansion Module	LSI	Not available for SPARC systems	<code>raidctl</code> for the Solaris OS
Sun Blade RAID 0/1 G2 Expansion Module	LSI	LSI Logic Configuration Utility (x86 servers only)	MegaRAID Storage Manager (MSM). Runs on Linux and Windows OS only. <code>raidctl</code> for the Solaris OS (x86 and SPARC)
Sun Blade RAID 5 Expansion Module	Adaptec	Adaptec RAID Configuration Utility (x86 servers only).	Sun StorageTek RAID Manager. Runs on Linux, Windows, and Solaris OS (x86 and SPARC)

Documentation for this software can be found on (<http://www.docs.sun.com>).

- The LSI BIOS utility and MegaRAID Storage Manager are both documented in the *Sun LSI 106x RAID User's Guide* (820-4933).

- The Adaptec BIOS Utility is documented in the *Sun Intel Adaptec BIOS RAID Configuration Utility User's Manual* (820-4708).
- The Adaptec RAID Manager, also known as the Sun StorageTek RAID Manager, is documented in the *Sun StorageTek RAID Manager Software User's Guide* (820-1177).

Note – You can configure hardware RAID either before or after loading the OS on your server module, with one exception:

If you plan to create a RAID volume for your OS, the recommended procedure is to create the RAID first and then install the OS.

You can do this with the LSI Logic RAID Configuration Utility (see [“To Configure LSI RAID for x86 Servers With Any OS From the BIOS Before the OS Is Installed”](#) on page 85) or with the Adaptec RAID Configuration Utility (see [“Configuring Adaptec RAID for x86 Servers With Any OS from the BIOS”](#) on page 91).

For SPARC servers with LSI host bus adapters, you can use `raidctl` to create a RAID volume before installing the Solaris OS (see [“To Configure RAID for SPARC Systems Before Installing the OS”](#) on page 85).

Setting Up Hardware RAID With LSI Host Bus Adapters

Hardware RAID configurations can be set up with a server blade's SAS host bus adapter. The type of controller varies from server module to server module (see [“Supported RAID Controllers and Host Bus Adapters \(HBA\)”](#) on page 18).

The X6220 and T6300 server modules that can be used with the Sun Blade 6000 Disk Module have an on-board LSI 1068E chip. Other server modules can use a RAID Expansion Modules (REM) that uses the same LSI 1068E chip.

The LSI 1068E SAS host bus adapter (an integrated RAID controller) simultaneously controls the disk drives on the server module itself (except for the X6450 servers that have no on-board disks) and the eight drives on the disk module, and can create RAID configurations involving all twelve disks. RAID 0 (striping), RAID 1 (mirroring), and RAID 1E (enhanced mirroring on three or more disks) are all supported by the firmware on the LSI 1068E chip.

Deciding Whether to Configure RAID Before or After OS Installation

You can configure hardware RAID either before or after loading the OS on your server module, with one exception: if you plan to create a RAID volume for your OS, the recommended procedure is to create the RAID first and then install the OS.

- For x86 servers, depending on which HBA you have, you can do this either with the LSI Logic RAID Configuration Utility (see “To Configure LSI RAID for x86 Servers With Any OS From the BIOS Before the OS Is Installed” on page 85), or with the Adaptec RAID Configuration Utility (see “Configuring Adaptec RAID for x86 Servers With Any OS from the BIOS” on page 91).
- For SPARC servers with LSI host bus adapters, you can use `raidctl` to create a RAID volume before installing the Solaris OS (see “To Configure RAID for SPARC Systems Before Installing the OS” on page 85).

Configuring RAID for the Windows and Linux OS After the OS Is Installed

LSI has developed a software tool called the MegaRAID Storage Manager (MSM) for configuring and managing RAID on platforms running the Windows and Linux operating systems. Refer to *Sun LSI 106x RAID User's Guide* (document number 820-4933) for a detailed description of MSM. You can download this document from:

(<http://docs.sun.com/app/docs/coll/blade6000dskmod>)

The MSM software is only available *after* the OS is installed.

Configuring RAID for the Solaris OS After the OS Is Installed

Solaris includes a RAID hardware utility, accessed with the command: `raidctl`. This utility is used for LSI controllers and has functionality that is similar to that of the LSI BIOS utility. `raidctl` is also covered in the *Sun LSI 106x RAID User's Guide* (820-4933).

Note – The man page for `raidctl` states that the utility is for the LSI 1064 SAS host bus adapter, but it can also be used with the LSI 1068E controller.

▼ To Configure RAID for SPARC Systems Before Installing the OS

To create a RAID volume where you can install your Solaris OS:

1. **Boot your server from a remote Solaris OS image.**
2. **Use `raidctl` to create a mirror.**
3. **Label the RAID volume.**
4. **Power-cycle your server.**
5. **Install the Solaris OS on your server.**

▼ To Configure LSI RAID for x86 Servers With Any OS From the BIOS Before the OS Is Installed

The LSI RAID configuration utility that is entered from the server's BIOS is valid for all x86 server blades and all supported operating systems.

1. **Power-cycle your server module.**
The BIOS screen appears.
2. **Watch for the LSI Logic Corp. screen which comes up soon after the opening screen.**



3. **Press Ctrl-C to start the LSI Logic Configuration Utility.**

```

LSI Corp Config Utility      v6.20.01.00 (2008.03.13)
Adapter List Global Properties
Adapter      PCI  PCI  PCI  PCI  FW Revision      Status      Boot
              Bus  Dev  Fnc  Slot              0.24.30.00-IR   Enabled     Order
SAS1068E    02  00  00  00              0.24.30.00-IR   Enabled     0

```

4. With the LSI 1068E adapter highlighted in the first screen, press Enter.

The main screen of the utility opens

```

LSI Corp Config Utility      v6.20.01.00 (2008.03.13)
Adapter Properties -- SAS1068E

Adapter          SAS1068E
PCI Slot         00
PCI Address(Bus/Dev/Func) 02:00:00
MPI Firmware Revision 0.24.30.00-IR
SAS Address      50000200:00422470
NVDATA Version   2D.03
Status           Enabled
Boot Order       0
Boot Support     [Enabled BIOS & OS]

RAID Properties
SAS Topology
Advanced Adapter Properties

Esc = Exit Menu      F1/Shift+1 = Help
Enter = Select Item  -/+Enter = Change Item

```

5. Use the arrow keys to select RAID Properties and press Enter.

The next screen allows you to choose the type of RAID

```
LSI Corp Config Utility      v6.20.01.00 (2008.03.13)
Select New Array Type -- SAS1068E

Create IM Volume           Create Integrated Mirror Array of 2
                           disks plus up to 2 optional hot spares.
                           Data on the primary disk may be migrated.

Create IME Volume         Create Integrated Mirrored Enhanced
                           Array of 3 to 10 disks including up
                           to 2 optional hot spares.
                           ALL DATA on array disks will be DELETED!


Create IS Volume          Create Integrated Striping array of
                           2 to 10 disks.
                           ALL DATA on array disks will be DELETED!

Esc = Exit Menu           F1/Shift+1 = Help
Enter = Choose array type to create
```

6. Select the type of volume you want to create, in this case an IME volume.

The next screen lists the disks you can choose to include in the volume and also disks you can choose as hot spares.

Note – Despite the fact that there are dual paths to disks 1 - 7, they are not shown in the BIOS configuration utility. The utility hides them so that you cannot include the same disk twice in a volume. That is, you cannot mirror a disk with itself.



Note – Disks in the server blade are always listed first, followed by the disks in the disk blade. Here we see a single disk in the server blade (in bay 0) followed by eight disks in the disk blade (0 through 7).

- Use the arrow keys to highlight the [No] item in the RAID Disk column for the disks you want to include in the volume. For each such disk, press the spacebar or type the + or - key to change the [No] to [Yes].
- Use the arrow keys to highlight the [No] item in the Hot Spr column for the disks you want to use as hot spares. For each such disk, press the spacebar or type the + or - key to change the [No] to [Yes].

Your selections might look like this:

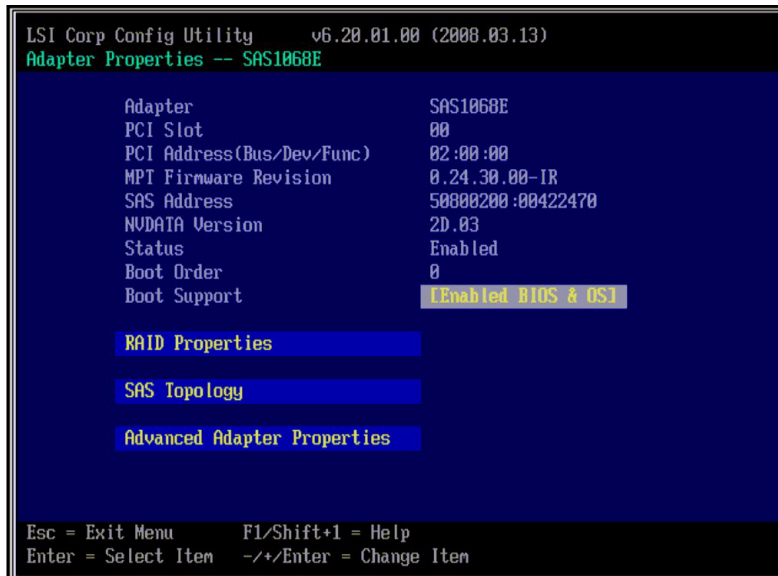
Slot Num	Device Identifier	RAID Disk	Hot Spr	Drive Status	Pred Fail	Size (MB)
0	SEAGATE ST973402SSUN72G 0400	[No]	[No]	-----	---	70007
0	SEAGATE ST973402SSUN72G 0400	[No]	[No]	-----	---	70007
1	SEAGATE ST973402SSUN72G 0400	[Yes]	[No]	-----	---	70007
2	SEAGATE ST973402SSUN72G 0400	[Yes]	[No]	-----	---	70007
3	SEAGATE ST973402SSUN72G 0400	[Yes]	[No]	-----	---	70007
4	SEAGATE ST973402SSUN72G 0400	[Yes]	[No]	-----	---	70007
5	SEAGATE ST973402SSUN72G 0400	[No]	[No]	-----	---	70007
6	SEAGATE ST973402SSUN72G 0400	[No]	[Yes]	Hot Spare	---	70007
7	SEAGATE ST973402SSUN72G 0400	[No]	[No]	-----	---	70007

9. When you are satisfied with your choices, type **c** to create the array.
A confirmation screen appears.



10. Select **Save changes then exit this menu**, and then press **Enter**.

The volume is created. When the utility finishes creating the volume, the main screen reappears.



11. Highlight RAID Properties and press Enter.
12. When the next screen appears, select View Existing Array and press Enter.
You see the volume you have created.

```

LSI Corp Config Utility      v6.20.01.00 (2008.03.13)
View Array -- SAS1068E
  Array                      1 of 1
  Identifier                 LSILOGICLogical Volume 3000
  Type                      IME
  Scan Order                 7
  Size(MB)                   139236
  Status                     Optimal

  Manage Array

Slot  Device Identifier      RAID  Hot  Drive  Pred  Size
Num  Num                    Disk Spr Status Fail (MB)
  1  SEAGATE ST973402SSUN72C 0400 Yes No  Ok    No    69617
  2  SEAGATE ST973402SSUN72C 0400 Yes No  Ok    No    69617
  3  SEAGATE ST973402SSUN72C 0400 Yes No  Ok    No    69617
  4  SEAGATE ST973402SSUN72C 0400 Yes No  Ok    No    69617
  6  SEAGATE ST973402SSUN72C 0400 No  Yes Hot Spare No    70007

Esc = Exit Menu      F1/Shift+1 = Help
Enter=Select Item  Alt+N=Next Array  C=Create an array  R=Refresh Display

```

Setting Up Hardware RAID With the Adaptec Host Bus Adapters

This section describes how hardware RAID configurations can be set up with a server blade's Adaptec SAS host bus adapter.

You can create up to 24 volumes (arrays) using either the HBA BIOS-based utility or the Sun StorageTek RAID Manager software. If a volume contains more than one disk, it is created as a RAID volume. You can also create a volume that contains only one disk. This is the mechanism you use to present a single disk to the OS.

Configuring Adaptec RAID After the OS Is Installed

The Sun StorageTek RAID Manager software is used for configuring and managing RAID on platforms running the Solaris, Linux, and Windows operating systems. Sun StorageTek RAID Manager Software and is documented in *Sun StorageTek RAID Manager Software User's Guide* (820-1177). You can download this document from (<http://docs.sun.com>.)

The Sun StorageTek RAID Manager software is only available *after* the OS is installed.

Note – The Adaptec BIOS utility is not available for SPARC systems so you must use the StorageTek RAID Manager to set up hardware RAID after the Solaris OS is installed.

Configuring Adaptec RAID for x86 Servers With Any OS from the BIOS

This section describes how to use the Adaptec RAID Configuration Utility (ARCU) to create a RAID array. You can also use the procedure in this chapter to create single-drive volumes. A volume, by definition, consists of a single drive.


On servers equipped with an Adaptec SAS host bus adapter, all drives must be part of a volume before they can be used by the BIOS or the operating system. Adaptec controllers support single-drive volumes that can be used for drives that are not part of a RAID array. This is especially important when you add drives to a system.

Note – The Adaptec ARCU GUI uses the term “array” for a RAID volume and the term “volume” for a single-disk volume. An Adaptec array is synonymous with an LSI volume. The concept of a single-disk volume does not exist with LSI adapters.

▼ To Create a RAID Array


1. **Power-cycle your server.**
2. **During power up, type Ctrl-A to bring up the ARCU.**

A message appears, and then the ARCU opens.



3. Select Array Configuration Utility.

The Array Configuration Utility appears



4. Select Initialize Drives.

A list of drives appears.



5. Select drives to initialize.

Use the arrow keys to scroll through the list and use the Insert key to select a drive.

Note – You can select many drives, and initialize them all at once, even if they are going to be in different volumes or arrays.

6. When you have selected all the drives to be initialized, press Enter.

A message warns you that initializing a drive erases array information from it.

7. Type *yes*.

The utility initializes the selected drives, and the main screen appears.



8. Select Create Array from the main menu.

A list of drives appears.



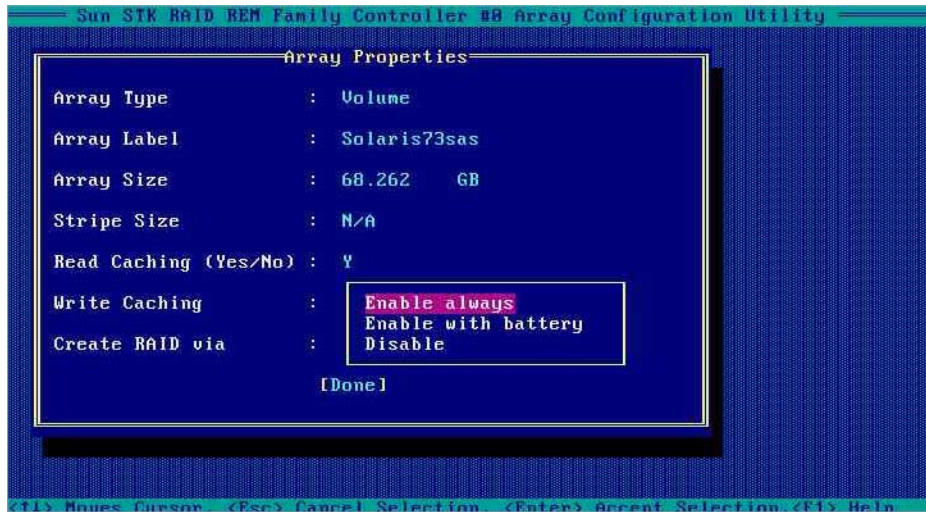
9. Select the drives to be included in the array or volume.

If you are going to create a one-disk volume, select the drive to be included in the volume. If you are going to create a RAID array, select all the drives to be included in the array. Repeat [Step 9](#) through [Step 13](#) for each subsequent array or volume.

Use the arrow keys to scroll through the list and use the spacebar to select a drive.

10. After you make your selections, press Enter.

The Array Properties view appears.




11. Make the following selections:

- Array Type – Select an array type from the drop-down list. If you selected a single drive, the array type is Volume.
- Array Label – Type a label.
- Stripe Size – Type a stripe size.
- Read Caching – Type **Y** or **N**.
- Write Caching – Select an option from the list.

12. Press Enter or click Done to proceed.

Depending on the Write Caching selection, a number of warnings might appear:

- If the Write Caching setting is Enable always, and if you do not have a battery, or if the battery is not charged sufficiently, several warnings might appear. Type **Yes** to proceed, or type **No** to return to the Array Properties screen.
- If the write cache is enabled, a write cache warning appears.



13. Click Enter to proceed.

The utility initializes the array.

ILOM For the Sun Blade 6000 Disk Module

This chapter contains these topics:

- “ILOM on the CMM” on page 97
- “ILOM 2.0 Proxy CLI Program For SAS-NEMs and Disk Modules” on page 98
 - “To Start the ILOM 2.0 Proxy CLI Program” on page 99
- “Navigation Using the CMM ILOM CLI” on page 99
- “CMM ILOM Management and Monitoring for SAS-NEMs and Disk Modules” on page 101
- “Upgrading CMM ILOM Firmware” on page 102

ILOM on the CMM

The Oracle Integrated Lights Out Manager (ILOM) on the Chassis Monitoring Module (CMM) provides minimal information for disk blades or SAS-NEMs.

The ILOM graphical user interface recognizes the presence of SAS-NEMs and disk blades but does not interact with them.

The ILOM CLI (command-line interface) can get FRU and SEEPROM information from both SAS-NEMs and disk blades, but does not provide any diagnostic information. Refer to the *Sun Blade 6000 Multi-Fabric Network Express Module User’s Guide* (820-1705) for details.

ILOM generates logs that record if a disk blade or SAS-NEM is inserted, removed, powered up, or powered down.

ILOM also monitors the Service LED on the Sun Blade 6000 Disk Module and SAS-NEM blades in the chassis. The Service LED is turned on for an over-temperature or over-voltage condition. If this light is turned on or off, the corresponding IPMI state sensor changes and a log entry is generated.

Note – A sudden over-temperature or over-voltage surge will power off the SAS-NEM or disk blade. When such a rapid shut down occurs, the Service LED is also turned off. However, this type of event is recorded in the log.

To see the CMM log for the SAS-NEMs and disk blades (as well as any server blades present), run this command from the CMM CLI:

```
-> show /CMM/logs/event/list
```

For ILOM version 2.0, if you want to monitor and interact with SAS-NEMs and disk blades you must use a special ILOM proxy program. ILOM 3.0 does not require a proxy program.

ILOM 2.0 Proxy CLI Program For SAS-NEMs and Disk Modules

If you are using CMM ILOM 2.0 (version 2.0.3.10 or later), you use a CLI-based proxy program to monitor and control your disk modules and SAS-NEM modules.

Note – There is no proxy CLI for ILOM 3.0.

You can connect to the ILOM CLI by either of two means:

- Connect a terminal or PC running a terminal emulator directly to the CMM serial port on your chassis.
- or
- Connect to the Ethernet network management port using a secure shell (SSH).

Instructions for setting up and using ILOM are documented in the *Sun Integrated Lights Out Manager 2.0 User's Guide*, available at:

(<http://docs.sun.com/app/docs/doc/820-1188>)

▼ To Start the ILOM 2.0 Proxy CLI Program

Prerequisite: To start the ILOM 2.0 proxy CLI, a SAS-NEM module must be installed in the chassis in either NEM slot 0 or NEM slot 1.

- **Connect to the ILOM 2.0 CLI and then start the proxy CLI with one of the following commands:**

```
-> start /CH/NEM0/SAS/cli  
-> start /CH/NEM1/SAS/cli
```

Note – The proxy CLI commands are case-sensitive.

When you start the proxy CLI, you see a list of all the SAS-NEMs and disk blades in the system. For example:

```
-> start /CH/NEM0/SAS/cli  
Are you sure you want to start /CH/NEM0/SAS/cli (y/n)? y  
Found SAS-NEM in NEM slot 0  
Found SAS-NEM in NEM slot 1  
Found STORAGE in BL slot 1  
Found STORAGE in BL slot 3  
Found STORAGE in BL slot 7  
Welcome to proxy CLI on slot 0  
proxy ->
```

There are two NEM slots, 0 and 1, and there are 10 slots for blades like the disk module, 0 through 9. The SAS-NEMs in the chassis are referred to as NEM0 and NEM1, while the disk blades are referred to as BL0, BL1, ... BL9.

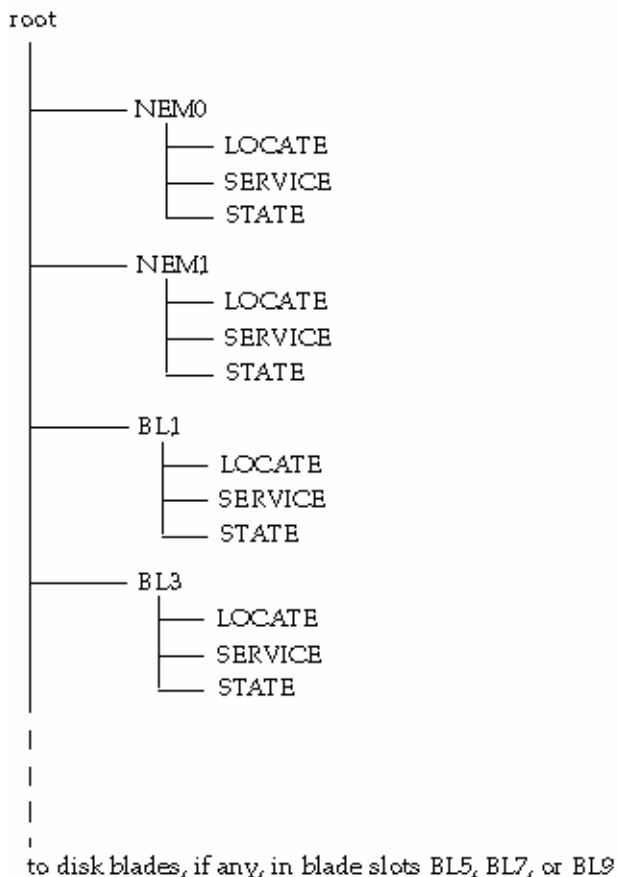
Note – Server modules in the chassis are not detected by the proxy CLI program.

Navigation Using the CMM ILOM CLI

With the CMM ILOM CLI (either with the proxy CLI for ILOM 2.0, or the standard CLI with ILOM 3.0), you navigate through a tree consisting of all the SAS-NEMs and disk modules (called *targets*) in the chassis. You can use standard Linux and Unix commands such as `cd` or `pwd` to navigate through the tree.

The tree consists of all the SAS-NEMs and disk blades in the chassis.

FIGURE 10-1 The CLI Tree of Targets



You can use these commands to navigate through the tree:

- `show` (with no argument): Shows contents of the current target. Lists targets directly beneath the current target, properties of the current target, and available proxy commands.
- `show` (with argument): Shows contents of the target specified. For example, `show BL3`
- `cd` (with argument): Changes current target to that of target specified. For example, `cd NEM1` or `cd ../NEM1`
- `pwd` (with no argument): Determines where you are in the tree of targets (because the prompt doesn't tell you).
- `cd /`: Returns to the root of the tree of targets.

You can use these commands on module targets (NEMx and BLx):

- `start`: Powers on the current target or the given target.
- `stop`: Powers off the current target or the given target.
- `reset`: Powers off, then powers on the current target or the given target.



Caution – The `reset` and `stop` commands do not check for host activity. You should only use these commands if you are sure that the device is not in use.

You can use this command on property targets:

- `set`: Changes the value of the specified property to the given value. For example, the Power LED reflects the `power_state` property, and you can change the `power_state` property with the `set` command.

You can use these commands at any time:

- `exit`: Leaves the proxy CLI and return to the ILOM CLI.
- `help`: Shows the help screen.
- `version`: Shows the current versions of the proxy program and the AMI MG9073 firmware.
- `load`: Loads new AMI MG9073 firmware.

CMM ILOM Management and Monitoring for SAS-NEMs and Disk Modules

The CMM ILOM CLI (ILOM 2.0 or 3.0) allows you to do the following:

- View and Navigate NEM and Sun Blade 6000 Disk Module chassis locations.
- View the values of the Locate and Service LEDs on the Sun Blade 6000 Disk Modules and SAS-NEMs. The Power and Service LEDs are both under control of the expander firmware running on the disk blade or SAS-NEM. You can view these LEDs from the CLI program, but you cannot change them directly.
- View the current version of AMI MG9073 firmware on Sun Blade 6000 Disk Modules.
- View IPMI sensor state values for the SAS-NEMs and Sun Blade 6000 Disk Modules.
- Stop and start the SAS-NEM or disk module by using the `stop` and `start` commands. You can achieve the same result by setting the `power_state` property to `off` or `on`: for example, set `power_state=off`.

- Reset the SAS-NEM or disk module.
- Light the Locate LED. The CMM ILOM CLI controls only one LED directly. You can turn the Locate LED on (value=fast blink) or off (value=off). ILOM will turn off the Locate LED after a specified timeout (default = 30 minutes).
- View and load new AMI MG9073s firmware to Sun Blade 6000 Disk Modules.

Upgrading CMM ILOM Firmware

ILOM 2.0.3.10, build 36968 for the Sun Blade 6000 System chassis is the minimum version required for use with disk blades and SAS-NEMs. This firmware can be downloaded from:

(<http://www.oracle.com/goto/blades>)

You can upgrade your ILOM firmware in two ways: by using the CLI or by using the ILOM web interface.

▼ To upgrade ILOM Firmware Using the CLI

Use a tftp server:

1. **Put the .ima file in the /tftpboot directory.**

2. **In the CLI, run this command:**

```
-> load -source tftp://<IPadd>/ilom.6000-2.0.3.10-r36968.ima
```

where *IPadd* is the IP address of your tftp server.

▼ To upgrade ILOM Firmware Using the ILOM Web Interface

1. **Log in as any user with Administrator privileges.**

2. **Select Maintenance --> Firmware Upgrade.**

The Firmware Upgrade page appears.

3. **Click Enter Upgrade Mode.**

A dialog box appears asking you to confirm that you want to enter Upgrade mode.

4. Click OK to enter Upgrade mode.

ILOM stops its normal operation and prepares for a flash upgrade.

5. Enter the path to the new ILOM .ima file in the Select Image File to Upload field or click Browse to locate and select the .ima file.

6. Click Upload.

After a short pause, a Firmware Verification screen appears that shows the current firmware version and the requested upgrade version.

7. Click Start Upgrade.

The process takes about 6 minutes. At the end, the service processor is reset.

Oracle's Sun Storage Common Array Manager (CAM)

This chapter contains the following sections:

- [“About Enclosure Management” on page 105](#)
 - [“Assigning Storage Resources” on page 106](#)
- [“CAM” on page 106](#)
- [“Obtaining CAM Software” on page 107](#)
- [“Using CAM With Disk Blades and SAS-NEMs” on page 108](#)
- [“Upgrading Expander Firmware Using CAM” on page 109](#)

About Enclosure Management

Your Sun Blade 6000 Disk Module supports a powerful set of enclosure management features accessible from a management client that supports SES-2 (SCSI Enclosure Services). The Sun Blade 6000 Disk Module expanders are SES-2 compliant. These enclosure management features are available through the management software, Oracle's Sun Storage Common Array Manager (CAM), which provides a system administrator with the following capabilities:

- Event and fault monitoring
- E-mail alert notification
- FRU identification and status
- Enclosure reset
- Enclosure firmware upgrade
- Fault isolation
- Service Advisor wizard for problem resolution

- Auto Service Request (ASR) uses fault telemetry 24/7 to automatically initiate a service request and begin the problem resolution process as soon as a problem occurs

Assigning Storage Resources

The storage resources of the Sun Blade 6000 Disk Module are assigned to a single blade server. Which server blade owns the disks is dependant on slot paring (described in [Chapter 1](#)). Although some versions of CAM support assigning storage to multiple hosts on a per-initiator basis (called “zoning”) for other storage products, the Sun Blade 6000 Disk Module cannot be configured to do this. Do not attempt to modify any SAS domain settings listed in CAM for the disk module.

CAM

Beginning with CAM version 6.1.2, CAM can be used to manage SAS-NEMs and disk blades in a chassis from one central point. It is written in Java™ code so you can run it on any platform.

CAM can monitor temperature and voltage on both the disk blade and the SAS-NEM. It can also display the topology of your system and report FRU IDs.

For the Sun Blade 6000 Disk Module, the most important functionality of CAM is firmware management. CAM can report the current versions of firmware and can upgrade the firmware on the disk blade, and the SAS-NEM. To obtain the latest CAM software go to:

(<http://www.oracle.com/us/products/servers-storage/storage/storage-software/031603.htm>).

In a Sun Blade 6000 chassis, one server blade can be designated to act as the CAM management host with the full version of CAM software. Alternatively, you can put the CAM software on any server on your network.

CAM Agent

The CAM Agent must be installed on every server blades designated as a data host that directly connect to a disk blade. Data hosts do not require the full version of CAM. The agent software is OS dependent. There are versions for Linux, Solaris, and Windows operating systems.

When a server blade connects to a disk blade, the CAM Agent plug-in is required for the CAM management host to recognize and manage the disk blade. With the CAM Agent plug-in, you can remotely manage the disk blade from the management host from anywhere on the network.

The CAM Agent communicates with the disk module through the server blade's SAS host bus adapter. Both LSI and Adaptec controllers are supported.

Problem Resolution

CAM provides two helpful tools for resolving problems that might occur with the system.

- **Auto Service Request (ASR)** monitors the system's health and performance and automatically notifies the Oracle Technical Support Center when critical events occur. Critical alarms generate an Auto Service Request case. These notifications enable Oracle Service to respond faster and more accurately to critical on-site issues.

To use ASR, you must provide Oracle online account information to enroll your CAM software to participate in the ASR service. After you enroll CAM with ASR, you can choose which systems you want to be monitored and enable them individually.

- **Service Advisor** is a troubleshooting wizard that provides information and procedures for replacing system components.

Obtaining CAM Software

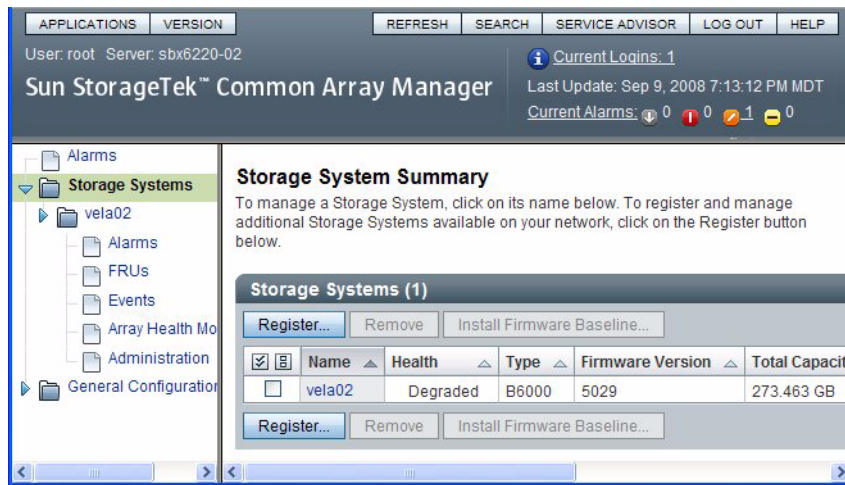
Support for Sun Blade 6000 Disk Modules and Sun Blade 6000 SAS-NEMs is available beginning with CAM version 6.1.2. To obtain the latest CAM software and expander firmware, go to:

(<http://www.oracle.com/us/products/servers-storage/storage/storage-software/031603.htm>)

Using CAM With Disk Blades and SAS-NEMs

CAM provides both a browser and command-line interface. When using the browser interface, you set up user accounts. When an authorized user logs in, information is laid out in pages. A handy navigation tree lists available information. See [FIGURE 11-1](#).

FIGURE 11-1 Example CAM Storage System Summary and Navigation Tree



You use the navigation tree on the left to move among pages within an application. You can click a link to get details about a selected item (such as alarms, FRUs, events, or array health). You can also sort and filter information displayed on a page. When you place your pointer over a button, tree object, link, icon, or column, a tool tip provides a brief description of the object.

Monitoring Component Health

CAM can monitor voltage and temperature on installed disk blades and SAS-NEMs and can raise alarms (including notification) when thresholds are exceeded. Choose the Array Health Monitoring item from the navigation tree to learn about this capability.

The monitoring capabilities of CAM are fully documented elsewhere. Refer to the *Sun StorageTek Common Array Manager User Guide* for your version of CAM at:

(<http://docs.sun.com/app/docs/prod/stor.armgr#hic>)

Note – CAM also has extensive online documentation that you can access using the Help button in the upper right of the browser interface.

Upgrading Expander Firmware Using CAM

Both Sun Blade 6000 Disk Modules and Sun Blade 6000 supported SAS-NEMs contain upgradable firmware. You should keep this firmware at the latest released versions. CAM includes firmware upgrade capabilities for these components. When using CAM to upgrade firmware, note the following:

- Sun Blade 6000 Disk Modules and Sun Blade 6000 supported SAS-NEMs must always be upgraded to the same firmware revision level. If required, CAM will upgrade them both.
- If your disk blade firmware is at a later version than CAM's baseline, CAM will still inform you that your firmware is not at baseline. If this is the case, **DO NOT** perform an upgrade. This will *downgrade* your disk blade firmware with the earlier version in CAM and could reduce functionality or disable the disk blade.

The firmware release history for the Sun Blade 6000 Disk Module and supported SAS NEMs is listed in the following table.

TABLE 11-1 Sun Blade 6000 Disk Module and SAS NEM Firmware

Firmware Revision	Base Firmware Delivery Platform	Supported Hardware	Description
5.02.10 (502A)	Common Array Manager 6.1.2	<ul style="list-style-type: none"> • Sun Blade 6000 Disk Module • Sun Blade 6000 Multi-Fabric Network Express Module • Sun Blade 6000 10GbE Multi-Fabric Network Express Module • Sun Blade 6000 Virtualized Multi-Fabric 10GbE Network Express Module 	Initial release.
5.02.14 (502E)	Common Array Manager 6.4.1	<ul style="list-style-type: none"> • Sun Blade 6000 Disk Module • Sun Blade 6000 Multi-Fabric Network Express Module • Sun Blade 6000 10GbE Multi-Fabric Network Express Module • Sun Blade 6000 Virtualized Multi-Fabric 10GbE Network Express Module 	Support was added for a second source to monitor disk module temperature and voltage.
5.04.03	Oracle Hardware Management Pack 2.0 (fwupdate utility)	<ul style="list-style-type: none"> • Sun Blade 6000 Disk Module • Sun Blade 6000 Multi-Fabric Network Express Module • Sun Blade 6000 10GbE Multi-Fabric Network Express Module • Sun Blade 6000 Virtualized Multi-Fabric 10GbE Network Express Module 	Support added for SAS-1/SAS-2 device coexistence in a Sun Blade 6000 Modular System chassis.

Note – At the release of this document, there is currently no support for using the Common Array Manager to upgrade disk module and SAS-NEM firmware to a level that allows SAS-1 device coexistence with server modules with SAS-2 REMs. For more information on upgrading disk module and SAS-NEM firmware to a level that allows SAS-1/SAS-2 coexistence in the chassis, see [Appendix C](#).

▼ To Upgrade Expander Firmware

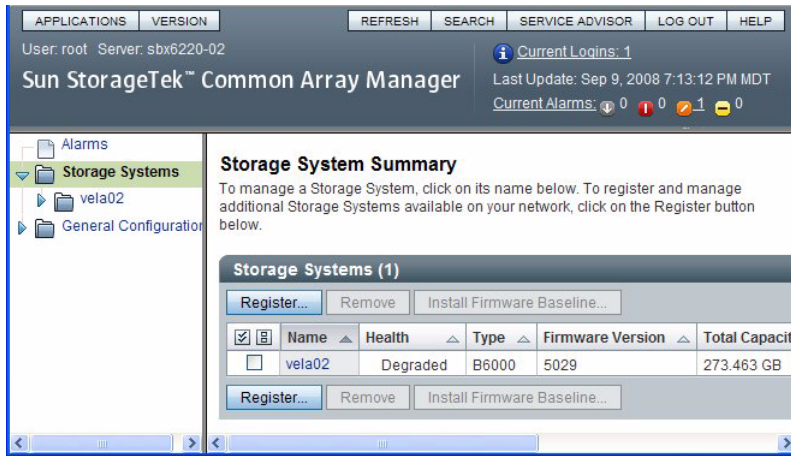
This procedure assumes that you have already registered your disk blades using the CAM interface. Open the CAM interface online help for instructions on registering disk blades.

Note – Registering the disk blades automatically makes CAM aware of the SAS-NEMs.

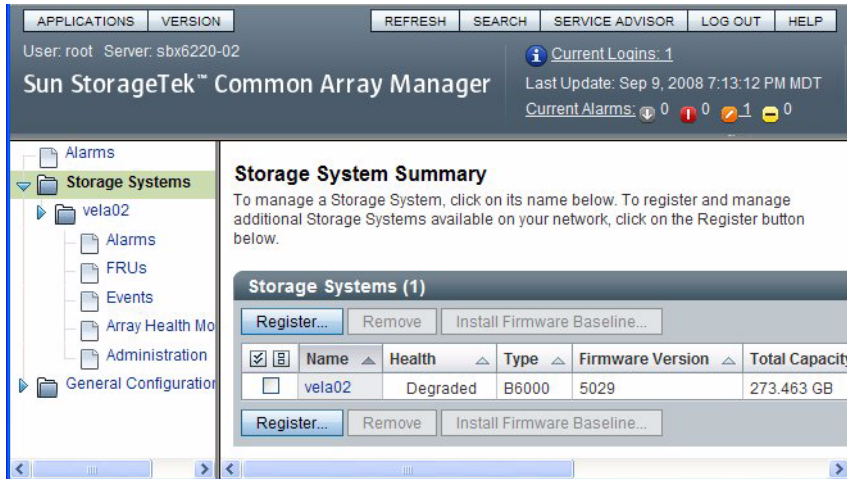
1. Open the CAM browser interface on the server that hosts the CAM software.

You see the registered storage systems listed in the opening summary page. In this example, there is only one registered storage system, a disk blade registered with the name “vela02.”

You can see that vela02’s health is degraded and also that there is a major (orange) alarm in the Current Alarms list.




2. Click the arrow to expand vela02 in the navigation tree.



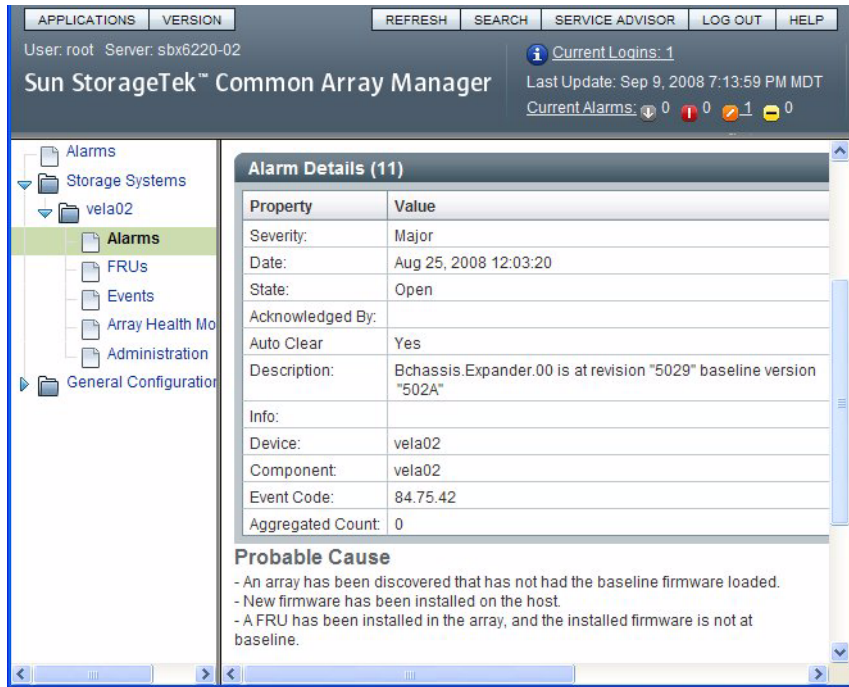
3. Select the Alarms item in the vela02 tree.

The Alarms page for vela02 opens. You can see one major severity (orange) alarm of type RevisionDeltaEvent.



4. Click the link under Alarm Details.

The Alarm Details page opens.



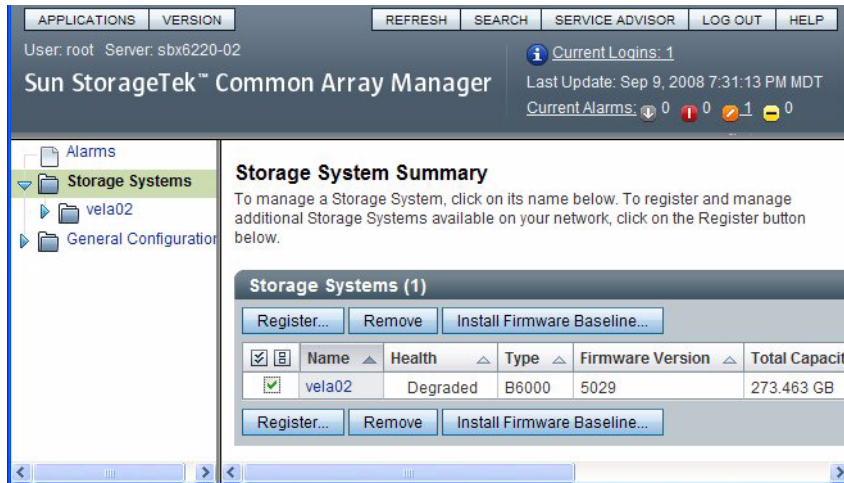
You can see in the Description field for this example that the installed version of expander firmware (5029) is not up to date. A more current version (502A) is available.

Note – You will get this same alarm if the firmware on *any* of the expanders (two on the disk blade and one on each of the SAS-NEMs) is not current for any reason.

The Probable Cause page lists various possible ways that the degraded condition that raised the alarm might have been produced.

5. Return to the Storage System Summary page by choosing Storage Systems in the navigation tree, and check the check box to the left of the vela02 name.

This selects the vela02 storage system and enables the firmware upgrade button.



Note – In this example, there is only one registered storage system. If there were more, the buttons would be enabled only for the checked system.

6. Click the **Install Firmware Baseline** button to upgrade the expander firmware. The Analyze and Install Array Firmware Baseline wizard opens.

Note – CAM upgrades the firmware on *every* expander that is not at the current level.



7. Click Next.

The next screen shows the current firmware versions and the baseline (correct) versions for each expander.

Sun StorageTek™ Common Array Manager

Analyze and Install Array Firmware Baseline

Steps Help Step 2: Analyze Arrays

1. Overview
2. Analyze Storage System Firmware
3. Review
4. Results

Firmware information for each array is displayed as the analysis is completed. The action column displays the recommended action to bring each array up to the current firmware baseline. For optimal performance, all arrays should have the same firmware version installed. This wizard will only install the components that are not at the baseline.

Storage Systems (1)

Name	Action	Current Firmware	Baseline
vela02	Install baseline, no disks Disk install is not required	Bchassis.Expander.00: 5029 Bchassis.Expander.01: 5029 nem.00: 5029 nem.01: 5029	Bchassis.Expander.00: 502A Bchassis.Expander.01: 502A nem.00: 502A nem.01: 502A

Previous Next Cancel



Caution – CAM will give you a message that your Sun Blade 6000 Disk Module firmware is not at baseline even if the firmware on your disk blade is at a later version than the version available in CAM. If your disk blade firmware is at a later version than CAM's baseline, **DO NOT** perform an upgrade. This will *downgrade* your disk blade instead, and could reduce functionality or disable the disk blade.

8. Accept the default Action and click Next.

The next screen shows a review of what you have selected.

Sun StorageTek™ Common Array Manager

Analyze and Install Array Firmware Baseline

Steps Help Step 3: Review

1. Overview
2. Analyze Storage System Firmware
3. Review
4. Results

Review the current selections and then click Finish to perform any specified firmware install.

⚠ All management operations for Storage System should stop prior to firmware install. Disk I/O must be quiesced if disk firmware is being installed.

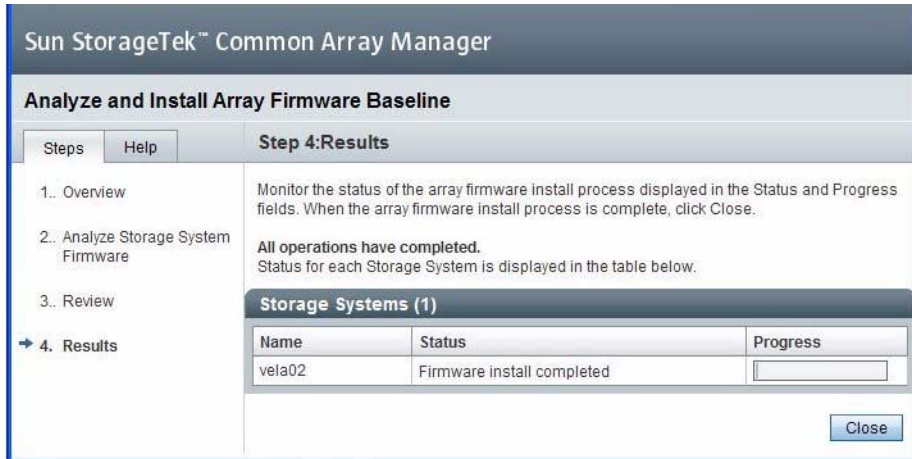
Storage Systems (1)

Name	Action	Current Firmware	Baseline
vela02	Install baseline, no disks Disk install is not required	Bchassis.Expander.00: 5029 Bchassis.Expander.01: 5029 nem.00: 5029 nem.01: 5029	Bchassis.Expander.00: 502A Bchassis.Expander.01: 502A nem.00: 502A nem.01: 502A

Previous Finish Cancel

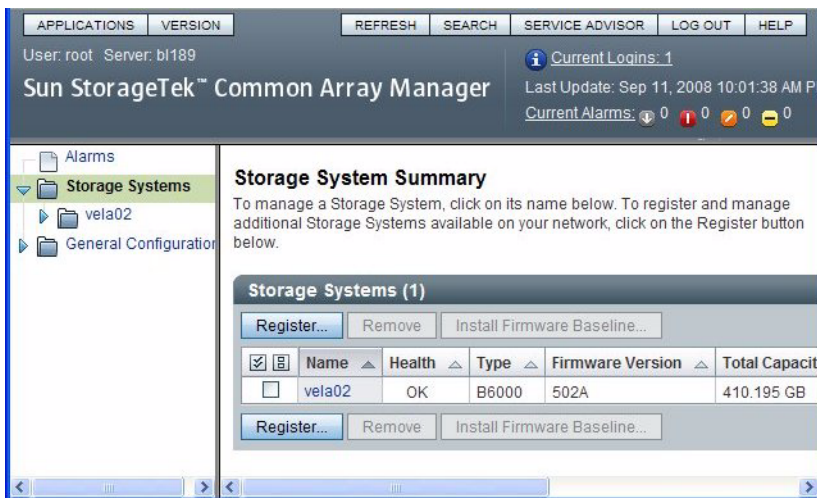
9. If you need to change the Action, click Previous. Otherwise, click Finish.

A series of screens appear, with the Status field showing the current expander being upgraded. When the process is completed, the Firmware Install completed message appears in the Status Field.



10. Click Close.

The Storage System Summary page reappears. You can see that the disk blade's Health has changed to OK and the alarm is turned off.



Using the SIA Application to Upgrade x86 Server Blade and Host Bus Adapter Firmware

Oracle's Sun Installation Assistant (SIA) is a tool for Oracle x86 servers and can perform a variety of tasks. You can use SIA to upgrade System BIOS and ILOM service processor firmware, as well as HBA firmware.

Using SIA to Upgrade x86 Systems and HBAs



Caution – For all systems: Never power off the system during an upgrade.

Local and Remote Media Options

On your Sun Blade server module, SIA can perform tasks using one of the following methods:

- **Locally, at the server using the SIA CD/DVD.** Server modules do not have internal CD/DVD drives. You must boot SIA from a USB CD/DVD drive or a thumb drive that is directly connected to the server module's USB dongle port. This method assumes you have set up a VGA console with keyboard and mouse as described in your Sun Blade server module installation documentation.

- **Remotely, using the remote console either through the chassis CMM or through the server module service processor.** This method allows you to boot SIA from a virtual CD-ROM. Refer to your server's service processor Lights Out Manager (Integrated LOM or Embedded LOM) documentation for information on the LOM remote console. (Note that there are multiple versions of LOM, so be sure to refer to the guide that matches your server's installed version.)

The SIA program provides a graphical user interface that guides you through the process of the updating firmware.

Note – The procedure for updating your SAS host bus adapter firmware with SIA is equally applicable for x86 server blades with on-board LSI chips, x86 server blades with RAID 0/1 Expansion Module cards (LSI), and x86 server blades with RAID 5 Expansion Module cards (Adaptec).

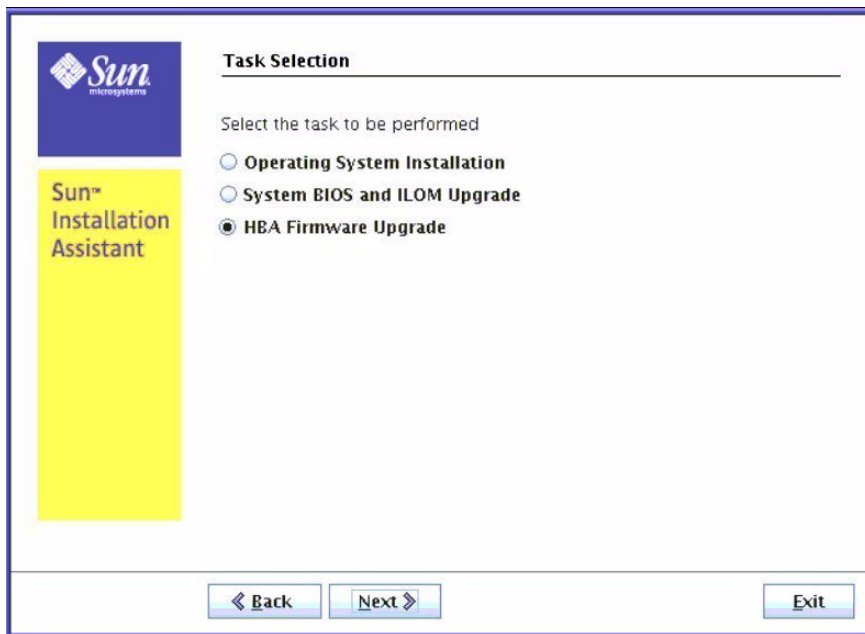
▼ To Perform a Firmware upgrade Task With SIA Using Local or Remote Media

1. Start SIA using one of these methods:

- Insert the SIA CD into the server dongle's CD/DVD drive and power on or reboot the server.
- Use a browser to log in to the server's ILOM service processor via your KVMs and use the **Remote Control --> Launch Redirection** feature to redirect the server to a network-accessible virtual CD-ROM that contains the SIA CD image. Reboot the server. Be sure to direct the server to boot from your virtual CD-ROM (typically done using the F8 menu during server boot).

2. Once started, follow the screens until you reach the Task Selection page.

An example of the Task Selection page is shown below. The tasks listed are server-dependant and only those supported by your server and installed hardware will be available.



3. Select the task to perform, either System BIOS and ILOM, or HBA firmware upgrade and click Next.

Follow the instructions in the successive screens.

Note – You might be able to perform both upgrade tasks during the same SIA session. If you finish the SP firmware and system BIOS upgrade without being asked to reboot, click the Back button repeatedly until you return to the Task Selection screen. If you finish the first task and are asked to reboot, make sure that you leave the SIA media in place and reboot to SIA.

Using the `lsiutil` Software

Some procedures in this manual require the use of a software program called `lsiutil`. This appendix provides information on where to obtain the software and how to use it.



Caution – The `lsiutil` software has many features and capabilities. Executing certain combinations of commands can leave your system in an unrecoverable state. When you use this software, be sure you execute the procedures exactly as they are documented here, command by command. Do not skip commands or steps and do not add commands or steps that are not in the documented procedure.

The following topics are covered in this chapter:

- “Where to Obtain the `lsiutil` Software” on page 121
- “Installing `lsiutil`” on page 122
- “To Use `lsiutil` From the Command Line” on page 122
- “To Use the Interactive `lsiutil` Menus” on page 122
- “Why Save LSI Host Bus Adapter Persistent Mappings?” on page 126
- “When to Save LSI Host Bus Adapter Persistent Mappings” on page 126
- “To Save a Snapshot of Your Host Bus Adapter Persistent Mappings” on page 127
- “To Restore a Snapshot of Your Host Bus Adapter Persistent Mappings” on page 131

Where to Obtain the `lsiutil` Software

To use `lsiutil` for the procedures in this appendix, you must have version 1.60 at minimum.

To obtain the latest `lsiutil` software, go to the disk blade site:

(<http://www.oracle.com/goto/blades>)

Installing `lsiutil`

Download `lsiutil` and unzip it in a temporary directory. It will create operating system-specific subdirectories. If you are running `lsiutil` on a SPARC blade, use the `Solaris` subdirectory. Otherwise, use the `Solaris x86` subdirectory.

▼ To Use `lsiutil` From the Command Line

Several of the procedures in this appendix use the `lsiutil` software at the command line.



Caution – To avoid harm to your system, you must use the commands exactly as they are presented in this document.

To run an `lsiutil` command from the command line, log in as `root` and then follow these steps:

1. **Change to the directory where you extracted the `lsiutil` zip file. For example,**

```
# cd directoryname
```
2. **Change to the subdirectory for your operating system. For example:**

```
# cd Solaris x86
```
3. **Change the permissions (read and execute access) on the `lsiutil` file by entering the following command:**

```
# chmod 755 lsiutil
```
4. **Enter the command with parameters. For example,**

```
# ./lsiutil -p1 -a 0 8
```

▼ To Use the Interactive `lsiutil` Menus

To open the interactive `lsiutil` menus, log in as `root` and then follow these steps:

1. Change to the directory that contains the appropriate version of `lsiutil` for your operating system. For example,

```
# cd directoryname/Solaris x86
```


2. If you have not done so already, change the permissions (read and execute access) on the `lsiutil` file by entering the following command:

```
# chmod 755 lsiutil
```

3. Start `lsiutil` by entering the command:

```
# ./lsiutil
```

The opening screen appears.

A screenshot of a terminal window titled "root@bl189:~/LSIUtilKit1.60/Linux". The terminal shows the execution of the `lsiutil` command. The output displays the version and date of the utility, the number of MPT ports found, and a table of port details. The table has columns for Port Name, Chip Vendor/Type/Rev, MPT Rev, Firmware Rev, and IOC. One port is listed with the name `/proc/mpt/ioc0` and vendor `LSI Logic SAS1068E B2`. The prompt "Select a device: [1-1 or 0 to quit]" is visible at the bottom of the terminal output.

```
root@bl189:~/LSIUtilKit1.60/Linux
File Edit View Terminal Tabs Help
[root@bl189 ~]# cd LSIUtilKit1.60/Linux
[root@bl189 Linux]# ./lsiutil

LSI Logic MPT Configuration Utility, Version 1.60, July 11, 2008

1 MPT Port found

   Port Name          Chip Vendor/Type/Rev   MPT Rev  Firmware Rev  IOC
1. /proc/mpt/ioc0    LSI Logic SAS1068E B2   105      01185b00      0

Select a device: [1-1 or 0 to quit]
```

4. Type **1** and press **Enter**.

The main menu appears. Because there are so many menu items (100), not all the items are shown.

```
root@bl189:~/LSIUtilKit1.60/Linux
File Edit View Terminal Tabs Help
Select a device: [1-1 or 0 to quit] 1

 1. Identify firmware, BIOS, and/or FCode
 2. Download firmware (update the FLASH)
 4. Download/erase BIOS and/or FCode (update the FLASH)
 8. Scan for devices
10. Change IOC settings (interrupt coalescing)
13. Change SAS IO Unit settings
16. Display attached devices
20. Diagnostics
21. RAID actions
22. Reset bus
23. Reset target
42. Display operating system names for devices
45. Concatenate SAS firmware and NVDATA files
59. Dump PCI config space
60. Show non-default settings
61. Restore default settings
66. Show SAS discovery errors
69. Show board manufacturing information
97. Reset SAS link, HARD RESET
98. Reset SAS link
99. Reset port
   e Enable expert mode in menus
   p Enable paged mode
   w Enable logging

Main menu, select an option: [1-99 or e/p/w or 0 to quit]
```

Tip – You do not have to see a command to enter it. If you know its number, just enter it. Alternatively, you can choose to see the entire menu if you want to, as shown in Step 4.

5. Enter **p to enable paged mode, or **e** to enable expert mode.**

The complete menu appears in expert mode.


```
root@bl189:~/LSIUtilKit1.60/Linux
File Edit View Terminal Tabs Help
42. Display operating system names for devices
45. Concatenate SAS firmware and NVDATA files
59. Dump PCI config space
60. Show non-default settings
61. Restore default settings
66. Show SAS discovery errors
69. Show board manufacturing information
97. Reset SAS link, HARD RESET
98. Reset SAS link
99. Reset port
   e Enable expert mode in menus
   p Enable paged mode
   w Enable logging

Main menu, select an option: [1-99 or e/p/w or 0 to quit] p

Enabled paged mode, 24 lines

Main menu, select an option: [1-99 or e/p/w or 0 to quit] e

Enabled expert mode in menus

  1. Identify firmware, BIOS, and/or FCode
  2. Download firmware (update the FLASH)
  3. Upload firmware
  4. Download/erase BIOS and/or FCode (update the FLASH)
  5. Upload BIOS and/or FCode
  6. Download SEEPROM
  7. Upload SEEPROM
  8. Scan for devices
  9. Read/change configuration pages
 10. Change IOC settings (interrupt coalescing)
 13. Change SAS IO Unit settings
 14. Change IO Unit settings (multi-pathing, queuing, caching)
 15. Change persistent mappings
 16. Display attached devices
 17. Show expander routing tables
 18. Change SAS WWID
 19. Test configuration page actions
 20. Diagnostics
 21. RAID actions
--more, hit RETURN--
```

Each time you press Enter you will see a new page of menu items until you reach the end.

Why Save LSI Host Bus Adapter Persistent Mappings?

Replacing a failed LSI host bus adapter on a server blade running the Solaris OS can lead to significant downtime if your replacement adapter does not know how the old adapter was addressing the disks in its server blade and in the disk blade. Among other problems, your system will not know where your boot drive is located.

The `lsiutil` software is used to export a snapshot of the addressing configuration (persistent mappings) of a healthy LSI host bus adapter to a file. Then, if the adapter fails at a later date, the persistent mappings snapshot can be reloaded on the replacement adapter and your system will operate as before.



Caution – You need to keep a persistent mappings snapshot for the LSI host bus adapters on all server blades in your chassis that are running the Solaris OS. A copy of each file must be kept on *external* media.

When to Save LSI Host Bus Adapter Persistent Mappings

You need to keep a snapshot of the persistent mappings of your LSI host bus adapters under any of these conditions:

- A new server blade, running the Solaris OS and paired with a disk blade, has been added to the chassis.

Make the snapshot of the LSI host bus adapter on this server blade after you have completed its installation and configuration. This means after you have finished creation of RAID volumes and chosen a boot volume or disk.

- You have replaced a disk blade in the chassis that is paired with a server blade running the Solaris OS.

Make the snapshot of the LSI host bus adapter on this server blade after you have replaced the disk blade.

- You have replaced a SAS-NEM in your chassis and the chassis contains server blades running both CAM (or a CAM agent) and the Solaris OS.

Make the snapshot of the LSI host bus adapter on these server blades after you have replaced the SAS-NEM

Saving and Restoring a Persistence Map

You can use `lsiutil` to both save and restore a snapshot of your persistence mapping.

▼ To Save a Snapshot of Your Host Bus Adapter Persistent Mappings

1. Run `lsiutil` in interactive mode (see [“To Use the Interactive `lsiutil` Menus”](#) on page 122).

```
root@bl189:~/LSIUtilKit1.60/Linux
File Edit View Terminal Tabs Help
[root@bl189 ~]# cd LSIUtilKit1.60/Linux
[root@bl189 Linux]# ./lsiutil

LSI Logic MPT Configuration Utility, Version 1.60, July 11, 2008


1 MPT Port found

   Port Name      Chip Vendor/Type/Rev  MPT Rev  Firmware Rev  IOC
1.  /proc/mpt/ioc0  LSI Logic SAS1068E B2   105      01185b00     0

Select a device: [1-1 or 0 to quit]
```

2. Type **1** and press **Enter**.

The basic (incomplete) interactive menu opens.



```
root@b1189:~/LSIUtilKit1.60/Linux
File Edit View Terminal Tabs Help
Select a device: [1-1 or 0 to quit] 1

 1. Identify firmware, BIOS, and/or FCode
 2. Download firmware (update the FLASH)
 4. Download/erase BIOS and/or FCode (update the FLASH)
 8. Scan for devices
10. Change IOC settings (interrupt coalescing)
13. Change SAS IO Unit settings
16. Display attached devices
20. Diagnostics
21. RAID actions
22. Reset bus
23. Reset target
42. Display operating system names for devices
45. Concatenate SAS firmware and NVDATA files
59. Dump PCI config space
60. Show non-default settings
61. Restore default settings
66. Show SAS discovery errors
69. Show board manufacturing information
97. Reset SAS link, HARD RESET
98. Reset SAS link
99. Reset port
   e Enable expert mode in menus
   p Enable paged mode
   w Enable logging

Main menu, select an option: [1-99 or e/p/w or 0 to quit]
```

Tip – You cannot see all the commands in the menu, but you can enter any command number if you know what it is. In this case you want to open the Persistence menu. This is done with command 15.

3. Type 15 and press Enter.

The Persistence menu opens.


```
root@bl189:~/LSIUtilKit1.60/Linux
File Edit View Terminal Tabs Help
Main menu, select an option: [1-99 or e/p/w or 0 to quit] 15
1. Show persistent mappings
6. Change Bus/Target or EnclosureId of an existing mapping
7. Save persistent mappings to a file
8. Load persistent mappings from a file
10. Clear all persistent mappings
11. Clear all non-present persistent mappings
12. Change (enable/disable) persistence
99. Reset port
e Enable expert mode in menus
p Disable paged mode
w Enable logging
Persistence menu, select an option: [1-99 or e/p/w or 0 to quit]
```

4. Type 7 and press Enter.

You are prompted for the name of the file where you want to store the host bus adapter persistent mappings snapshot.


Note – If you expect to boot from DOS when you restore the persistent mappings snapshot, you must restrict the file name to 8 characters.

5. Enter a file name. For example, `HBA33_persistence_file`.



```
root@bl189:~/LSIUtilKit1.60/Linux
File Edit View Terminal Tabs Help
Main menu, select an option: [1-99 or e/p/w or 0 to quit] 15
 1. Show persistent mappings
 6. Change Bus/Target or EnclosureId of an existing mapping
 7. Save persistent mappings to a file
 8. Load persistent mappings from a file
10. Clear all persistent mappings
11. Clear all non-present persistent mappings
12. Change (enable/disable) persistence
99. Reset port
  e Enable expert mode in menus
  p Disable paged mode
  w Enable logging
Persistence menu, select an option: [1-99 or e/p/w or 0 to quit] 7
Enter persistent mapping filename: HBA33_persistence_file
```

The file is saved to the current directory. `lsiutil` confirms this by displaying a statement of the number of persistent entries saved.



```
root@bl189:~/LSIUtilKit1.60/Linux
File Edit View Terminal Tabs Help
Main menu, select an option: [1-99 or e/p/w or 0 to quit] 15
 1. Show persistent mappings
 6. Change Bus/Target or EnclosureId of an existing mapping
 7. Save persistent mappings to a file
 8. Load persistent mappings from a file
10. Clear all persistent mappings
11. Clear all non-present persistent mappings
12. Change (enable/disable) persistence
99. Reset port
  e Enable expert mode in menus
  p Disable paged mode
  w Enable logging
Persistence menu, select an option: [1-99 or e/p/w or 0 to quit] 7
Enter persistent mapping filename: HBA33_persistence_file
4 persistent entries saved
Persistence menu, select an option: [1-99 or e/p/w or 0 to quit]
```

6. Type 0 (zero) and press Enter three times to exit `lsiutil`.

```
root@bl189:~/LSIUtilKit1.60/Linux
File Edit View Terminal Tabs Help
Persistence menu, select an option: [1-99 or e/p/w or 0 to quit] 7
Enter persistent mapping filename: HBA33_persistence_file
4 persistent entries saved
Persistence menu, select an option: [1-99 or e/p/w or 0 to quit] 0
Main menu, select an option: [1-99 or e/p/w or 0 to quit] 0
      Port Name          Chip Vendor/Type/Rev  MPT Rev  Firmware Rev  IOC
1. /proc/mpt/ioc0      LSI Logic SAS1068E B2    105      01185b00      0
Select a device: [1-1 or 0 to quit] 0
[root@bl189 Linux]#
```



Caution – You *must* save the persistence table snapshot to external media as you will not be able to reinstall it from a local disk if your LSI host bus adapter fails.

You must have a snapshot for *every* server blade in the chassis. Label them carefully.

▼ To Restore a Snapshot of Your Host Bus Adapter Persistent Mappings

The stored snapshot files are used when a host bus adapter fails and must be replaced. They are not needed for any other purpose.

Note – Unless your OS is installed on a single disk on your server blade, or a RAID volume with at least one member disk installed on your server blade (always true for SPARC systems, for example), you need to reboot with an external OS (a bootable DOS disk is preferred, if possible) to perform this procedure.

To restore the persistent mappings from a previously saved snapshot file (for example, `HBA33_persistence_file`), do the following:

1. **Make sure that you saved the persistence mappings snapshot file in the same directory as the `lsiutil` software you are using. Copy the file from external media if necessary.**
2. **Run `lsiutil` in interactive mode (see “[To Use the Interactive `lsiutil` Menus](#)” on page 122).**

```
root@bl189:~/LSIUtilKit1.60/Linux
File Edit View Terminal Tabs Help
[root@bl189 ~]# cd LSIUtilKit1.60/Linux
[root@bl189 Linux]# ./lsitutil

LSI Logic MPT Configuration Utility, Version 1.60, July 11, 2008

1 MPT Port found

   Port Name          Chip Vendor/Type/Rev   MPT Rev  Firmware Rev  IOC
1. /proc/mpt/ioc0    LSI Logic SAS1068E B2   105      01185b00     0

Select a device: [1-1 or 0 to quit]
```

3. Type 1 and press Enter.

The basic (incomplete) interactive menu opens.

Tip – You cannot see all the commands in the menu, but you can enter any command number if you know what it is. In our case we want to open the Persistence menu. This is done with command 15.

4. Type 15 and press Enter.

The Persistence menu opens.

```
root@bl189:~/LSIUtilKit1.60/Linux
File Edit View Terminal Tabs Help


Main menu, select an option: [1-99 or e/p/w or 0 to quit] 15

 1. Show persistent mappings
 6. Change Bus/Target or EnclosureId of an existing mapping
 7. Save persistent mappings to a file
 8. Load persistent mappings from a file
10. Clear all persistent mappings
11. Clear all non-present persistent mappings
12. Change (enable/disable) persistence
99. Reset port
  e Enable expert mode in menus
  p Disable paged mode
  w Enable logging

Persistence menu, select an option: [1-99 or e/p/w or 0 to quit]
```


5. Type 10 and press Enter.

This clears any persistence map that is present in the host bus adapter.




```
root@bl189:~/LSIUtilKit1.60/Linux
File Edit View Terminal Tabs Help
1. Show persistent mappings
6. Change Bus/Target or EnclosureId of an existing mapping
7. Save persistent mappings to a file
8. Load persistent mappings from a file
10. Clear all persistent mappings
11. Clear all non-present persistent mappings
12. Change (enable/disable) persistence
99. Reset port
e Enable expert mode in menus
p Enable paged mode
w Enable logging

Persistence menu, select an option: [1-99 or e/p/w or 0 to quit] 10
Clearing all persistent entries...
Persistence menu, select an option: [1-99 or e/p/w or 0 to quit] 8
```

6. Type 8 and press Enter.

You are prompted for the name of a persistent mappings snapshot file to load.

Type the name of the snapshot file. For example, **HBA33_persistence_file**.



```
root@bl189:~/LSIUtilKit1.60/Linux
File Edit View Terminal Tabs Help
1. Show persistent mappings
6. Change Bus/Target or EnclosureId of an existing mapping
7. Save persistent mappings to a file
8. Load persistent mappings from a file
10. Clear all persistent mappings
11. Clear all non-present persistent mappings
12. Change (enable/disable) persistence
99. Reset port
e Enable expert mode in menus
p Enable paged mode
w Enable logging

Persistence menu, select an option: [1-99 or e/p/w or 0 to quit] 10
Clearing all persistent entries...
Persistence menu, select an option: [1-99 or e/p/w or 0 to quit] 8
Enter persistent mapping filename: HBA33_persistence_file
```

7. Press Enter.

The persistence map is loaded. `lsiutil` confirms that four persistence entries were loaded.

```
root@bl189:~/LSIUtilKit1.60/Linux
File Edit View Terminal Tabs Help
1. Show persistent mappings
6. Change Bus/Target or Enclosureid of an existing mapping
7. Save persistent mappings to a file
8. Load persistent mappings from a file
10. Clear all persistent mappings
11. Clear all non-present persistent mappings
12. Change (enable/disable) persistence
99. Reset port
e Enable expert mode in menus
p Enable paged mode
w Enable logging

Persistence menu, select an option: [1-99 or e/p/w or 0 to quit] 10

Clearing all persistent entries...

Persistence menu, select an option: [1-99 or e/p/w or 0 to quit] 8

Enter persistent mapping filename: HBA33_persistence_file
4 persistent entries loaded

Persistence menu, select an option: [1-99 or e/p/w or 0 to quit]
```

8. Type 0 (zero) and press Enter three times to exit lsiutil.

Note – The adapter persistent mappings snapshot file you just loaded is *still valid*. Make sure that you save a copy of it on external media.

Upgrading the Disk Module to Allow a SAS-2 Compatible Server Blade in the Chassis

This appendix provides information on what is required to allow a server module with a SAS-2 RAID Expansion Module (REM) to function in a Sun Blade 6000 Modular System chassis that has SAS-1 NEMs, or SAS-1 NEMs and a Sun Blade 6000 Disk Module (which is a SAS-1 device).

Note – The easiest way to identify SAS-1 and SAS-2 devices is by the maximum data transfer speeds supported by the device. SAS-1 devices have a maximum data transfer speed of 3 Gb per second. SAS-2 devices have a maximum data transfer speed of 6 Gb per second. Refer to your device documentation for a list of specifications.

Server with a SAS-2 REM Might Hang When Inserted into a Chassis with SAS-1 Devices

If a server module with a SAS-2 REM is inserted into a Sun Blade 6000 Modular System chassis that has SAS-1 NEMs, or SAS-1 NEMs and Sun Blade 6000 Disk Module, the server's SAS-2 REM might hang. To prevent this, you need to upgrade the SAS expander firmware of your SAS-1 NEMs and Sun Blade 6000 Disk Modules to a version that supports SAS-1/SAS-2 coexistence. The minimum SAS expander

firmware required for SAS-1 NEMs and the Sun Blade 6000 Disk Module is 5.04.03. **This upgrade must be done before you insert a SAS-2 server blade into the chassis.**

Available SAS-1 NEMs and disk modules that would require this upgrade include Oracle's:

- Sun Blade 6000 Disk Module (B18-AA)
- Sun Blade 6000 Multi-Fabric Network Express Module (X4212A)
- Sun Blade 6000 10GbE Multi-Fabric Network Express Module (X4236A)
- Sun Blade 6000 Virtualized Multi-Fabric 10GbE Network Express Module (X4238)

An example of a SAS-2 server module is the Sun Blade X6270 M2. This server includes a SAS-2 REM. Available SAS-2 REMs include:

- Sun Storage 6 Gb SAS REM RAID HBA (SGX-SAS6-R-REM-Z)
- Sun Storage 6 Gb SAS REM HBA (SGX-SAS6-REM-Z)

Refer to your product documentation to find out if your server blade includes a SAS-2 REM.

For detailed instructions on performing the upgrade, refer to the *SAS-1/SAS-2 Compatibility Upgrade Guide for the Sun Blade 6000 Modular System* (821-1800) which can be found in:

(<http://docs.sun.com/coll/blade6000dskmod>)

Index

Numerics

1068E

verifying correct version on X6220 blades, 26

A

Adaptec host bus adapters

about, 73

Adaptec RAID

BIOS configuration for x86 servers with any OS, 91

configuring after the OS is installed, 91

creating RAID volume, 92

setting up with BIOS utility, 91

Adaptec volumes

single-drive volume, 94

admin user, 42

administrator role, 42

ALOM CLI mode, 42

ALOM CMT shell

admin user, 42

C

CAM

about enclosure management, 105

agent, 106

monitoring component health, 108

obtaining CAM software, 107

overview, 106

problem resolution, 107

updating enclosure firmware, 109

using with disk blades and SAS-NEMs, 108

CLI mode

ALOM, 42

compatibility for SAS-1/SAS-2 devices, 135

D

disk blades

controlling, 9

front panel, 5

inserting, 9

replacing, 10

disk drive rules, 48

disk drives

mixing, 48

supported, 17

updating firmware, 24

disk modules and server modules in a chassis, 7

driver support for x86 server blades, 34

drivers

Linux, 36

required, 22

where to get, 24

Windows, 37

E

expander firmware

updating, 109

F

firmware

minimum required versions, 41

required, 22

updating disk drive firmware, 24

updating expander firmware, 109

- upgrading for SAS-1/SAS-2 compatibility, 135
- upgrading LSI host bus adapter firmware with SIA, 118
- where to get, 24

front panel LEDs, 5

H

HBA driver support for a disk blade, 34

I

ILOM for disk modules and SAS-NEM modules, 97

ILOM proxy program

- for disk modules and SAS-NEM modules, 98
- navigation, 99
- starting, 99
- what you can see, 101
- what you can set or change, 101

installation, overview, 9

L

LEDs

- disk drive status, 6
- front panel, 5
- locate, 5
- module activity, 6
- service action required, 5

Linux drivers, 36

locate button, 5

locate LED, 5

LSI host bus adapters

- restoring a snapshot of your persistent mappings, 131
- saving a snapshot of your persistent mappings, 127
- why save a snapshot of your persistent mappings, 126

LSI RAID

- configuring for Solaris OS after OS is installed, 84
- configuring for SPARC systems before installing OS, 85
- configuring for Windows and Linux after OS is installed, 84
- configuring for x86 servers with any OS before OS is installed, 85
- setting up with BIOS utility, 83

lsiutil

installing, 122

using from the command-line, 122

using the interactive menus, 122

where to obtain the software, 121

M

mixing disk drives, 48

multipathing, 48

configuring after installation, 59

failover, 65

handling multiple paths at the OS level, 57

hardware RAID hides dual paths from the OS, 58

installing the RHEL 5.x OS on multi-path disks, 59

overview, 53

using Adaptec host bus adapters, 57

using LSI host bus adapters, 54

with Linux, 58

with the Solaris OS, 59

with Windows 2008, 62

N

Network Express Modules, mixing rules, 21

Network Express Modules, supported, 21

O

operating systems

special considerations, 16

overview

CAM, 106

installation, 9

Sun blade 6000 Disk Module, 2

P

persistent mappings

restoring a snapshot for your LSI host bus adapters, 131

saving a snapshot for your LSI host bus adapters, 127

why save a snapshot for your LSI host bus adapters, 126

persistent mappings snapshot

when to save, 126

product updates, where to get, xi

R

- RAID configuration before or after OS installation, 84
- RAID controllers, 18, 19
- RAID options, 81
- related documentation, xi
- replacing disk blades, 10
- replacing disk drives, 49
- replacing older X6220 server blades, 26
- requirements for all server blades in chassis, 15

S

- SAS expanders
 - updating firmware, 109
- SAS host bus adapter, 6
- SAS-1/SAS-2 compatibility issue, 135
- server modules and disk modules in a chassis, 7
- servers
 - special considerations, 16
- SIA
 - downloading for each x86 server blade, 28
 - local and remote media options, 117
 - performing a firmware update task, 118
 - using to upgrade x86 HBA firmware, 117
- Solaris network install servers
 - updating for SPARC servers, 45
 - updating for x86 servers, 45
- special considerations
 - operating systems, 16
 - servers, 16
- Sun Blade 6000 Disk Module, overview, 2
- supported disk drives, 17

T

- target IDs
 - allowed range with LSI controllers, 70
 - assigning with Adaptec host bus adapters, 74
 - collision with rebuilt volumes, 72
 - how assigned with LSI controllers, 69
 - mapping with LSI controllers, 70
- terminology, 1
- typographic conventions, xii

U

- updating enclosure firmware, 109

- updating expander firmware, 109
- updating Solaris network install servers
 - for SPARC servers, 45
 - for x86 servers, 45
- upgrading
 - blade server Windows drivers, 37
 - chassis ILOM, 27
 - drivers on x86 server blades, 34
 - enclosure device driver in Windows, 40
 - firmware on SPARC server blades, 41
 - firmware on x86 server blades, 27
 - host bus adapter firmware on x86 server blades, 34
 - LSI expander firmware, 44
 - LSI host bus adapter firmware with SIA, 118
 - older servers, 26
 - Solaris network install servers, 45
 - SP firmware and system BIOS on x86 server blades, 31
 - SPARC blade disk controller firmware, 43
 - SPARC blade system firmware, 41

V

- verifying
 - Adaptec adapter firmware version by rebooting, 33
 - Adaptec Adapter firmware version without rebooting, 33
 - correct 1068E version, 26
 - HBA firmware versions on x86 server blades, 31
 - LSI adapter firmware version, 31
 - SP firmware and system BIOS versions on x86 server blades, 28
- verifying correct 1068E version, 26
- volume IDs
 - assigning with Adaptec host bus adapters, 77
 - collision with rebuilt volumes, 72
 - how assigned with LSI controllers, 71

W

- where to get drivers and firmware, 24
- Windows drivers, 37

X

- X6220 blades
 - verifying correct 1068E version, 26

Z

zoning with multiple hosts not supported, 106