



SGI® Modular InfiniteStorage™ (MIS)
1.5 Platform User Guide

007-5916-001

COPYRIGHT

© 2013 Silicon Graphics International Corp. All rights reserved; provided portions may be copyright in third parties, as indicated elsewhere herein. No permission is granted to copy, distribute, or create derivative works from the contents of this electronic documentation in any manner, in whole or in part, without the prior written permission of SGI.

LIMITED RIGHTS LEGEND

The software described in this document is “commercial computer software” provided with restricted rights (except as to included open/free source) as specified in the FAR 52.227-19 and/or the DFAR 227.7202, or successive sections. Use beyond license provisions is a violation of worldwide intellectual property laws, treaties and conventions. This document is provided with limited rights as defined in 52.227-14.

The electronic (software) version of this document was developed at private expense; if acquired under an agreement with the USA government or any contractor thereto, it is acquired as “commercial computer software” subject to the provisions of its applicable license agreement, as specified in (a) 48 CFR 12.212 of the FAR; or, if acquired for Department of Defense units, (b) 48 CFR 227-7202 of the DoD FAR Supplement; or sections succeeding thereto. Contractor/manufacturer is SGI, 46600 Landing Parkway, Fremont, CA 94538.

TRADEMARKS AND ATTRIBUTIONS

Silicon Graphics, SGI, the SGI logo, InfiniteStorage, and Supportfolio are trademarks or registered trademarks of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries worldwide.

Fusion-MPT, Integrated RAID, MegaRAID, and LSI Logic are trademarks or registered trademarks of LSI Logic Corporation. InfiniBand is a registered trademark of the InfiniBand Trade Association. Intel and Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries. Internet Explorer and Windows are registered trademarks of Microsoft Corporation. Java and Java Virtual Machine are trademarks or registered trademarks of Sun Microsystems, Inc. Linux is a registered trademark of Linus Torvalds, used with permission by SGI. Novell and Novell Netware are registered trademarks of Novell Inc. PCIe and PCI-X are registered trademarks of PCI SIG. Red Hat and all Red Hat-based trademarks are trademarks or registered trademarks of Red Hat, Inc. in the United States and other countries. Sharp is a registered trademark of Sharp corporation. SUSE LINUX and the SUSE logo are registered trademarks of Novell, Inc. VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions.

All other trademarks mentioned herein are the property of their respective owners.

Record of Revision

Version	Description
001	July 2013 Original printing for the new hardware architecture of the MIS 1.5 platform. Owners of the MIS 1.0 platforms should refer to <i>SGI Modular Infinite Storage (MIS) Platform, User Guide, version 1.0</i> , document number 007-5818-004.

:

Contents

0	Introduction	. xix
	What's New	. xix
	Audience	. xix
	Important Information	. xix
	Environmental Requirements	. xx
	Safety Precautions	. xxii
	ESD Precautions	. xxii
	Safety & Emissions	. xxii
	Electromagnetic Compatibility	. xxiii
	Safety Certification	. xxiii
	Chapter Descriptions	. xxiv
	Related Publications	. xxv
	Conventions	. xxvi
	Product Support	. xxvi
	Software Downloads	. xxvi
	CRU/FRU	. xxvii
	Purchasable Support & Maintenance Programs	. xxvii
	Reader Comments	. xxviii
1	System Overview	1
	New Hardware Features of MIS 1.5	5
	Bridge Midplane	5
	SAS Interposers and SAS Midplane	5
	MIS Enclosure	6
	Front Grille and Control Panels	7
	Rear Panel Components	8
	Cable Management Arm	11
	MIS Common Modules	12

Power Supply Modules 12
Fan Assembly Modules 14
StorBrick Modules 14
MIS Server Platform or JBOD Unit 17
Server Module 21
Layout of Server CPUs, and PCIe Risers HBAs 23
Turbo Boost 24
Boot Drive Module 25
MIS JBOD I/O Module 26
System Diagrams 28
2 System Interfaces 33
Control Panel 33
MIS Server Control Panel 33
Powering an MIS Platform 36
Turning the machine on. 36
Turning the machine off 36
MIS JBOD Control Panel 37
Disk Drive Light Pipes 38
Power Supply LEDs 38
Power Supply Monitoring 39
Requirements 40
Viewing Existing Filters. 40
Setting Filters 40
Clearing Filters 41
Fan Base Interface 41
MIS-S9D Proprietary Network Interface 42
BMC Integrated Web Console 43
System Debug Log 44
Server Health 45
Sensor Readings 45
Event Log 47
Power Statistics 49
Configuration Tab 49

Alerts	50
Alert Email.	51
Server Power Control	54
Intelligent Management Platform Interface (IPMI 2.0)	56
3 System Software	59
Overview	59
Section Guide	60
Downloading & Installing Software.	62
Required Downloads	63
Installing MegaRAID Storage Manager for CentOS	63
Installing Zones 1.4.2 for CentOS	65
Starting Zones 1.4.2 for CentOS	65
Installing MegaRAID Storage Manager for Linux.	65
Installing Zones 1.4.2 for Linux	67
Starting Zones 1.4.2 for Linux.	67
Installing MegaRAID Storage Manager for Windows	67
Installing Python for Windows	68
Installing Zones 1.4.2 for Windows	68
Starting Zones 1.4.2 for Windows.	69
Operating Zones 1.4.2 for CentOS, Linux, & Windows	69
Zoning MIS Server Platforms	69
Zones interface.	70
Opening a Session	71
144 Drives Setup	73
Saving Zoning Configuration	75
Zoning MIS JBOD Platforms	76
Phy-based Zoning	79
Loading a CSV File to Zones	81
Creating the Drive Groups in MegaRAID (CentOS, Linux & Windows)	83
Formatting the Drives Using Command Line in CentOS	86
Verify Drives Seen	86
Adding a New Drive	86
Formatting Drives	87

Mount the New Drive 88
Formatting the Drives Using YaST2 in Linux 88
Verify Drives Seen 90
Partitioning Drives 90
Windows Server Manager 93
Verify Drives Seen 94
Formatting the Drives 94
CLI Zoning Tool, version 1.4 99
Installing CLI Zoning Tool on CentOS or Linux Host 100
Installing CLI Zoning Tool on Windows Host 100
Basics of CLI Zoning 101
Preparing to Zone using the CLI Zoning Tool 105
Editing the ShackCLI.ini file for CentOS & Linux 105
Editing the ShackCLI.ini file for Windows 107
CLI Zoning Tool Main Menu 108
Version 1.4.0 Default Menu 108
Editing the .csv File for the CLI Zoning Tool 111
Version 1.4.0 Expanded Menu 113
9) Display StorBrick(s) Settings 113
10) Display CLI Settings 114
11) Enter StorBrick CLI (Must select a single StorBrick) 115
12) Reboot Fan Base - will not reset StorBricks 115
13) Reset Fan Base - Danger. This will reset StorBricks 115
14) Force StorBrick(s) into Debug Mode 115
15) Exit StorBrick(s) Debug Mode 115
16) Display PHY Error Counters for selected StorBricks 116
17) Display PHY Information for selected StorBricks 118
18) Display StorBrick UUID for selected StorBricks 120
19) Display StorBrick Firmware Revision Levels for Selected StorBricks 120
20) Change state of JBOD-IO ID LED - off, on blink 121
21) Clear Persistent Flash 121
22) Display Fan Base Status 121
23) Display JBOD-IO Status 124

	24) Change Polling of Expanders by Fan Base or JBOD-IO125
	25) Display Disk Drive Status126
	26) Change Power State of a Disk Drive126
	27) Power Cycle test126
	28) Update StorBrick UUID126
	29) Write/Read Buffer Test126
	Disk RAID Support127
	RAID Configuration Notes128
	Spare Drives130
4	System Maintenance131
	Detecting Component Failures132
	Sliding the Chassis Forward/Backward133
	Removing the Front or Rear Chassis Cover134
	Replacing a Power Supply135
	Replacing a Fan Module136
	Replacing a Disk Drive138
	Removing the Drive139
	Re-installing the Drive139
	Installing the Rack Air Ducts140
	Replacing the Front Bezel Grille141
5	Troubleshooting145
	No Video146
	Losing the System's Setup Configuration146
	Safe Power-Off146
	"Missing" Firmware Files147
	Cannot Receive Email Alerts Using BMC148
	Path not found errors after installing Python for Windows150
	Intel BIOS POST error messages and handling description150
	Beep Codes154
	Not all beep codes signal an error155
	Fans don't spin when power button pressed.155
	Fans spin when power button is pressed, but no video display is seen155

	BIOS or board logo appear, but OS load screen never appears	155
	OS load screen appears, but nothing further	155
	OS fully loads, but errors are seen	156
	Cannot boot after new software or drivers were installed	156
	Cannot boot after settings were changed in BIOS	156
	System Gets Stuck in a Reboot Cycle	156
	Shortening the Cable Management Arm	156
A	Technical Specifications	163
B	Zone Permission Groups Rules	167

Figures

Figure 1-1	SGI Destination Rack (D-Rack)	4
Figure 1-2	Bridge Midplane	5
Figure 1-3	SAS Interposers and SAS Midplane	6
Figure 1-4	MIS Case and Front Grille	6
Figure 1-5	New Bi-directional Rolling Rail Mounts.	7
Figure 1-6	Single Control Panel	8
Figure 1-7	Dual Control Panel	8
Figure 1-8	Rear View – MIS Single Server Platform	9
Figure 1-9	Rear View – MIS Single Server Platform (four power supplies)	9
Figure 1-10	Rear View – MIS Dual Server Platform	9
Figure 1-11	Rear View – MIS Dual Server Chassis with Single Compute Server Module Platform	10
Figure 1-12	Rear View – MIS JBOD Unit (single I/O module)	10
Figure 1-13	Rear View – MIS JBOD Unit (dual I/O modules)	10
Figure 1-14	Cable Management Arm	11
Figure 1-15	Power Supply Module (rated at 1100 Watts)	13
Figure 1-16	Power Supply Numbering	13
Figure 1-17	Fan Assembly Module (each contains two impellers)	14
Figure 1-18	StorBrick Modules for 3.5" or 2.5" 15mm Drives (left) and 2.5" 9.5 mm Drives (right)	15
Figure 1-19	3.5" 15mm Drive and Carrier	16
Figure 1-20	3.5" 15mm Drive Carrier (top view, with thumb latch)	16
Figure 1-21	Two 2.5" 9.5mm Drives and Carrier	16
Figure 1-22	2.5" 9.5mm Drive Carrier (isometric view with dual thumb-latches)	17
Figure 1-23	MIS 1.5 Single Server Platform	18
Figure 1-24	MIS 1.5 Dual Server Platform	19
Figure 1-25	MIS 1.5 JBOD Platform	20
Figure 1-26	MIS Single Server Module	21

Figure 1-27	Single Server Module – Component View 22
Figure 1-28	MIS Dual Server Module – Half Height 22
Figure 1-29	Dual Server Module – Component View 23
Figure 1-30	CPU and PCIe Riser Layout 24
Figure 1-31	CPU Riser and HBA layout 25
Figure 1-32	Boot Drive Module 26
Figure 1-33	I/O Module for MIS JBOD Unit 26
Figure 1-34	MIS JBOD Midplane I/O Connector (right & left views) 27
Figure 1-35	JBOD I/O Module Port Diagram 27
Figure 1-36	System-Level Block Diagram – Server 29
Figure 1-37	System-Level Block Diagram – JBOD 30
Figure 1-38	System Layout 31
Figure 1-39	SAS Storage Topology, PCAs and Cables 32
Figure 2-1	MIS Control Panel 34
Figure 2-2	Disk Drive Light Pipes 38
Figure 2-3	Power Supply LEDs 38
Figure 2-4	Fan Base Service Page 42
Figure 2-5	MIS-S9D Proprietary Network Interface 43
Figure 2-6	BMC Web Console – Logout, Refresh, and Help Buttons 44
Figure 2-7	BMC Web Console – System Debug Log 45
Figure 2-8	BMC Web Console – Server Health 46
Figure 2-9	BMC Web Console – Event Log 48
Figure 2-10	BMC Web Console – Power Statistics 49
Figure 2-11	BMC Web Console – Alerts 51
Figure 2-12	BMC Web Console – Alert Email Settings 52
Figure 2-13	BMC Web Console – Alert Email Result. 53
Figure 2-14	BMC Web Console – Power Control and Status 55
Figure 3-1	Zones 1.4.2 – Current Zoning Configuration 70
Figure 3-2	Zones 1.4.2 – Bottom Row of Buttons 70
Figure 3-3	Zones 1.4.2 – Step 1: Opening a Session 71
Figure 3-4	Zones 1.4.2 – New Session Alias 72
Figure 3-5	Zones 1.4.2 – Warning Box, Alias Rules 72

Figure 3-6	Zones 1.4.2 – Top Menu Bar and Adapter Tabs	73
Figure 3-7	Zones 1.4.2 – Setup Menu Options	74
Figure 3-8	Zones 1.4.2 – Warning 144 Drive Configuration	74
Figure 3-9	StorBrick Diagram for 144-Drive Systems	75
Figure 3-10	Zones 1.4.2 – Brick Selection Window	76
Figure 3-11	Zones 1.4.2 – JBOD Zoning Interface	77
Figure 3-12	Zones 1.4.2 – JBOD Open Session Window.	78
Figure 3-13	Zones 1.4.2 – JBOD New Session Alias	79
Figure 3-14	Zones 1.4.2 – JBOD Phy-based Zoning Interface	80
Figure 3-15	Zones 1.4.2 – JBOD Brick Selection Window	81
Figure 3-16	Zones 1.4.2 – Choose CSV File Window	82
Figure 3-17	Zones 1.4.2 – CSV Load: Primary Only or Primary and Secondary.	82
Figure 3-18	Zones 1.4.2 – CSV Warning Window	83
Figure 3-19	MegaRAID – Create a Virtual Drive.	84
Figure 3-20	MegaRAID – Create Virtual Drive Mode	84
Figure 3-21	Create Virtual Drive – Drive Group Settings	85
Figure 3-22	Create Virtual Drive – Summary.	85
Figure 3-23	YaST2 – Server Manager GUI	89
Figure 3-24	YaST2 – Warning Message	90
Figure 3-25	YaST2 – Drives Have Appeared.	90
Figure 3-26	YaST2 – Add Button.	91
Figure 3-27	YaST2 – Select Partition Size	91
Figure 3-28	YaST2 – Format & Mount the Drive.	91
Figure 3-29	YaST2 – Check for Partition	92
Figure 3-30	YaST2 – Click Finish	92
Figure 3-31	YaST2 – Disk Mounting (in process)	93
Figure 3-32	Windows Server Manager – Disk Management	93
Figure 3-33	Server Manager – Disk Management	94
Figure 3-34	Server Manager – Initialize Disks	95
Figure 3-35	Server Manager – Select GPT (GUID Partition Table)	95
Figure 3-36	Server Manager – Disks Initialized and Online	96
Figure 3-37	Server Manager – New Simple Volume	96
Figure 3-38	Server Manager – New Simple Volume Wizard.	97

Figure 3-39	New Simple Volume Wizard – Volume Size 97
Figure 3-40	New Simple Volume Wizard – Assign Drive Letter or Path 98
Figure 3-41	New Simple Volume Wizard – Format Partition 98
Figure 3-42	New Simple Volume Wizard – Settings Confirmation 99
Figure 3-43	Server Manager – New Simple Volume 99
Figure 3-44	MIS-S9D Proprietary Network Interface	105
Figure 3-45	Block diagram of MIS Server StorBrick SB0	112
Figure 3-46	RAID 1 with One Drive per StorBrick	128
Figure 3-47	RAID 1 with Two Drives Spanning a StorBrick	128
Figure 3-48	RAID 5 or 6 with One Drive per StorBrick	129
Figure 3-49	Loss of a Drive with Multiple Drives on a StorBrick.	130
Figure 3-50	Three Drive Loss in RAID 6 Requires StorBrick Replacement	130
Figure 4-1	MIS-S9D Proprietary Network Interface	133
Figure 4-2	Front & Rear Chassis Covers	134
Figure 4-3	Power Supply Module	135
Figure 4-4	Fan Module	136
Figure 4-5	Fan Base Service Page	137
Figure 4-6	Hard Drive Carrier	139
Figure 4-7	Rack Air Duct Placement	141
Figure 4-8	Front Bezel Grille Control Panel Sockets and S9D Network Interface	142
Figure 4-9	Chassis Front Cables	142
Figure 4-10	Control Panel Ribbon Cables Attached to Front Bezel	143
Figure 4-11	S9D Network Interface Cable.	143
Figure 4-12	Front Bezel Grille Attach Tabs	144
Figure 5-1	BMC Alerts Email Settings Page	148
Figure 5-2	BMC Alert Configuration Success Pop-up	148
Figure 5-3	BMC Alerts Configuration Page	149
Figure 5-4	Cable Management Arm	157
Figure 5-5	Remove Segments from the End of an Empty Arm	158
Figure 5-6	Pry the Joint Apart	158
Figure 5-7	Pull the Segments from the Arm	159
Figure 5-8	Shortened Arm Ready for Installation	159
Figure 5-9	Remove Segments from the Middle of a Full Arm	160

Figure 5-10	Pry Open the Segment Joint160
Figure 5-11	Pull Apart the Cable Management Arm161
Figure 5-12	Pry Apart the Joint at the Other End161
Figure 5-13	Removing the Segment from the Cables.162
Figure 5-14	Join the Spliced Ends of the Cable Management Arm162
Figure B-1	Zone Permission Groups – Example 1168
Figure B-2	Zone Permission Groups – Example 2169

Tables

Table 0-1	MIS Server Platform Region and EMC Compliance References . . .	xxiii
Table 0-2	MIS Server Platform Region and Functional Safety Listing Mark . . .	xxiv
Table 2-1	MIS Server Platform Control Panel Buttons and LEDs	35
Table 2-2	Disk Drive LEDs	38
Table 2-3	Power Supply LEDs	39
Table 2-4	Server Power Control Actions	54
Table 2-5	System Information Details	55
Table 3-1	Zoning Configuration101
Table 3-2	Phy Based Zoning Table for StorBrick 0103
Table 3-3	Zone Group Implementation104
Table 3-4	CLI Zoning Tool Menu Options and Descriptions109
Table 3-5	Zone Group Implementation111
Table 5-1	Intel BIOS POST Error Messages150
Table 5-2	Summary of LSI MegaRAID Card Beep Codes154
Table A-1	Technical Specifications163

Introduction

This guide describes the features and components of the SGI® Modular InfiniteStorage™ (MIS) platform, version 1.5. With two main configurations possible for the enclosure (storage server or JBOD—Just a Bunch Of Disks) this guide covers the different configurations, their respective components, interface panels, indicator lights and meanings, software, zoning, maintenance, and troubleshooting.

What's New



Version 1.5 of the MIS platform features new interior hardware with increased ease-of-use and serviceability, with improved air-flow. This manual details information for the new system, and does not require you to be familiar with version 1.0.

Audience

This guide is written for owners/users of the MIS platform, version 1.5. It is written with the assumption that the reader has a good working knowledge of computers, servers, networking, hardware, software and RAID arrays.

Important Information

The following section details several safety precautions that should be observed at all times. First, a fully loaded MIS Platform can weigh up to 220lbs. Second, electricity is a major concern when handling components, especially Electrostatic Discharge (ESD), detailed later in this section. Please read these sections carefully prior to using the MIS Platform.



Environmental Requirements

The following are the environmental requirements for the MIS platforms.

Operating temperature: **41° to 95° F [5° to 35° C] (up to 5,000 ft. [1,500 m])**

- Derate max temperature (95° F [35° C]) by 1.8° Fahrenheit [1° Celsius] per 1,000 ft. [305 m] of altitude above 5,000 ft. [1525 m]
- Temperature rate of change must not exceed 18° F [10° C] per hour

Operating humidity: **8% to 80% non-condensing**

- Humidity rate of change must not exceed 10% relative humidity per hour

Operating altitude: **up to 10,000 ft. [up to 3,050 m]**

Shipping temperature: **-40° to 140° F [-40° to 60° C]**

- Temperature rate of change must not exceed 36° F [20° C] per hour

Shipping humidity: **10% to 95% non-condensing**

Shipping altitude: **up to 40,000 ft. (up to 12,200 m)**

Storage temperature: **41° F [5° C] to 113° F [45° C]**

– **Temperature rate of change must not exceed 36° F [20° C] per hour**

Storage humidity: **8% to 80% non-condensing**

Storage altitude: **up to 40,000 ft. [up to 12,200 m]**

Safety Precautions

Do NOT wear loose clothing, such as neckties or unbuttoned shirt sleeves, while working on the unit which can be pulled into a cooling fan or tangled in cabling.

Remove any jewelry and any metal objects from your body, which are excellent electrical conductors, and can harm you and/or cause short circuits if they come into contact with printed circuit boards or powered areas.

Be aware of the locations of the power on/off switch on the chassis as well as the room's emergency power-off switch, disconnection switch or electrical outlet. If an electrical accident occurs, you can then quickly remove power from the system.

Do NOT work alone when working with high voltage components.

When working around exposed electrical circuits, another person should be nearby, who is familiar with the power-off controls, to switch off the power if necessary.

Use only one hand when working with powered-on electrical equipment. This is to avoid making a complete circuit, which will cause electrical shock. Use extreme caution when using metal tools, which can easily damage any electrical components or circuit boards with which they come into contact.

Do NOT use mats designed to decrease static electrical discharge as protection from electrical shock. Instead, use rubber mats that have been specifically designed as **electrical insulators**.

The power supply power cords must include a grounding plug and must be plugged into grounded electrical outlets.

Do NOT attempt to transport/move a fully loaded MIS system. An MIS system can weigh up to 220lbs. when fully loaded. If the system must be moved, first remove the drives from the chassis. When lifting the system, two people (one at each end) should lift slowly with feet spread apart to distribute the weight. Always follow safe lifting practices when moving heavy objects. More information on moving large objects, requiring a two-person team, is available in the Centers for

Disease Control's, "Ergonomic Guidelines for Manual Material Handling"
(<http://www.cdc.gov/niosh/docs/2007-131/pdfs/2007-131.pdf>)

Power should always be disconnected from the system when removing or installing system components that are not hot-swappable, such as server boards and memory modules. When disconnecting power, you should first do a clean shut down of the operating system, then power down the system, and then unplug all power cords (the unit has more than one power supply cord).

ESD Precautions



Caution: Electrostatic Discharge (ESD) is generated by two objects with different electrical charges coming into contact with each other. An electrical discharge is created to neutralize this difference, which can damage electronic components and printed circuit boards.

The following measures are generally sufficient to neutralize this difference before contact is made to protect your equipment from ESD:

- Use a grounded wrist strap designed to prevent static discharge.
- Keep all components and printed circuit boards (PCBs) in their antistatic bags until ready for use.
- Touch a grounded metal object before removing the board from the antistatic bag.
- Do not let components or PCBs come into contact with your clothing, which may retain a charge even if you are wearing a wrist strap.
- Handle a board by its edges only; do not touch its components, peripheral chips, memory modules or contacts.
- When handling chips or modules, avoid touching their pins.
- Put the server board and peripherals back into their antistatic bags when not in use.
- For grounding purposes, make sure your computer chassis provides excellent conductivity between the power supply, the case, the mounting fasteners and the server board.

Safety & Emissions

The following is a list of agency approvals for MIS on safety and emissions.

Electromagnetic Compatibility

Table 0-1 lists the region and compliance reference for EMC (Electromagnetic Compatibility) compliance.

Table 0-1 MIS Server Platform Region and EMC Compliance References

Region	Compliance Reference
Australia/ New Zealand	AS/NZS CISPR 22 (Emissions)
USA/Canada	Industry Canada ICES-003 Issue 4 FCC CFR47, Part 15 Subpart B
CENELEC Europe	EN55022 Emissions EN55024 Immunity
International	CISPR 22/ CISPR 24
Japan	VCCI V-3 Certification
Korea	KCC KN22, KN24 Verification
Taiwan	BSMI CNS 13438 Verification
Russia	EAC

Safety Certification

National Recognized Testing Laboratory (NRTL) provides the safety certification for ITE products. NRTL's offer various product markings depending on the type of products being tested and satisfactory tests results. Underwriters Laboratories (UL) and Canadian Standards Association (CSA) are typical (NRTL's) that provides safety certification services for Information

Technology Equipment (ITE). Table 0-2 lists the region and compliance reference for safety compliance.

Table 0-2 MIS Server Platform Region and Functional Safety Listing Mark

Region	Compliance Reference
USA/Canada	CSA 60950 / UL60950/ 60950-1 cert to CAN/CSA STD C22.2 No. 60950-1
IEC (Europe)	IEC60950-1 – CB Certification, CE Mark
Russia	EAC

Chapter Descriptions

Chapter 1, “System Overview,” introduces new features of the MIS 1.5, describes the hardware components of the MIS enclosures, the common modules in unit, and the major differences between the MIS Server Platform and MIS JBOD Unit. Included at the end of the chapter are a system block diagram and a system layout diagram.

Chapter 2, “System Interfaces,” describes the hardware and software interfaces used to operate the MIS Server and MIS JBOD. This includes the front control panel, disk drive LEDs, power supply LEDs, and the BMC Web Console.

Chapter 3, “System Software,” covers the software used on the MIS Platforms, including installation information for the tools, using the MegaRAID tool, and the available zoning software. In this chapter are step-by-step instructions for Zones 1.3 tool, its features and their function, plus warnings and error codes. Screen shots are given for a walk-through of the step-by-step instructions for zoning. Information on using the CLI Zoning tool 1.4 is also given, though Zones 1.4.2 can handle the zoning of both servers and JBODs at this time. Some considerations for RAID Configurations are given at the end of the chapter.

Chapter 4, “System Maintenance,” describes how to use the BMC for detecting component failures, and service instructions for modules that are customer replaceable units (CRUs). The service instructions include how to move the chassis forward or backwards in the rack, how to remove the case front and rear covers, how to replace a power supply, how to replace a fan module, how to replace a storage drive in a StorBrick, how to replace a boot drive, and how to install the air ducts to provide proper air flow to the chassis.

Chapter 5, “Troubleshooting,” describes some problem-solving techniques, plus when to contact customer support.

Appendix A, “Technical Specifications,” gives the technical specifications for the MIS enclosures.

Appendix B, “BIOS Error Codes,” details the beep codes used when a problem is detected by the BMC environmental controls.

Appendix B, “Zone Permission Groups Rules,” explains bit assignment rules for the Master, Initiator, and Zone Permission Groups, as well as some advanced Zoning concepts.

Related Publications

The following documents are relevant to the MIS Platform:

- *MegaRAID SAS Software User Guide*, publication number 51530-0x.
- *MegaRAID 6Gb/s SAS RAID Controllers User Guide*, publication number 41450-0x.
- *Intel Server Boards and Server Platforms Server Management Guide*, publication number 37830-002
- *Intel® Remote Management Module 4 and Integrated BMC Web Console User Guide* (pdf) Revision 2.6 (also refer to the **Help** files located within the BMC itself)
- SGI InfiniteStorage series documentation (<http://techpubs.sgi.com>)
- Man pages (<http://www.linuxmanpages.com/>)

Various formats of SGI documentation, release notes, and man pages are available. The SGI Technical Publications Library (<http://docs.sgi.com/>) contains the most recent and most comprehensive set of online books, release notes, man pages, and other information. Refer to the SGI Supportfolio™ web page for documents which access requires a support contract (as do the MegaRAID books cited above). See “Product Support” on page xxvi. You can also view man pages by typing `man <title>` on a command line in Linux.

Conventions

The following conventions are used throughout this document:

Convention	Meaning
Command	This fixed-space font denotes literal items such as commands, files, routines, path names, signals, messages, and programming language structures.
<i>variable</i>	The italic typeface denotes variable entries and words or concepts being defined. Italic typeface is also used for book titles.
[]	Brackets enclose optional portions of a command or directive line.
GUI element	This font denotes the names of graphical user interface (GUI) elements such as windows, screens, dialog boxes, menus, toolbars, icons, buttons, boxes, fields, and lists.

Product Support

SGI provides a comprehensive product support and maintenance program for its products, as follows:

- If you are in North America, contact the Technical Assistance Center at +1 800 800 4SGI (4744) or contact your authorized service provider.
- If you are outside North America, contact the SGI subsidiary or authorized distributor in your country. International customers can visit <http://www.sgi.com/support/> Click on the “Support Centers” link under the “Online Support” heading for information on how to contact your nearest SGI customer support center.

Software Downloads

SGI provides all the software that has been approved for use on the MIS platforms, including the SGI approved versions of 3rd party software. These SGI approved versions have been tested for quality assurance and hardware/software/firmware compatibility with the MIS. To obtain software, either contact SGI Customer Support for the most recent approved versions, or download from the Supportfolio web site: <http://support.sgi.com> (requires Supportfolio account and password for access).

Warning: Do NOT download or install software directly from 3rd party sites. This may cause **software conflicts that result in data loss.**

CRU/FRU

Some of the components on the MIS Platform are customer-replaceable units (CRUs), meaning that these modules were designed to be repaired/replaced by you, the customer. These include fan assemblies, power supplies, storage drives, and boot drives, all of which are hot-swappable. However, many of the other components on the MIS Platform should be serviced by SGI field technicians ONLY, so as not to violate the warranty agreement. The components are field-technician replaceable units, or FRUs. It is important to note that our CRUs can be easily installed and replaced by customers, which enables a speedy recovery of proper system operation.

For additional information about CRUs, please see:

- [Customer Replaceable Units \(CRUs\) Installation Policy](http://www.sgi.com/services/support/cru/policy.html)
(<http://www.sgi.com/services/support/cru/policy.html>)
- [Customer Replaceable Units \(CRU\) and Customer Obligations](http://www.sgi.com/services/support/cru/obligations.html)
(<http://www.sgi.com/services/support/cru/obligations.html>)

Purchasable Support & Maintenance Programs

SGI provides several comprehensive product support and maintenance programs for its products. SGI also offers services to implement and integrate Linux applications in your environment.

- Refer to <http://www.sgi.com/services/>
- If you are in North America, contact the Technical Assistance Center at +1-800-800-4SGI (4744), or contact your authorized service provider.
- If you are outside North America, contact the SGI subsidiary or authorized distributor in your country. See <http://www.sgi.com/global/index.html> for more information.

Reader Comments

If you have comments about the technical accuracy, content, or organization of this document, please contact SGI. Be sure to include the title and document number of the manual with your comments. (Online, the document number is located in the front matter of the manual. In printed manuals, the document number is located at the bottom of each page.)

You can contact SGI in any of the following ways:

- Send e-mail to the following address: techpubs@sgi.com
- Contact your customer service representative, and ask that an incident be filed in the SGI incident tracking system.
- Send mail to the following address:

SGI
Technical Publications
46600 Landing Parkway
Fremont, CA 94538

SGI values your comments, and will respond to them promptly.

System Overview

The SGI Modular InfiniteStorage Platform, Version 1.5, is a high-density, integrated storage server platform. The MIS 1.5 uses a 4U rackmount system, and can be either a compute and storage server, or “Just a Bunch Of Disks” expansion storage unit (MIS JBOD unit). The MIS Server Platform can be single or dual server.

Features of the MIS 1.5 Platform include:

- Up to 72 (3.5" or 2.5" 15mm) and a maximum of 144 (2.5" 9.5mm) storage drives in the Server Platform.
- Up to 81 (3.5" or 2.5" 15mm) and a maximum of 162 (2.5" 9.5mm) storage drives in the JBOD unit.
- All fit in a standard size 4U chassis: height 6.94" (176mm), width 16.9" (429.2mm), depth 36" (914.4mm).
- Storage drives can be 3.5" or 2.5" (15mm or 9.5mm), SAS or SATA, rotational or SSD drives. Up to four JBOD units per MIS Server Platform (requires server have 2 JBOD HBAs).



Warning: Rotational SAS drives and rotational SATA drives cannot be included in the same enclosure due to vibration conflicts.

New features of the modular design of the MIS 1.5 Platform include:

- Higher capacity fan modules for more efficient cooling.
- Interior foam strips and fan modules maximize air flow and thermal efficiency.
- Server modules feature new CPU air duct for improved thermals.
- Hot-swappable server modules reduce service actions from four hours to fifteen minutes.

- Point-to-point, straight-through connections to StorBricks take up less space, improve signal quality and air flow.
- Cable-to-PCBA interposers create better signal quality and easier routing.

The MIS Server Platform features:

- Up to 2 server modules per platform.
 - Two Intel® Xeon® E5-2600 series processors per server motherboard, S2600JF family.
 - Intel Turbo Boost Technology 2.0 that automatically allows processor cores to run faster than the base operating frequency, if the cores are operating below power, current, and temperature specification limits (< 35°C ambient).
 - Dual GbE networking onboard.
- Up to 8 DDR3 DIMMs (new: 8 GB, 16 GB, or 32 GB) for a single-server motherboard configuration, and up to 16 DIMMs for a dual-server motherboard configuration.
- Up to 6 HBAs for single servers—these can be full-height (4.25") and half-depth (3.375"), up to 4 externally facing and/or up to 4 internally facing (total external and internal not to exceed 6). For single-server platforms only, optional 4 port¹ 8Gb Fibre Channel PCIe cards.
- Up to 4 HBAs per server module (for up to a total of 8 across both server modules) in a dual server system—these can be half-height (2.12"), half-depth (3.375"), 2 internal max and/or a maximum of 2 external.
- Unique Battery Back-Up PCIe technology allows the optional addition of BBUs without the consumption of any of the available PCIe slots.
- Two boot drive modules per server, with up to two boot drives per module.

Chapter 1: System Overview will first explore “New Hardware Features of MIS 1.5” on page 5 detailing

- the new “Bridge Midplane” on page 5
- and new “SAS Interposers and SAS Midplane” on page 5.

Following is information on the “MIS Enclosure” on page 6 including

- “Front Grille and Control Panels” on page 7,

¹ The four port card cannot be used in the dual server as it only supports half height cards. The dual server must use the 2 port Fibre Channel card. The four port can be used in the single server.

-
- “Rear Panel Components” on page 8,
 - and the “Cable Management Arm” on page 11

Next, the “MIS Common Modules” on page 12 are covered, including:

- “Power Supply Modules” on page 12,
- new “Fan Assembly Modules” on page 14,
- and “StorBrick Modules” on page 14.

“MIS Server Platform or JBOD Unit” on page 17, discusses the presence of:

- “Server Module” on page 21, and its available features,
- “Layout of Server CPUs, and PCIe Risers HBAs” on page 23,
- and associated “Boot Drive Module” on page 26,
- or the presence of a ninth StorBrick and “MIS JBOD I/O Module” on page 26.

Finally, a new “System Diagrams” on page 28 shows a logical layout of the chassis components and connections, and its corresponding “System Layout” in Figure 1-38, shows the physical layout of the chassis and components.



Figure 1-1 SGI Destination Rack (D-Rack)

New Hardware Features of MIS 1.5

Bridge Midplane

The new Bridge Midplane (Figure 1-2) features a power design that allows for hot-swappable server modules, decreasing service time from four hours to fifteen minutes. The elimination of a shutdown requirement in servicing server modules also means that information can be kept available even while servicing the machines. Many service issues, now more than before, can be addressed while keeping the machines on and available.




Figure 1-2 Bridge Midplane

SAS Interposers and SAS Midplane

The new SAS Interposers and SAS Midplane work together to provide point-to-point, straight-through connections to StorBricks that take up less space, improve signal quality, and increase air flow through the system. Cable-to-PCBA interposers and flat 3M twin axial cable create better signal quality and easier routing (Figure 1-3).




Figure 1-3 SAS Interposers and SAS Midplane

MIS Enclosure

The MIS enclosure, whether it is a server or JBOD, consists of a chassis, case (with front bezel grille, control panel and rear ports).



Figure 1-4 MIS Case and Front Grille

The SGI MIS chassis features a front bezel display with an internal EMI grille (Figure 1-4). The unique bi-directional sliding rail mounts (Figure 1-5) allow the unit to be slid forwards or backwards 20" to access disk drives and other serviceable components. This also makes for an overall safer rack design, as chassis do not need to be extended to their full length to be serviced.




Figure 1-5 New Bi-directional Rolling Rail Mounts

Front Grille and Control Panels

Next to the bezel grille, up to two control panels can be present on the MIS Platform, one for each server in the MIS Server Platform, or one for each I/O unit on the MIS JBOD. Figure 1-6 shows a single control panel and Figure 1-7 shows two control panels.

Each control panel has a Power LED, Power button, Status LED, Reset Button, Locator LED, Locator button, Network Activity LED, Boot Drive Activity LED, and NMI Reset button (non-functional; not supported). Indicator light meanings and button functions are explained in, “Control Panel” in Chapter 2.




Figure 1-6 Single Control Panel




Figure 1-7 Dual Control Panel

Rear Panel Components

The appearance of the rear panel on the MIS chassis will depend on what modules are installed. An MIS Platform can have up to four power supplies, each with their own AC power inputs. They are high-efficiency, hot-swappable power supplies rated at 1100 Watts.

All rear panels feature clearly silk-screened labels next to the port in question. The MIS Server Platform (single server) in Figure 1-8, features a single server module with two USB ports, a video port, and two NIC ports. Figure 1-9 show a MIS Server Platform with the optional four power supplies. The MIS Server Platform (dual server) rear panel as shown in Figure 1-10 has a second server module with its own set of USB ports, video port, and NIC ports. Figure 1-11 shows an MIS Server Platform which features dual server construction with only one compute server module installed, and the option to upgrade to include a second compute server module later. Figure 1-12 shows the rear panel of the MIS JBOD with one I/O module and Figure 1-13 shows two, each with 4 miniSAS, one Ethernet, and one serial port (note that the first I/O module is on the bottom and the second on the top).




Figure 1-8 Rear View – MIS Single Server Platform




Figure 1-9 Rear View – MIS Single Server Platform (four power supplies)




Figure 1-10 Rear View – MIS Dual Server Platform




Figure 1-11 Rear View – MIS Dual Server Chassis with Single Compute Server Module Platform




Figure 1-12 Rear View – MIS JBOD Unit (single I/O module)




Figure 1-13 Rear View – MIS JBOD Unit (dual I/O modules)

Cable Management Arm

The cable management arm is an optional addition to the MIS chassis that helps keep cables organized in the rack. It attaches to a plate on the rear of the MIS chassis and to the rack in which the chassis is mounted. The cable management arm comes in a length designed for 26" deep racks, such as the D-Rack. It can be shortened to accommodate other rack sizes (Figure 1-14). Instructions on shortening the cable management arm are in Chapter 5, "Troubleshooting."




Figure 1-14 Cable Management Arm

MIS Common Modules

This section describes the internal modules of the MIS 1.5 platforms that both the server and JBOD share in common. Designed to deliver a high level of reliability, scalability and manageability, the MIS 1.5 platform continues to make use of modules to contain key components, and now uses interposers, midplanes and bridges to replace a significant number of cables, creating better air flow and thermals throughout the system. Whether the unit is an MIS 1.5 Server or a JBOD, both chassis contain the following hot-swappable modules:

- Two or four power supplies (if four, two are redundant) (Figure 1-15);
- Six (newly-designed, higher-capacity) fan assemblies (Figure 1-17);
- Capacity drives installed in StorBricks (Figure 1-18).

The power supply modules are high-efficiency, hot-swappable, and rated at 1100 Watts, AC Input: 100–240 VAC (50-60Hz), single or three phase. There are six hot-swappable fan modules housing one fan with two counter-rotating impellers. This fan design features easier replacement, and greater speeds of operation for a greater cooling effect. And instead of the conventional disk architecture, the unique StorBrick modules (innovative, highly-dense drive modules used to house drive bays) allow the platform to maximize storage density.

Each MIS Server has eight StorBricks modules, and each MIS JBOD has nine, with the ninth module taking the place of the compute server module. Each StorBrick module holds up to nine 3.5" or 2.5" (15mm), SAS or SATA, rotational or SSD drives, or, using the dual-slot drive option, eighteen 2.5" (9.5mm), SAS or SATA, rotational or SSD drives.

Power Supply Modules

The **1100 Watt AC to DC** power factor corrected (PFC) power supply converts standard AC mains power into a main output of 12VDC for powering intermediate bus architectures (IBA). The power supplies meet international safety standards and display the CE-Mark for the European Low Voltage Directive (LVD).

Two to four power supplies provide power for the SGI MIS server. Power supplies are configured for N+N support, meaning only two power supplies are used at any given time, with the remainder in standby (if not off or in fault). The power supplies provide 12V DC main power and 5V DC standby power.

Power supplies are hot-swappable, and can be replaced under full load, as long as there are *two or more* functioning power supplies present—this requires the redundant power option be installed.

Power supplies are numbered 0-3 from the bottom up, on the rear panel of the enclosure (Figure 1-16).

Warning: Removal of one power supply when there are *only two* power supplies present may cause an **unexpected shut-down that may result in data loss.**




Figure 1-15 Power Supply Module (rated at 1100 Watts)




Figure 1-16 Power Supply Numbering

The high efficiency is achieved by using state-of-the-art silicon power devices in conjunction with soft-transition topologies minimizing switching losses and a full digital control scheme. Synchronous rectifiers on the output reduce the losses in the high current output path. The rpm of the internal fan is digitally controlled to keep all components at an optimal operating temperature regardless of the ambient temperature and load conditions. Quick-acting 16-amp input fuses (5 × 20 mm) inside the power supply protect against severe defects (fuses are not accessible from the outside and are therefore not serviceable parts; instead, power supplies are replaced by swapping out the module entirely).

Important: The maximum over current protection is 50A per power supply. The overcurrent protection must be provided by the facility power grid. All national and local electrical codes apply. Check with your local licensed electrician for details.

Fan Assembly Modules




Figure 1-17 Fan Assembly Module (each contains two impellers)

The fan modules of the MIS 1.5 feature a higher capacity fan for better thermal performance. There are six fan assemblies mounted in the middle of the MIS chassis used to cool the system. Each hot-swappable fan assembly contains two impellers. Air flows from the front to the back of the enclosure. The fan baseboard distributes power and control signals to the fan assemblies. Firmware on the fan baseboard monitors the fan speeds and temperatures within the enclosure. Fan speeds are kept at maximum speed to continuously provide optimal cooling for the enclosure.

StorBrick Modules

Each StorBrick Module contains up to nine 3.5" or 2.5" 15mm drives (Figure 1-18), or eighteen 2.5" 9mm drives (Figure 1-21 on page 16), mounted in the StorBrick using proprietary drive carriers (Figure 1-19). A sliding thumb-latch securely fastens the drive carriers in place (Figure 1-20, thumb-latch is pictured in blue, but is grey on the actual product). StorBricks use

SAS-2 protocol, which enables the system to use SAS or SATA drives (rotational), or SAS *and* SATA solid-state drives.




Figure 1-18 StorBrick Modules for 3.5" or 2.5" 15mm Drives (left) and 2.5" 9.5 mm Drives (right)




Figure 1-19 3.5" 15mm Drive and Carrier



Figure 1-20 3.5" 15mm Drive Carrier (top view, with thumb latch)




Figure 1-21 Two 2.5" 9.5mm Drives and Carrier




Figure 1-22 2.5" 9.5mm Drive Carrier (isometric view with dual thumb-latches)

MIS Server Platform or JBOD Unit

The key difference between the MIS Server Platform and the MIS JBOD Unit is the presence of the compute server module (Figure 1-26 or Figure 1-28) and boot drives (Figure 1-32) in the Server Platform, or a ninth StorBrick (Figure 1-18) and I/O modules (Figure 1-33) and associated midplane (Figure 1-34) in the JBOD.




Figure 1-23 MIS 1.5 Single Server Platform




Figure 1-24 MIS 1.5 Dual Server Platform




Figure 1-25 MIS 1.5 JBOD Platform

Server Module

The MIS Server Platform can be single- or dual-server (Figure 1-26 or Figure 1-28) depending on whether it has one or two compute server modules. Each compute server module can have:

- Up to two Intel® Xeon® E5-2600 series processors per motherboard, with Intel Turbo Boost Technology 2.0: if the cores are operating below power, current, and temperature specs (< 35°C ambient) limits, they automatically run faster than base operating speed.
- 8 DDR3 DIMMs (now 8 GB, 16 GB, or 32 GB) for a single-server board configuration. Up to 16 DIMMs for a dual-server board configuration,
- Up to 6 HBAs for single servers—these can be full-height (4.25") and half-depth (3.375"), up to 4 externally facing and/or up to 4 internally facing (total external and internal not to exceed 6). For single-server platforms only, optional 4 port 8Gb Fibre Channel PCIe cards.
- Up to 4 HBAs per server module (for up to a total of 8 across both server modules) in a dual server system—these can be half-height (2.12"), half-depth (3.375"), 2 internal max and/or a maximum of 2 external. (See Figure 1-30 on page 24)
- Up to three PCIe riser cards for a single server (dual servers have a mandatory 3 PCIe risers).
- Up to four battery back up units for a single server module. Up to three battery back up units per server module for a dual server platform, for a maximum total of six. (Unique BBU PCIe technology allows the inclusion of BBUs without the consumption of any of the available PCIe slots.)




Figure 1-26 MIS Single Server Module




Figure 1-27 Single Server Module – Component View




Figure 1-28 MIS Dual Server Module – Half Height




Figure 1-29 Dual Server Module – Component View

Layout of Server CPUs, and PCIe Risers HBAs

Figure 1-30 shows the CPU and riser layout. The first CPU handles Riser 1 and 2. The second CPU manages Riser 3. HBAs populated on Riser 3 are internal facing SAS HBAs only, which connect to the StorBricks (Figure 1-39).




Figure 1-30 CPU and PCIe Riser Layout

Turbo Boost

Mode 1. (Turbo Mode Off) System performance and operation in non-turbo mode and at specific environmental conditions – The Intel CPUs have various power states and thermal protection states built-in that dictate the performance characteristics of an individual socket. The largest impact when operating in a system that was designed to support the TDP power level is the operating temperature of the individual socket and its thermal management device(s) in that specific product for a set of specified environmental conditions. The processors will go into a “throttling” mode whenever the temperature reaches a limit that is coded in to the SKU by Intel. When the Intel CPU invokes this throttling mechanism the operating frequency of the CPU is reduced and will go through a predefined performance reduction algorithm to attempt to bring the processor back within its thermal operating window. With a MIS operating with turbo mode off at an ambient temperature of 35°C or less, this throttling condition should not exist.

Mode 2. (Turbo Mode On) Enabling Turbo mode will provide the opportunity for the customer to recognize performance improvements by increasing the processor clock speed up to the Turbo limits. It is important to understand that any performance increase is variable and dependent on these environmental variables and the current activity state of the unit. This performance improvement will vary significantly from data center to data center and between specific applications. For a MIS operating with turbo mode enabled at an ambient temperature of 35°C or

less, there may be some very heavy workloads that cause the processor to throttle, but even during this throttling condition, the processor frequency will still be at or above its rated frequency.

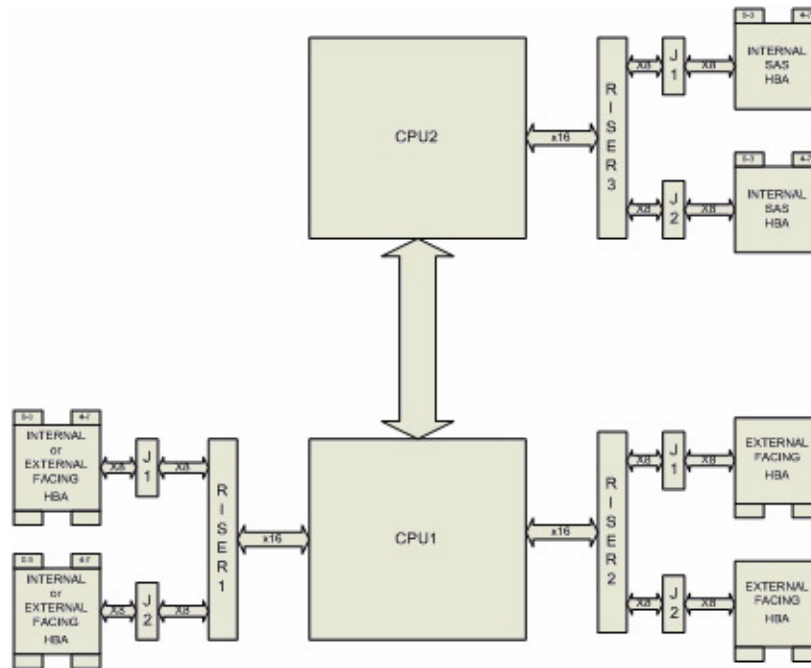


Figure 1-31 CPU Riser and HBA layout

Boot Drive Module

Each MIS Server Platform features two boot drives per server module (up to four total – mirrored using LSI software RAID 1). These drives are SAS or SATA, rotational or SSD, up to 300GB, used to store server data and the server operating system. Supported operating systems include:

- Microsoft® Windows® 2008 R2 SP1 (not shipped with product),
- Red Hat® Enterprise Linux (RHEL) 6.2,
- SUSE LINUX® Enterprise Server 11 SP1,
- VMware® ESXi 5.0, or
- CentOS 6.3 (not shipped with product).




Figure 1-32 Boot Drive Module

MIS JBOD I/O Module




Figure 1-33 I/O Module for MIS JBOD Unit

JBOD I/O modules slide into a midplane (Figure 1-34), which connect to the SAS controllers.




Figure 1-34 MIS JBOD Midplane I/O Connector (right & left views)

Connecting to a JBOD is normally done through a network connection to the JBOD I/O module. A diagram of the JBOD I/O module rear panel shows where you can connect to the module using Ethernet cable (Figure 1-35), HD miniSAS, or a serial connection. The ports labeled 0-3 are the HD miniSAS slots to the host system or other JBODs. There is also a serial port available. Whether networked to an MIS Server or not, the JBOD I/O requires DHCP for its IP address.




Figure 1-35 JBOD I/O Module Port Diagram

The status LEDs include a chassis locator LED, status/alert LED, and network activity lights.

System Diagrams

Figure 1-36 shows the system-level block diagram for a fully populated dual-server and Figure 1-37 on page 30 for a fully populated JBOD.

Important: Due to thermal requirements of proper air flow through the system, drive slots 0, 1, and 2 must always be populated in every StorBrick, with drives all the same model and size form factor (all 3.5" drives, all 2.5"-15mm drives in carriers with adapters, or all 2.5"-9mm drives in carriers with adapters). Additional drives can then be added in groups of 8 for either single- or dual-server platforms, and groups of 9 for the JBOD, starting with slot 3 of StorBrick 0 through StorBrick 7 (or SB8 for JBODs). The next set of additional drives is added to slot 4 in all the StorBricks, etc.

Figure 1-38 on page 31 is a map of the physical location of the modules, including top and rear views. In the MIS Server 1.5, the computer server module(s) take(s) the place of the JBOD's 9th StorBrick and I/O modules. Power supplies remain in the same location regardless of platform type (server or JBOD). Figure 1-39 on page 32 shows the SAS storage topology, PCAs and cables.

Note: Boot drives CM1-0 and CM1-1 are installed in dual-server platforms only.

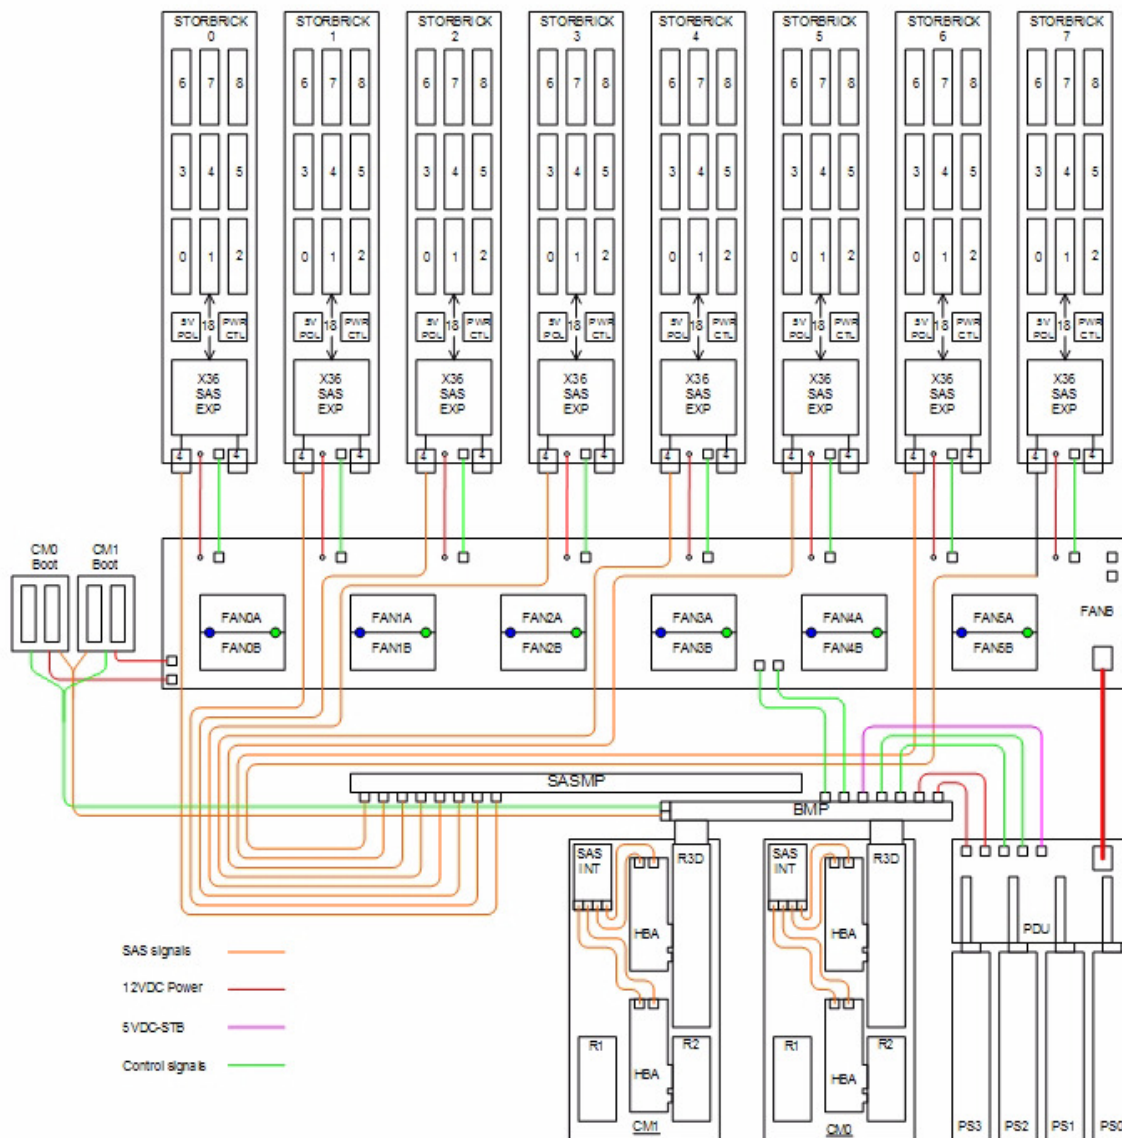


Figure 1-36 System-Level Block Diagram – Server

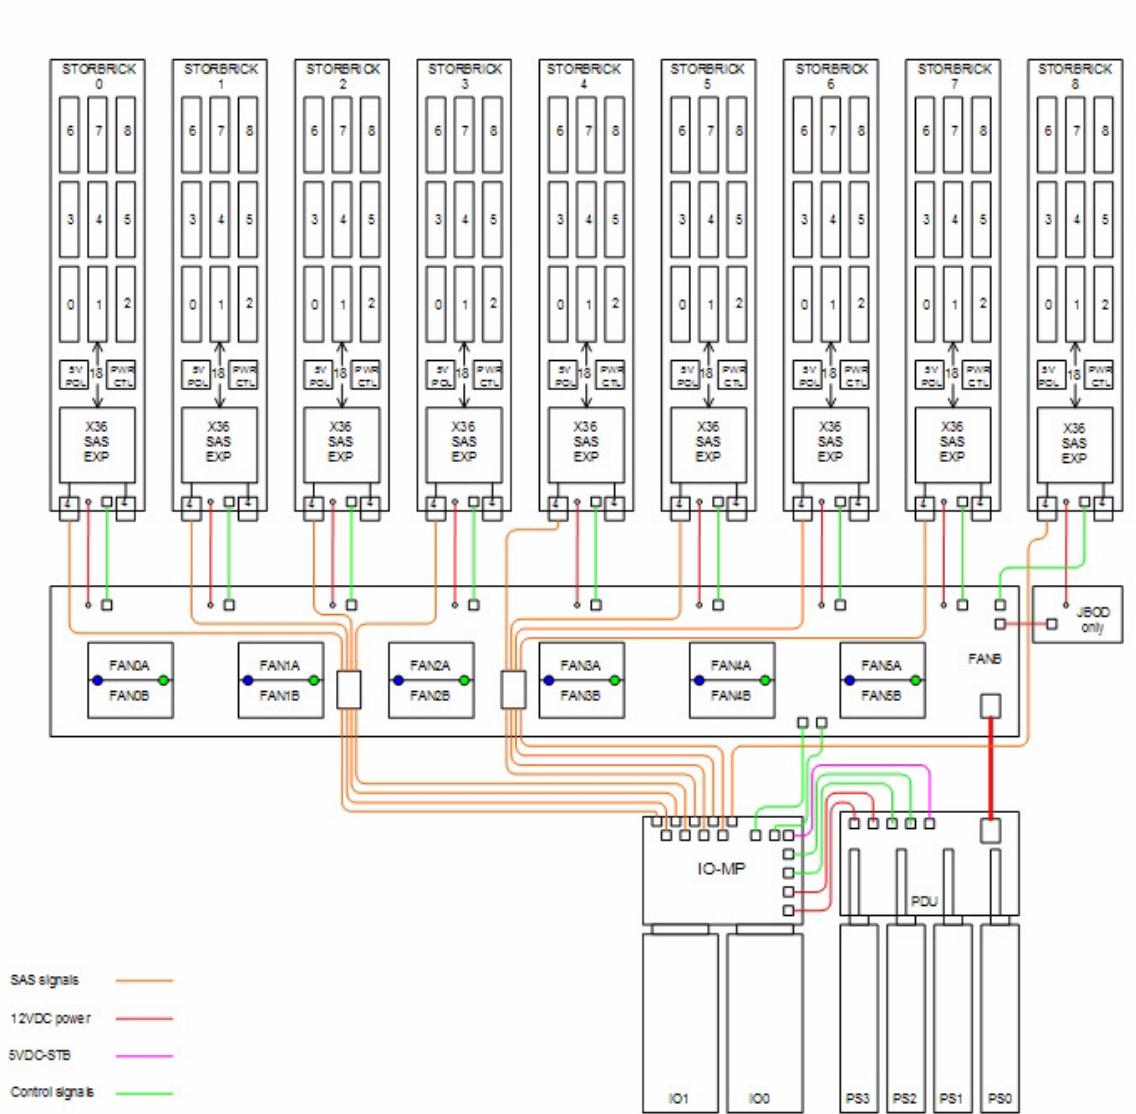


Figure 1-37 System-Level Block Diagram – JBOD




Figure 1-38 System Layout




Figure 1-39 SAS Storage Topology, PCAs and Cables

Chapter 2

System Interfaces

This chapter describes the hardware and software interfaces of the MIS platforms. Both the MIS server platform and MIS JBOD storage unit have a front control panel, disk drive LED codes, and power supply LED codes. The control panel lights and buttons have different meanings and functions, depending on whether the machine is the MIS Server Platform or MIS JBOD unit. The disk drive LED codes and power supply LED codes remain the same whether the system is a server platform or JBOD unit. Additionally, there are programs used to initialize and monitor the MIS machines. This chapter details the hardware interfaces, their functions and indications, as well as the Baseboard Management Controllers (BMC) Web Console. These programs provide power management features, environmental monitoring, etc.

Note: SGI provides features beyond those of IPMI 2.0, for instance, chassis intrusion detection, which will gracefully power down the system if the case cover is left off for more than 15 minutes.

Control Panel

MIS Server Control Panel

The control panel (Figure 2-1) interface consists of five indicator lights and four buttons. More information on remote functionality and the BMC Web Console is presented at the end of this chapter.




Figure 2-1 MIS Control Panel

The NMI Reset button (Non-Maskable Interrupt) is not supported and is non-functional.

Table 2-1 MIS Server Platform Control Panel Buttons and LEDs

LED/Button	Description
Power LED	Green LED lit means power is on.
Power button	<p>If the system is off, pushing this button powers on the system. If the operating system is running, pushing this button shuts down the operating system and powers off the system gracefully. If the operating system is hung, holding the button down for 10 seconds or more will power off the system for a hard reset.</p> <hr/> <p>Note: The power button DOES NOT completely turn off the system AC power, 5V standby power is active whenever the system is plugged in.</p> <hr/>
Status LED	This indicator will be lit whenever there is AC power available to the power supplies, whether the unit is on or off. Green means the system is in good working order. Yellow indicates a problem with the system, and service is required.
Reset button	The service reset button. When pushed, this button causes the server to reboot and, if the problem is cleared by the reset, returns the Status LED to green.
Locator LED	Blue LED is lit on the front and the back to help locate the unit in a rack or bay.
Locator button	<p>The Locator LED will be lit blue when the Locator button is pushed. There is a corresponding LED on the back of the server that will be blue. When the Locator button is pushed again, the LED will go off. This function may also be activated remotely using the Intel BMC Web Console and pressing the virtual button, or using the Linux IPMI Terminal Tool:</p> <pre>-H <ip address> -P <password> -U <user> chassis identify and identify off to turn off</pre>

Table 2-1 MIS Server Platform Control Panel Buttons and LEDs **(continued)**

LED/Button	Description
NIC Activity LED	The green LED will be active whenever there is any network traffic occurring on the base board NIC ports.
Boot Drive Activity LED	The LED is lit whenever the boot drives are being accessed.
NMI Reset button	Not supported, non-functional.

Powering an MIS Platform

Turning the machine on

To power-on an MIS Server Platform, hold down the power button firmly until the fans spin up and the hardware lights come on. This will initialize the machine's BIOS and start internal diagnostics, including hardware status, RAID Volume information, missing hardware alerts (e.g., missing batteries), and missing configuration alerts. This process will be repeated for each expander in the machine, and finally (if conditions allow), the machine's operating system will load and welcome/login screen appear.

If the machine is powering on and there is a fault (e.g., the machine was in the middle of a rebuild when the power was cut, a degraded RAID set is in the system, battery is exhausted, etc.), these warnings will appear during start-up.

If the platform is an MIS dual-server and both servers are powered down, performing the above steps only powers on the server with which you are working. Though the fans and drives will be powered on, the second server will remain powered off until the second server is powered on, then all power will be turned on.

For a JBOD Unit, the power button on the front panel will turn on the power to that I/O module. If a second module is installed and powered off, it will remain off until it too, is powered on.

Turning the machine off

There are different ways to shut down an MIS machine. The preferred method is to go into the machine's operating system and select **Shut Down**. This will prompt the user to enter a password before allowing the shut-down process. Other ways to power off the machine include:

- Using the OS GUI power-off button at the console screen, if a keyboard/mouse/video monitor is connected.
- When logged in via an ssh session and executing a “shutdown” or “poweroff” command.
- When logged in to the BMC and using the power control page to power off the sever.
- Using the remote console screen GUI power-off button, if a KVM RMM4Lite session is established through the BMC.

If the platform is an MIS dual-server and **both** servers are powered up, performing the above steps **only** powers off the server with which you are working. The fans, drives and second server will remain powered on until the second server is powered off, then all power (but standby) will be turned off.

For a JBOD Unit, the power button on the front panel will turn off the power to that I/O module. If a second module is installed and powered on, it, the fans and the drives will remain on until it, too, is powered off.

One additional way to shut down the platform is to hold down the **Power** button on the front of the unit (Figure 2-1) until the machine powers off. However, this forces the machine off and does not allow the machine to go through normal shut-down procedures, requiring a recovery process at start-up.

If power is lost from outside the machine (power outage), the machine will recognize the loss of power and execute an emergency shutdown procedure. If there is a battery back-up unit installed, it will protect integrity of cache in the event of power or server failure.

MIS JBOD Control Panel

The control panel (Figure 2-1) for the MIS JBOD is exactly the same as the MIS Server Platform. However, some of the buttons on the JBOD do not have the same function as they do on the MIS Server. Since there is no boot drive module in a JBOD, the Boot Drive Activity LED, located next to the Network Activity LED, is present but inactive.

Important: When there are two I/O modules on a JBOD, the top control panel connects to the bottom I/O module on the back of the unit, and vice versa, the bottom control panel accesses the top I/O module.

Disk Drive Light Pipes

Figure 2-2 shows the disk drive Light Pipes.



Figure 2-2 Disk Drive Light Pipes

Table 2-2 describes the meaning of disk drive Light Pipes.

Table 2-2 Disk Drive LEDs

Bi-color LED	Blue LED	Drive Status
Off	Off	Drive is off and can be removed.
Green	Off	Drive is on.
Yellow	Off	Service required.
Off/Green/Yellow	On	Indicates drive location.

Power Supply LEDs

There are two LEDs located on the face plate of the power supply, one green on top for AC power, and one bi-color yellow/green below indicating DC power (Figure 2-3). Table 2-3 describes the function of the power supply LEDs




Figure 2-3 Power Supply LEDs

Table 2-3 Power Supply LEDs

LED Signaling	Operating Condition
AC Solid Green	AC Line within range
AC Off	AC Line UV condition
DC Blinking Yellow (1:1)	PSON High
DC Blinking Yellow/Green (1:2)	Hot-Standby Mode
DC Solid Yellow	V_1 or V_{SB} out of regulation, Over Temperature shutdown, Output over voltage shutdown (V_1 or V_{SB}), Output over current shutdown (V_1 or V_{SB}), Fan error
DC Blinking Yellow/Green (2:1)	Over temperature warning
DC Blinking Yellow/Green (1:1)	Minor fan regulation error (>5%, <15%)

Power Supply Monitoring

The fan base firmware will monitor all Power Supply Units (PSUs) installed in the system and generate SEL entries regardless of whether or not the alerts are enabled. This functionality cannot be turned off. To receive alerts, in addition to these instructions for entering new filters, correct configuration of the Intel BMC must also be performed. This can be done via the Intel Integrated BMC Web Console. Specific instructions for the BMC configuration is not included here because it is site specific. Command strings are IPMI commands.

When the fan base firmware monitors the PSU it sends alerts to the BMC. This alone does not send the email. You will need to set up the BMC Web Console to have those emails delivered, and set filters 19 and 20.

The S2600JF motherboard allows 20 filters to be set. Intel uses the first 13. The examples that set filters 19 and 20 are how we set up the BMC so that it realizes that the PSU alert that came in should also have an email sent.

Platform Event Filter 5 sets the alert for PSU failure detected and PSU predictive failure. A SEL entry will be recorded on both assertion and deassertion. An alert will only be sent when asserted. This filter is configured by Intel.

PEF filter 19 sets the alert for PSU input lost (AC/DC). A SEL entry will be recorded on both assertion and deassertion. An alert will only be sent when asserted. This filter is configured by SGI.

PEF filter 20 sets the alert for PSU presence detected. A SEL entry will be recorded on both assertion and deassertion. An alert will only be sent when deasserted. This filter is configured by SGI.

Requirements

- fanbase_release_01.04.a.01 or later
- Shackleford_EFI_BMC-1.17.4151_BIOS-01.08.0003_ME-02.01.07.112_FRUSDR-1.10.02 or later

Viewing Existing Filters

Example using ipmitool to read filter 19:

```
ipmitool [-I lanplus -H <BMC_IP_ADDRESS> -U <BMC_USERNAME> -P <BMC_PASSWORD>] raw 0x04 0x13 0x06 0x13 0x00
```

Example using ipmitool to read filter 20:

```
ipmitool [-I lanplus -H <BMC_IP_ADDRESS> -U <BMC_USERNAME> -P <BMC_PASSWORD>] raw 0x04 0x13 0x06 0x14 0x00
```

Example using ipmitool to read all filters:

```
ipmitool [-I lanplus -H <BMC_IP_ADDRESS> -U <BMC_USERNAME> -P <BMC_PASSWORD>] pef list
```

Setting Filters

Example using ipmitool to set filter 19:

```
ipmitool [-I lanplus -H <BMC_IP_ADDRESS> -U <BMC_USERNAME> -P
<BMC_PASSWORD>] raw 0x04 0x12 0x06 0x13 0xC0 0x01 0x01 0x00 0xFF
0xFF 0x08 0xFF 0x6F 0x08 0x00 0x00 0x00 0x00 0x00 0x00 0x00
0x00 0x00
```

Example using ipmitool to set filter 20:

```
ipmitool [-I lanplus -H <BMC_IP_ADDRESS> -U <BMC_USERNAME> -P
<BMC_PASSWORD>] raw 0x04 0x12 0x06 0x14 0xC0 0x01 0x01 0x00 0xFF
0xFF 0x08 0xFF 0xEF 0x01 0x00 0x00 0x00 0x00 0x00 0x00 0x00
0x00 0x00
```

Clearing Filters

Example using ipmitool to clear filter 19:

```
ipmitool [-I lanplus -H <BMC_IP_ADDRESS> -U <BMC_USERNAME> -P
<BMC_PASSWORD>] raw 0x04 0x12 0x06 0x13 0x00 0x00 0x00 0x00 0x00
0x00 0x00 0x00 0x00 0x00 0x00 0x00 0x00 0x00 0x00 0x00 0x00
0x00 0x00
```

Example using ipmitool to clear filter 20:

```
ipmitool [-I lanplus -H <BMC_IP_ADDRESS> -U <BMC_USERNAME> -P
<BMC_PASSWORD>] raw 0x04 0x12 0x06 0x14 0x00 0x00 0x00 0x00 0x00
0x00 0x00 0x00 0x00 0x00 0x00 0x00 0x00 0x00 0x00 0x00 0x00
0x00 0x00
```

Fan Base Interface

The fan baseboard is a circuit board at the heart of the chassis design. As such, it can be used to monitor hardware and issue commands to the hardware without having to use an operating system or even a server module (which is why it is useful for JBODs). Instead, commands are issued and executed by the fan base firmware. Connecting to the fan base can be done in several ways, either through the MIS-S9D proprietary network interface, through an MIS Server NIC card, or through a JBOD I/O card Ethernet port and pointing a browser to the Fan Base Service Page IP address: 192.168.0.3/service.xml for servers (Figure 2-4) and <DHCP address>/service.xml for JBODs.

MIS Server Chassis													
Cover is on, uptime 16868006(ms mod 2 ³²)				All FANS <input type="button" value="On"/>				<input type="button" value="Shutdown"/>					
R E A R	D8	D5	D2	StorBrick 5 Good	Boot Drive CM0-1	FAN 5 Good <input type="button" value="On"/> <input type="button" value="Off"/> RPM : 12273 : 11650	StorBrick 4 Good	D0	D3	D6	F R O N T		
	D7	D4	D1		Boot Drive CM0-0			D1	D4	D7			
	D6	D3	D0		Boot Drive CM1-1			D2	D5	D8			
	D8	D5	D2	StorBrick 6 Good	Boot Drive CM1-0	FAN 4 Good <input type="button" value="On"/> <input type="button" value="Off"/> RPM : 11977 : 10704	StorBrick 3 Good	D0	D3	D6			
	D7	D4	D1		D1			D4	D7				
	D6	D3	D0		D2			D5	D8				
	D8	D5	D2	StorBrick 7 Good		FAN 3 Good <input type="button" value="On"/> <input type="button" value="Off"/> RPM : 10836 : 12334	StorBrick 2 Good	D0	D3	D6			
	D7	D4	D1		D1			D4	D7				
	D6	D3	D0		D2			D5	D8				
	Compute Module 1 -				FAN 2 Good <input type="button" value="On"/> <input type="button" value="Off"/> RPM : 11753 : 11021				StorBrick 1 Good				
	Compute Module 0 CM0 P1 59.0°C Ok - CM0 P2 55.0°C Ok				FAN 1 Good <input type="button" value="On"/> <input type="button" value="Off"/> RPM : 11384 : 11014								
	Power Supply 3 Good				FAN 0 Good <input type="button" value="On"/> <input type="button" value="Off"/> RPM : 10987 : 10763								StorBrick 0 Good
Power Supply 2 Good								D0				D3	D6
Power Supply 1 o								D1				D4	D7
Power Supply 0 o								D2				D5	D8

Figure 2-4 Fan Base Service Page

From here you can view StorBrick status, power supply status, fan status, and set fans in a “**Safe to Service**” mode (i.e., off) for maintenance. For further instructions, see “Replacing a Fan Module” on page 136.

MIS-S9D Proprietary Network Interface

The MIS-S9D interface is to be used when accessing the fan base for servicing, or when zoning (see “Operating Zones 1.4.2 for CentOS, Linux, & Windows” on page 69). It is located at the front of the chassis at the upper right corner (Figure 2-5). The chassis must be slid out forward at least one inch in order to connect a network crossover cable. (See “Sliding the Chassis Forward/Backwards” on page 125.)



Figure 2-5 MIS-S9D Proprietary Network Interface

Ensure the MIS system is powered on. Use an Ethernet cable to connect a server/laptop. The network port connected to the server/laptop must be set to 192.168.0.xxx (10 will do). The static IP address of the fan base is set to 192.168.0.3, verify connectivity to the fan base with a ping command to 192.168.0.3 from the server/laptop. Verify it responds. If not it will be necessary to power cycle the MIS Server or JBOD.

BMC Integrated Web Console

The control panel and various other LEDs are used to monitor the overall status of the system and its components. Underlying the light-guided diagnostics provided through the various LEDs on the control panel, power supplies, motherboard, etc. are the BMC/IPMI interfaces. The MIS server supports the platform management features (environmental monitoring, power management, etc.) provided by the Intel BMCs and IPMI 2.0. Moreover, the BMCs have features beyond those of IPMI 2.0 (for instance, detection of chassis intrusion).

The BMC Integrated Web Console is a web-based program provided by Intel, and is used to give general system information such as system diagnostics, server health, environmental reporting, and event logs. Additionally, the BMC-IWC provides a remote virtual control panel for the MIS Server, allowing for remote locating and reboot.

For more information, see the **Help** files within the BMC, or view the platform management documentation for the Intel S2600JF motherboard online: [Intel® Remote Management Module 4 and Integrated BMC Web Console User Guide](#) (pdf).

This section gives you a description of BMC Web Console pages that **differ** from the Intel documentation, or otherwise require more detail.

To view the **Help** files within the BMC, Click the **Help** button in the upper right corner (Figure 2-6).




Figure 2-6 BMC Web Console – Logout, Refresh, and Help Buttons

System Debug Log

The System Debug Log page is not supported at this time. It allows administrators to collect system debug information. The files are compressed, encrypted, and password protected. The file is not meant to be viewable by the end user. De-encryption of these files is unavailable at this time.

Select either the “System Debug Log” or the “System & BMC Debug Log” and press the **Run** button. It may take some time for the debug information to be collected. Once the debug log dump is finished you can click the debug log filename to save the results as a `.zip` file on your client system. However, this file cannot be used at this time.

System Debug Log Type

The System Debug Log data is mainly used by the system manufacturer for analysis. Baseboard Management Controller (BMC) status, BMC configuration settings, BMC Sensor readings, Power supply data, System Event Log, sensor readings, SMBIOS tables, CPU machine check registers and PCI configuration space information. The System & BMC Debug Log contains regular System Debug Log plus the BMC debug log.

Last Log

Shows the time of the last data collection. Collection times older than three minutes will be marked as an “Old” debug log.

Encryption

The resulting zip file will be encrypted for privacy, and may only be extracted for analysis by an authorized representative of the system manufacturer.

Generate Log

Click the Generate Log button to collect recent Debug Log data. The resulting compressed archive will be downloaded to your system by clicking on the debug log link. You may also choose to download the data at a later time using the debug log link. Note that it is recommended that fresh data always be downloaded for analysis.

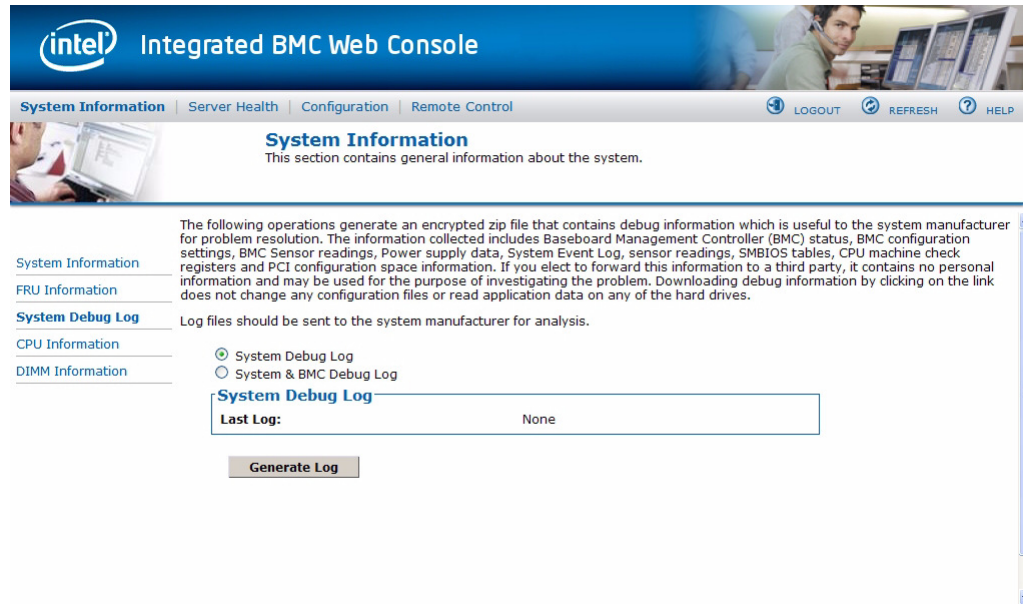


Figure 2-7 BMC Web Console – System Debug Log

Server Health

The Server Health tab shows you data related to the server's health, such as sensor readings, the event log, and power statistics as explained in the following sub sections. Click on the Server Health tab to select the various pages. By default, this tab opens the Sensor Readings page.

Sensor Readings

The Sensor Readings page displays system sensor information including status, health, and reading. By default the sensor readings are updated every 60 seconds but this can be changed by entering a value in the Set auto-refresh in seconds selection box and then pressing the **Set** button.

Sensor Selection drop-down box allows you to select the type of sensor readings to display in the list. The default is set to All Sensors, with other options: Temperature Sensors, Voltage Sensors, Fan Sensors, Physical Security, Processor, Power Units, Memory, Event Logging Disable, System Event, Button/Switch, Module/Board, Watchdog Sensor, Management Subsystem Health, Node Manager, and SMI.



Server Health
This section shows you data related to the server's health, such as sensor readings and the event log.

Sensor Readings
This page displays system sensor information, including readings and status. You can toggle viewing the thresholds for the sensors by pressing the Show Thresholds button below.

Refreshing readings every 60 seconds
Select a sensor type category:

Sensor Readings: 102 sensors

	Status	Health	Reading
All Sensors			
All Sensors			
Temperature Sensors	All deasserted	OK	0x0000
Voltage Sensors	All deasserted	OK	0x0000
Fan Sensors	All deasserted	OK	0x0000
Physical Security	All deasserted	OK	0x0000
Processor	All deasserted	OK	0x0000
Power Unit	reports the System Event Log (SEL) is full	OK	0x0010
Memory	All deasserted	OK	0x0000
Event Logging Disabled	All deasserted	OK	0x0000
System Event	All deasserted	OK	0x0000
Button / Switch	All deasserted	OK	0x0000
Module / Board	All deasserted	OK	0x0000
Watchdog 2	All deasserted	OK	0x0000
Management Subsystem Health	All deasserted	OK	0x0000
Node Manager	All deasserted	OK	0x0000
SMI	All deasserted	OK	0x0000
SSB Temp	Normal	OK	27 degrees C
SSB Temp	Normal	OK	46 degrees C
BB BMC Temp	Normal	OK	36 degrees C
P1 VR Temp	Normal	OK	31 degrees C
IB QDR Temp	Normal	OK	35 degrees C
Powerville Temp	Normal	OK	53 degrees C
Sys Fan 1A	All deasserted	OK	0x0000
Sys Fan 1B	All deasserted	OK	0x0000
Sys Fan 2A	All deasserted	OK	0x0000

Figure 2-8 BMC Web Console – Server Health

Click **Show Thresholds** to expand the list, showing low and high threshold assignments. Use scroll bar at bottom to move display left and right.

- CT: Critical threshold
- NC: Non-critical threshold

Click **Hide Thresholds** to return to original display, hiding the threshold values, showing only the name, status and reading for selected sensors. Click **Refresh** to refresh the selected sensor readings.

Event Log

The Event Log is a table of the events from the system's event log. The BMC provides a centralized, non-volatile repository for critical, warning, and informational system events called the System Event Log or SEL. By having the BMC manage the SEL and logging functions, it helps ensure that 'post-mortem' logging information is available should a failure occur that disables the system processor(s).

The BMC allows access to SEL from in-band and out-of-band mechanisms. There are various tools and utilities that can be used to access the SEL, such as the Intel® SELView utility and multiple open-sourced IPMI tools.

Items that are monitored using the SEL, such as power supply status, can generate alerts through the BMC Integrated Web Console.

Event Log

Below is a table of the events from the system's event log. You can choose a category from the pull-down box to filter the events, and also sort them by clicking on a column header.

Select an event log category:

All Events

All Events

Sensor-Specific Events

BIOS Generated Events

System Management Software Events

Event Log: 3639 event entries

Event ID	Time	Sensor Name	Sensor Type	Description
3638	05/11/2012 20:39:07	SSB Temp	Temperature	reports the sensor is in a low, critical, and going lower state - Asserted
3637	05/11/2012 20:39:07	SSB Temp	Temperature	reports the sensor is in a low, critical, and going lower state - Asserted
3636	05/11/2012 20:39:01	SSB Temp	Temperature	reports the sensor is in a high, but non-critical, and going higher state - Deasserted
3635	05/11/2012 20:39:00	SSB Temp	Temperature	reports the sensor is in a high, but non-critical, and going higher state - Asserted
3634	05/11/2012 20:38:41	SSB Temp	Temperature	reports the sensor is in a low, critical, and going lower state - Deasserted
3633	05/11/2012 20:38:41	SSB Temp	Temperature	reports the sensor is in a low, critical, and going lower state - Deasserted
3632	05/11/2012 20:38:40	SSB Temp	Temperature	reports the sensor is in a low, critical, and going lower state - Asserted
3631	05/11/2012 20:38:40	SSB Temp	Temperature	reports the sensor is in a low, but non-critical, and going lower state - Asserted
3630	05/11/2012 20:38:32	SSB Temp	Temperature	reports the sensor is in a low, but non-critical, and going lower state - Deasserted
3629	05/11/2012 20:38:32	SSB Temp	Temperature	reports the sensor is in a low, critical, and going lower state - Deasserted
3628	05/11/2012 20:38:31	SSB Temp	Temperature	reports the sensor is in a low, critical, and going lower state - Asserted
3627	05/11/2012 20:38:31	SSB Temp	Temperature	reports the sensor is in a low, but non-critical, and going lower state - Asserted

Figure 2-9 BMC Web Console – Event Log

You can choose a category from the pull-down box to filter the events, and also sort them by clicking on a column header. The filters available are All Events, Sensor-Specific Event, BIOS Generated events, and System Management Software Events. Use this page to view and save the Event log. **Event Log Category** selects the type of events to display in the list. **Event Log List** is a list of the events with their ID, time stamp, sensor name, sensor type, and description. Click **Clear Event Log** to clear the event logs. Click on **Save Event Log** to download the event logs to local system.

The system event log hexadecimal codes can be translated using Intel's [System Event Log Troubleshooting Guide for EPSD Platforms Based on Intel Xenon Processors E5 4600/2600/2400/1600/1400 Product Families](#) (pdf).

Power Statistics

This section is not used by MIS platforms.

Power Usage Summary

System power statistics in watts	
Minimum:	0W
Current:	0W
Maximum:	0W
Average:	0W

Figure 2-10 BMC Web Console – Power Statistics

Configuration Tab

The Configuration tab of the BMC Web Console is used to configure various settings, such as alerts, users, or network. It contains the following menu options in the left navigation pane: IPv4 Network, IPv6 Network, Users, Login, LDAP, VLAN, SSL, Remote Session, Mouse Mode, Keyboard Macros, Alerts, Alert Email, Node Manager.

Alerts

The Alerts page allows you to configure which system events generate Alerts and the external network destinations to which they should be sent.

When one of the selected system events occurs, an alert is generated and sent to the configured destination(s). Each LAN channel can have up to two destinations.

Globally Enable Platform Event Filtering

Global control for enabling or disabling platform event filtering. When filtering is globally disabled through this setting, alerts will not be sent. This can be used to prevent sending alerts until you have fully specified your desired alerting policies.

Select Events

Select one or more system events that will trigger an Alert. Clearing all events disables Alerts. These events correspond to the IPMI preconfigured Platform Event Filters.

LAN Channel

Select which LAN Channel to configure destinations for. Each LAN Channel has its own set of up to two destinations. Alert destinations can be one of two types:

- SMTP Trap
- Email (requires Alert Email to be configured)

The **Check All** button selects all events to generate Alerts. The **Clear All** button unchecks all events so no Alerts will be generated. Click the **Save** button to save any changes made.

Send Test Alert

To test whether an alert will reach its destination, set the LAN Channel field to the desired channel and configure at least one destination. Then click **Send Test Alerts** to send a simple test alert to the destination(s) for that Channel.

Configuration
Use these pages to configure various settings, such as alerts, users, or network.

Alerts

Configure which system events generate Alerts and the external network destinations they should be sent to.

IPv4 Network
IPv6 Network
Users
Login
LDAP
VLAN
SSL
Remote Session
Mouse Mode
Keyboard Macros
Alerts
Alert Email
Node Manager

Globally Enable Platform Event Filtering: Enabled Disabled

Select the events that will trigger alerts:

<input checked="" type="checkbox"/> Temperature Sensor Out of Range	<input checked="" type="checkbox"/> Watchdog Timer
<input checked="" type="checkbox"/> System Restart	<input checked="" type="checkbox"/> Voltage Sensor Out of Range
<input checked="" type="checkbox"/> Fan Failure	<input checked="" type="checkbox"/> Chassis Intrusion
<input checked="" type="checkbox"/> Power Supply Failure	<input checked="" type="checkbox"/> Memory Error
<input checked="" type="checkbox"/> BIOS: Post Error Code	<input checked="" type="checkbox"/> FRB Failure
<input checked="" type="checkbox"/> Node Manager Exception	<input checked="" type="checkbox"/> Hard Drive Failure

LAN Channel to Configure: Baseboard Mgmt

Alert Destination #1:

SNMP Send SNMP Alerts to IP: 102.38.137.30

Email Send Email to: brewer@sgi.com

Alert Destination #2:

SNMP Send SNMP Alerts to IP: 0.0.0.0

Email Send Email to:

Figure 2-11 BMC Web Console – Alerts

Alert Email

Alert Email Settings allows you to configure how Alerts are sent by email to an external SMTP Mailserver. Each LAN Channel has a separate configuration, selected through the drop-down menu. The **SMTP Server IP** is the IP address of the remote SMTP Mailserver that Alert email should be sent to. The **Sender Address** is the string to be put in the `From:` field of outgoing Alert emails. **Local Hostname** is a name for the local machine that is generating the alert, and this name

is included in the outgoing Alert email. The **Local Hostname** is a string of maximum 18 alpha-numeric characters. Spaces and special characters are not allowed.

Once these settings are saved, you can go to the previous Alerts screen and click the button **Send Test Alert**. If the SMTP Server IP address is correct, an email containing hexadecimal code should be sent to the emails configured on the Alerts page. To understand what these hex codes mean, see Intel’s [System Event Log Troubleshooting Guide for EDSP Platforms Based on Intel Xeon Processor E5 4600/2600/2400/1600/1400 Product Families](http://download.intel.com/support/motherboards/server/sb/g90620001_s1400_s4600_systemeventlog_troubleshooti.pdf) (http://download.intel.com/support/motherboards/server/sb/g90620001_s1400_s4600_systemeventlog_troubleshooti.pdf).

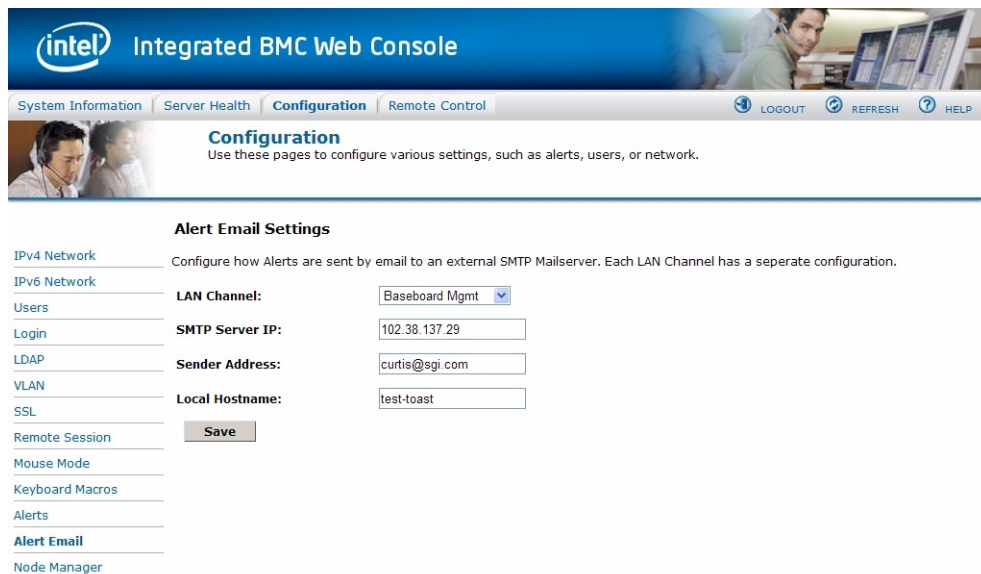


Figure 2-12 BMC Web Console – Alert Email Settings

In this example, the mail server’s IP address is 102.38.137.29. When the **Send Test Alert** button is pushed, it creates an alert email from the Sender Address, with the heading “**Alert from <Local Hostname>**” (in this example, Figure 2-13, the local hostname is test-toast).

Alert from <test-toast>

AlertEmail [curtis@sgi.com]

Extra line breaks in this message were removed.

To: Pamela Curtis

Event that generated this alert:

```

RID:FFFF TS:02/08/2013 10:41:31 SN:unnamed sensor 0x00 (no SDR)
ST:Unknown ED:Unknown ET:Asserted EC:Non-Recoverable RID:FFFF
RT:02 TS:5114D65B GID:0020 ER:04 ST:00 S#:00 ET:00 ED:00 00 00
EX:3F 00 00 00 00 00 00 00

```

Figure 2-13 BMC Web Console – Alert Email Result

The resulting hex code in the email can be translated using the code tables given in Intel's *System Event Log Troubleshooting Guide for EDSP Platforms Based on Intel Xeon Processor E5 4600/2600/2400/1600/1400 Product Families*. In our example, the hex code includes the record identification (**RID**), a time stamp for when the event was generated in MM/DD/YYYY HH:MM:SS format (**TS**), the number of the sensor that generated the event (**SN**), the sensor type code (**ST**), event directory (**ED**), event type – Asserted or Deasserted (**ET**), event code (**EC**), record identification used for SEL record access (**RID**), record type in Hex code (**RT**), time stamp for when the event was logged (**TS**), generator identification in hex including RqSA and LUN if the event was generated from the IPMB software and software identification if the event was generated from system software (**GID**):

- 0001 – BIOS POST for POST errors, RAS Config/State, Timestamp Synch, OS Boot events
- 0033 – BIOS SMI Handler
- 0020 – BMC Firmware
- 002C – ME Firmware
- 0041 – Server Management Software
- 00C0 – HSC Firmware HSBP A
- 00C2 – HSC Firmware HSBP B

In the example above (Figure 2-13), you'll notice that the **GID** is 0020, indicating correctly that the alert was sent from the BMC Firmware.

Continuing with the hex code contained within the email, ER is the IPMI version in use: 04 = IPMI version 2.0, 03 = IPMI version 1. Sensor Type (**ST**) and Sensor Number (**S#**) come next, followed by the Event Trigger (**ET**), Event Data (**ED**) and finally Event X (**EX**).

Server Power Control

This page shows the power status of the server and the following power control operations can be performed:

Table 2-4 Server Power Control Actions

Option	Details
Reset Server	Selecting this option will hard reset the host without powering off.
Force-enter BIOS Setup	Check this option to enter into the BIOS setup after resetting the server.
Power Off Server	Selecting this option will immediately power off the host.
Graceful Shutdown	Selecting this option will soft power off the host.
Power On Server	Selecting this option will power on the host.
Power Cycle Server	Selecting this option will immediately power off the host, then power it back on after one second.




Figure 2-14 BMC Web Console – Power Control and Status

All power control actions are done through the BMC and are **immediate** actions. It is suggested to gracefully shut down the operating system via the KVM interface or other interface before initiating power actions through the Virtual Front Panel.

The Virtual Front Panel is a graphic representation of the front panel, providing remote functionality virtually.

Table 2-5 System Information Details

Button	Details
Power Button	Power button is used to Power ON or Power OFF
Reset Button	Reset Button is used to reset the server while system is ON
Chassis ID Button	When Chassis ID button is pressed then the chassis ID LED, on the front and rear of the unit are lit (solid blue). If the button is pressed again the chassis ID LED turns off.

Table 2-5 System Information Details

Button	Details
NMI Button	At present, NMI button is disabled.
Status LED	Status LED will reflect the system status and will automatically sync with BMC every 60 seconds. If any abnormality occurs in system, then status LED will change accordingly. Thermal fault means <i>fault</i> occurred in one of Thermal sensors present in BMC. Fan fault means <i>fault</i> occurred in one of the system fans. System fault means a <i>fault</i> occurred because of system errors. Power fault means <i>fault</i> occurred in one of Power sensors. Here, <i>fault</i> means a sensor value crossed upper-non-critical, or upper-critical value, or lower-non-critical value, or lower-critical value.
Power LED	Power LED shows system power status. If LED is green then System is ON. If LED is grey then System is OFF.
Chassis ID LED	The Chassis ID LED will be lit blue when the Chassis ID LED button is pushed. This is the same as the Locator LED on the physical control panel (Table 2-1). When the Locator button is pushed again, the LED will be lit. There is a corresponding LED on the back of the server that will be lit blue as well. This function can be done physically through the Control Panel (Figure 2-1).

Intelligent Management Platform Interface (IPMI 2.0)

IPMI 2.0 is a standardized computer system interface used by system administrators for out-of-band management of computer systems and monitoring of their operation. The development of this interface specification was led by Intel Corporation. IPMI is a message-based, hardware-level interface specification: an IPMI subsystem operates independently of the operating system (OS) (i.e., out-of-band). This allows administrators to manage a system remotely in the absence of an operating system, before an OS has booted (allowing e.g. BIOS settings to be

remotely monitored or changed), when the system is powered down, or after OS or system failure, which is the key characteristic of IPMI compared with in-band system management (such as, by remote login using SSH). Multiple disparate servers can all be managed together using IPMI because of its standardized interface and protocol.

System administrators can use IPMI messaging to:

- monitor platform status (e.g., system temperatures, voltages, fans, power supplies, and chassis intrusion),
- query inventory information,
- review hardware logs of out-of-range conditions,
- perform recovery procedures such as issuing requests from a remote console through the same connections (e.g., system power-down and rebooting, or configuring watchdog timers)
- and define an alerting mechanism for the system to send a simple network management protocol (SNMP) platform event traps (PETs).

Programs such as IPMITool and IPMIUtil are industry freeware standards that can both be used to configure and issue IPMI commands. More information on both tools can be found at <http://ipmiutil.sourceforge.net/docs/UserGuide> and <http://ipmitool.sourceforge.net/manpage.html>. IPMITool comes with Linux and IPMIUtil is available for download through <http://support.sgi.com> (requires customer account user name and password).

More Information on IPMI can be found through <http://www.intel.com/content/www/us/en/servers/ipmi/what-is-ipmi.html> and http://en.wikipedia.org/wiki/Intelligent_Platform_Management_Interface.

System Software

Overview

Important: Do **not** install any software from 3rd-party sites. All software required to monitor and operate MIS platforms is available through <http://support.sgi.com> or by contacting your SGI support representative.

This chapter contains information on each of the software sets necessary for zoning MIS Servers and JBODs:

- “Downloading & Installing Software,” on page 62,
- “Operating Zones 1.4.2 for CentOS, Linux, & Windows,” on page 69,
- “Creating the Drive Groups in MegaRAID (CentOS, Linux & Windows),” on page 83,
- “Formatting the Drives Using YaST2 in Linux,” on page 88,
- “Windows Server Manager,” on page 93,
- “CLI Zoning Tool, version 1.4,” on page 99,
- “Disk RAID Support,” on page 127.

Zoning is required when multiple SAS connections are operational, in order to stop drives from being affected by other non-owner SAS controllers (HBAs). Zoning allows the various SAS connections to be accessible only to the drives that they own. Essentially, zoning allows an administrator to control which HBAs can access which drives. When **open zoning** is enabled, all the SAS connections can access all of the drives. For dual-slotted SAS drives, **both** slots will be exposed, so the drives will show up **twice**. This situation will cause conflict between the HBAs.

Caution: Do not assign two HBAs to the same drive sets as this will cause data collisions.

Zoning can be either hard or soft. In hard zoning, each device is assigned to a particular zone, and this assignment does not change. In soft zoning, device assignments can be changed by the network administrator to accommodate variations in the demands on different servers in the network. New drive sets are “zoned in” to existing drive sets using the zoning software described later in this chapter.

Phy-based zoning allows you to split the drives between two I/O *cards* (though it also may be used when only one I/O card is present).

Note: Zones 1.4.2 for CentOS, Linux, & Windows *no longer requires* the presence of an LSI MegaRAID *card*. The MegaRAID Storage Manager program, however, is still required.

There are two main tools for zoning: the SGI Zones application and the SGI CLI Zoning Tool. Both zoning tools require the presence of other programs in order to operate. The Zones application offers a GUI interface, and can now zone JBODs (version 1.3 and later). The CLI Zoning Tool can zone in the absence of an operating system, but is a command-line only application, now with expanded features (version 1.4 and later).

Important: The Zones program is installed and run on the server you wish to zone. JBODs are zoned through their SAS connection to an MIS server running the Zones or CLI programs, or directly through the hardware using the CLI Zoning tool installed on a laptop, and an Ethernet crossover cable connected to the MIS-S9D proprietary network interface (Figure 3-44).

Even though MIS Servers and JBODs are shipped with zoning from SGI manufacturing, it is assumed that this zoning will be eliminated and replaced with your desired zoning. Any RAID sets will need to be cleared and set to unconfigured good before zoning (see “Creating the Drive Groups in MegaRAID (CentOS, Linux & Windows),” on page 83 for more information).

Section Guide

Inside the first section you will find instructions on downloading and installing the software needed. The next two sections give instructions on how to use the software to zone MIS Platforms. The final section contains information on the different RAID configurations available, advantages and disadvantages of each configuration, as well as best practices.

- “Downloading & Installing Software” and “Required Downloads,” on page 63,

- › “Installing MegaRAID Storage Manager for CentOS,” on page 63
- › “Installing Zones 1.4.2 for CentOS,” on page 65
- › “Installing MegaRAID Storage Manager for Linux,” on page 65,
- › “Installing Zones 1.4.2 for Linux,” on page 67,
- › “Installing MegaRAID Storage Manager for Windows,” on page 67,
- › “Installing Python for Windows,” on page 68,
- › “Installing Zones 1.4.2 for Windows,” on page 68,

The second section presents the new “Operating Zones 1.4.2 for CentOS, Linux, & Windows,” on page 69, including step-by-step instructions for “Zoning MIS Server Platforms,” “144 Drives Setup,” “Zoning MIS JBOD Platforms,” “Phy-based Zoning,” and “Formatting the Drives,” on page 94.

- “Operating Zones 1.4.2 for CentOS, Linux, & Windows,” on page 69,
 - › “Zoning MIS Server Platforms,” on page 69,
 - › “Opening a Session,” on page 71,
 - › “144 Drives Setup,” on page 73,
 - › “Saving Zoning Configuration,” on page 75,
 - › “Zoning MIS JBOD Platforms,” on page 76,
 - › “Phy-based Zoning,” on page 79,
 - › “Loading a CSV File to Zones,” on page 81,

Sections three, four, and five cover “Creating the Drive Groups in MegaRAID (CentOS, Linux & Windows),” on page 83, “Formatting the Drives Using YaST2 in Linux,” on page 88, and “Windows Server Manager,” on page 93 to finish up the necessary steps for configuring MIS platforms.

- “Creating the Drive Groups in MegaRAID (CentOS, Linux & Windows),” on page 83
- “Formatting the Drives Using Command Line in CentOS,” on page 86
 - › “Verify Drives Seen,” on page 86
 - › “Adding a New Drive,” on page 86
 - › “Formatting Drives,” on page 87
- “Formatting the Drives Using YaST2 in Linux,” on page 88

- › “Verify Drives Seen,” on page 90,
- › “Partitioning Drives,” on page 90,
- “Windows Server Manager,” on page 93
 - › “Windows Server Manager,” on page 93,
 - › “Formatting the Drives,” on page 94,

The next section details the “CLI Zoning Tool, version 1.4,” the command-line-only tool that is run from a host machine connected to the MIS Server or JBOD through the MIS-S9D proprietary network interface. The CLI Zoning Tool requires the presence of Python version 2.6 or 2.7 on the host machine in order to run. Python is standard on most Linux-based machines, and will need to be installed for Windows-based machines for the CLI Zoning tool to run. (See “Installing Python for Windows,” on page 68). “Preparing to Zone using the CLI Zoning Tool,” on page 105 gives instruction on what is necessary before “Editing the ShackCLI.ini file for CentOS & Linux,” on page 105 and “Editing the .csv File for the CLI Zoning Tool,” on page 111.

- “CLI Zoning Tool, version 1.4,” on page 99,
 - › “Preparing to Zone using the CLI Zoning Tool,” on page 105
 - “Editing the ShackCLI.ini file for CentOS & Linux,” on page 105,
 - “Editing the ShackCLI.ini file for Windows,” on page 107,
 - › “CLI Zoning Tool Main Menu,” on page 108
 - › “Editing the .csv File for the CLI Zoning Tool,” on page 111,

Finally, “RAID Configuration Notes,” on page 128 explains the special considerations in creating RAID arrays for use on StorBricks, that is, there are RAID configurations, namely 6+2 and 7+2, that ensure there is no single point of failure on the StorBricks within the MIS system. This section details how to manage those concerns.

- “Disk RAID Support,” on page 127,
 - › “RAID Configuration Notes,” on page 128
 - › “Spare Drives,” on page 130.

Downloading & Installing Software

All software downloads can be found by going to <http://support.sgi.com>, logging in with your SGI account credentials, and navigating to the software download section. This should be your primary

source for downloading any software used by the MIS platforms. If difficulties arise, contact your SGI Support Technician for assistance.

Zoning the MIS Server and JBOD platforms requires certain software. Which software you need to download and install to zone your system may depend on the operating system on the MIS Server. MIS Servers and JBODs can also be zoned in the absence of an operating system using the CLI (Command Line Interface) Zoning Tool. However, the CLI zoning tool requires a cable connection through the MIS-S9D proprietary network interface (Figure 3-44) in order to zone. Otherwise, the operating system installed on the boot drives is what determines which software to download and use to zone.

Commands to zone are executed from the boot drive module in the MIS Server platform to zone the StorBrick module also housed in the MIS Server. From here, any JBOD that is connected to the MIS Server can also be zoned, and included in the zoning of the server. Download the software for the operating system used by the MIS Server.

Required Downloads

No MIS Operating System – SGI CLI Zoning Tool 1.4 on host machine² with Python 2.6 or 2.7 installed

MIS Linux and Linux-based systems (CentOS, Red Hat) – LSI MegaRAID Storage Manager for Linux, SGI Zones 1.4.2 for Linux & Windows

MIS Windows systems – LSI MegaRAID Storage Manager for Windows, Python 2.7 for Windows, SGI Zones 1.4.2 for Linux & Windows

Installing MegaRAID Storage Manager for CentOS

Note: MegaRAID Storage Manager is not necessary for zoning using the CLI Zoning Tool.

The MegaRAID Storage Manager is used to prepare the drives for zoning, prior to using the Zones tool, and for creating the drive groups after zoning has been performed using the Zones tool.

² If the host laptop/server running the CLI Zoning Tool 1.4 is a Windows machine, Python for Windows is also required.

Currently, there is no specific version of MegaRAID Storage Manager for CentOS, but since CentOS is Linux-based, MegaRAID Storage Manager for Linux is used instead.

1. Go to <http://support.sgi.com>. Log in with your credentials.
2. Search for MegaRAID Storage Manager for Linux, or contact SGI Customer Support for the latest supported version.
3. Select the latest version for Linux and save the `.tar` file in a directory easily found from a command line interface, such as `/home/<user_name>`.
4. Open a shell prompt by selecting **Applications** from the desktop main menu, then selecting **System Tools** and **Terminal**, or right-click on the desktop and choose **Open Terminal** from the menu presented.
5. Change to the directory where the `.tar` file was saved (i.e., `# cd /home/<user_name>/`).
6. Untar the `.tar` file (e.g., `# tar -zxvf 13.04.03.01_Linux(x86)_MSM.tar`).
7. You should see an output of files and a new directory called `disk`. Change to the `disk` directory (i.e., `cd disk`).
8. Run the installation by typing `./install`.
9. Press `Y` to accept the License Agreement. A menu with the following options comes up:
Press 0 to exit from the installation
Choose[1-5]:
 - (1) - Complete
This option will install all the program features
 - (2) - client
This option will only install components required to remotely view and configure servers
 - (3) - StandAlone
This option will only install components required for local server management
 - (4) - Local
This option will only install components required for local configuration
 - (5) - Server
This option will only install components required for remote server management
10. Select option 1 - Complete.
11. When the installation is finished, close the command line interface, and access MegaRAID Storage Manager from the **Applications** drop-down menu on the desktop.

Installing Zones 1.4.2 for CentOS

Zones 1.4.2 for Linux (proprietary SGI software) is compatible with CentOS 6.3 and later, and is used to zone drives on the MIS Server platform (and eventually MIS JBOD units as well). To install the software complete the following instructions.

1. Go to <http://support.sgi.com>, log in with your credentials, and download Zones 1.4.2 for Linux (or contact SGI Customer Support for the latest supported version) and save in the `/opt` directory (if the `/opt` directory doesn't exist, make it).
2. Open a shell prompt by selecting **Applications** from the desktop main menu, then selecting **System Tools** and **Terminal**, or right-click on the desktop and choose **Open Terminal** from the menu presented.
3. Change to the `/opt` directory and unzip `Zones_1.4.2.zip` (e.g., `# unzip Zones_1.4.2.zip`). A `Zones` folder will appear, thus creating the `/opt/Zones` directory.

Warning: *Zones must be installed in the `/opt/Zones/` directory or it will not work.*

4. Change directory into the new `Zones` folder (e.g., `# cd /opt/Zones`)
5. Look for the folders `JBOD` and `MIS`. During copying, these folders names may change to lower case versions `jbod` and `mis`. If so, rename them back to all-capitalized versions.

Starting Zones 1.4.2 for CentOS

Run the Zones program by opening a shell prompt by selecting **Applications** from the desktop main menu, then selecting **System Tools** and **Terminal**, or right-click on the desktop and choose **Open Terminal** from the menu presented. Change to the `/opt/Zones` directory by typing `cd /opt/Zones`. Run Zones by typing `# python Start.py` from the command line and pressing **Enter**.

Installing MegaRAID Storage Manager for Linux

Note: MegaRAID Storage Manager is not necessary for zoning using the CLI Zoning Tool.

The MegaRAID Storage Manager is used to prepare the drives for zoning, prior to using the Zones tool, and for creating the drive groups after zoning has been performed using the Zones tool.

1. Go to <http://support.sgi.com>, log in with your credentials, and search for MegaRAID Storage Manager for Linux, or contact SGI Customer Support for the latest supported version.
2. Select the latest version for Linux and save the `.tar` file to a directory that is easy to navigate to using a command line interface (e.g., `/home/<user_name>`).
3. Change directory to where you have saved your `.tar` file (e.g., `# cd /home/<user_name>`)
4. Untar the `.tar` file, `# tar -zxvf <filename.tar>` (e.g., `-zxvf 13.04.03.01_Linux(x86)_MSM.tar`)
5. You should see an output of files and a new directory called `disk`. Change to the `disk` directory (i.e., `cd disk`)
6. Run the installation by typing `./install`.
7. Press `Y` to accept the License Agreement. A menu with the following options comes up:
Press `0` to exit from the installation
Choose[1-5]:
(1) - Complete
 This option will install all the program features
(2) - client
 This option will only install components required to remotely view and configure servers
(3) - StandAlone
 This option will only install components required for local server management
(4) - Local
 This option will only install components required for local configuration
(5) - Server
 This option will only install components required for remote server management
8. Select option 1 - Complete.
9. When the installation is finished, issue `./startupui.sh` command to start MegaRAID Storage Manager GUI (i.e., `# ./startupui.sh`).

MegaRAID Storage Manager GUI for Linux will appear.

Installing Zones 1.4.2 for Linux

Zones 1.4.2 for Linux is proprietary SGI software, used to zone drives on the MIS Server platform (and eventually MIS JBOD units as well). To install the software complete the following instructions.

1. Go to <http://support.sgi.com>, log in with your credentials, and download Zones 1.4.2 for Linux, or contact SGI Customer Support for the latest version.
2. Unzip `Zones_1.4.2.zip` from the command line (e.g., `# unzip Zones_1.4.2.zip`). A `Zones` folder will appear.
3. Copy the `Zones` folder into the `/opt` directory, thus creating the `/opt/Zones` directory.

Warning: *Zones must be installed in the `/opt/Zones/` directory or it will not work.*

4. Change directory into the new `Zones` folder (e.g., `# cd /opt/Zones`)
5. Look for the folders `JBOD` and `MIS`. During copying, these folders names may change to lower case versions `jbod` and `mis`. If so, rename them back to all-capitalized versions.

Starting Zones 1.4.2 for Linux

Run the `Zones` program from the `/opt/Zones` directory by typing `# python Start.py` from the command line and pressing enter.

Installing MegaRAID Storage Manager for Windows

Note: MegaRAID Storage Manager is not necessary for zoning using the CLI Zoning Tool.

The MegaRAID Storage Manager is used to prepare the drives for zoning, and to create the drive groups after zoning.

1. Go to <http://support.sgi.com>, log in with your credentials, and search for MegaRAID Storage Manager for Windows, or contact SGI Customer Support for the latest supported version.
2. Click **Accept** on the license agreement page and install.

3. Once installed, select MegaRAID Storage Manager from the list of Programs available, and click to run.

MegaRAID Storage Manager GUI for Windows will appear.

Installing Python for Windows

Both Zones 1.4.2 for Windows and the CLI Zoning Tool require Python be installed on the machine that will perform the zoning. Zones 1.4.2 for Windows uses Python 2.7. The CLI Zoning Tool will work on Python 2.6 or 2.7.

1. Go to <http://support.sgi.com>, log in with your credentials, and select the version of Python required, or contact SGI Customer Support for the latest approved version.
2. Download and start the installation.
3. The installation will ask which directory in which to install python (the default is `c:\Python##\` where ## is the version number).
4. On the **Customize Python** pane, click **Next**.
5. On the **Complete Python Installation** pane, click **Finish**.

Note: It may be necessary to set the variable path for Python in Windows. If errors occur after installation saying “cannot find path,” see Chapter 5, “Troubleshooting,” for more information.

Installing Zones 1.4.2 for Windows

Zones 1.4.2 for Windows is proprietary SGI software, used to zone drives on the MIS Server and JBOD Platforms.

1. Go to <http://support.sgi.com>, log in with your credentials, and download Zones 1.4.2 for Windows, or contact SGI Customer Support for the latest released version.
1. Extract the `Zones.zip` files to the destination folder `c:\Program Files (x86)\`. This will create the directory `c:\Program Files (x86)\Zones\`.
2. In the directory `c:\Program Files (x86)\Zones\`, look for the extracted file `pygtk-all-in-one-2.22.6.win32-py2.7.msi`. Run by double-clicking on it.

Starting Zones 1.4.2 for Windows

To run Zones, go to `c:\Program Files(x86)\Zones\`, and double-click on the icon `Start.py`.

Operating Zones 1.4.2 for CentOS, Linux, & Windows

This updated release of Zones features the familiar graphic user interface, and is streamlined to make zoning faster and easier to execute. Zones 1.4.2 for CentOS & Linux and Zones 1.4.2 for Windows both function the same, with only a slight difference in appearance. The following instructions work for Linux (CentOS) and Windows versions, though only images from the Linux version appear here. For instructions on how to run Zones 1.4.2, see “Starting Zones 1.4.2 for CentOS,” on page 65, “Starting Zones 1.4.2 for Linux,” on page 67, or “Starting Zones 1.4.2 for Windows,” on page 69.

Zoning MIS Server Platforms

Upon opening, Zones will display the current zoning configuration, if present and if the previous zoning was done by the Zones tool.

Note: Zones will not display the previous zoning configuration if zoning was previously done using the CLI Zoning Tool.

Drives that are currently installed in the system will be highlighted and a check mark will be in their corresponding box. In the following example (Figure 3-1), **Drive 0** is zoned for every StorBrick (**Brick 0** through **Brick 7**) on **Adapter 0**.

Zones interface

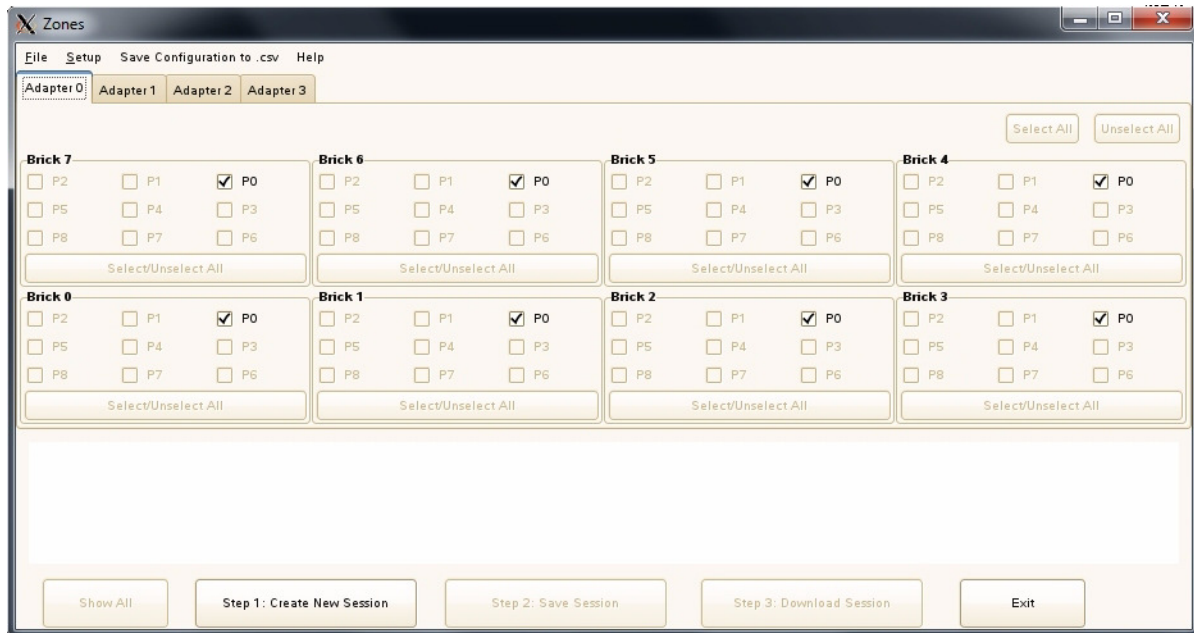


Figure 3-1 Zones 1.4.2 – Current Zoning Configuration

To begin zoning, click the button on the bottom of the home screen: **Step 1: Create a New Session** (Figure 3-2).




Figure 3-2 Zones 1.4.2 – Bottom Row of Buttons

This will give you the option to **Open a CSV Session**, **Open a New Session**, or **Cancel** the create new session operation (Figure 3-3). Select: **Open New Session**. (How to use CSV Sessions is covered in “Loading a CSV File to Zones,” on page 81.)

Opening a Session




Figure 3-3 Zones 1.4.2 – Step 1: Opening a Session

This will bring up a **New Session** prompt, asking you to create an alias name for the session (Figure 3-4). This alias name will be used to create a time-stamped folder with that alias, necessary to store the files used to zone the machine: `YYMMDD_HHMMSS_alias`. After entering a session alias, click **OK**.

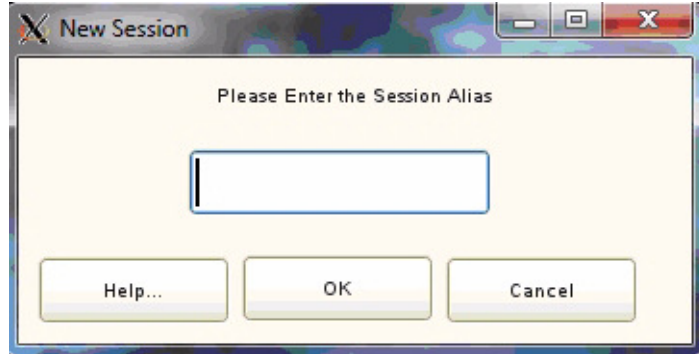


Figure 3-4 Zones 1.4.2 – New Session Alias

Note: Aliases have a 64 character limit, and may not contain spaces or non-alpha-numeric characters. If used, a warning message will appear (Figure 3-5).

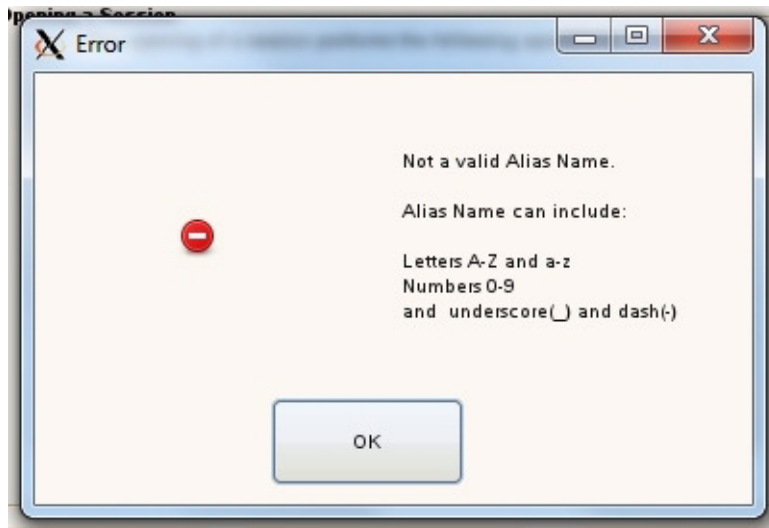


Figure 3-5 Zones 1.4.2 – Warning Box, Alias Rules

The home screen will now display the number of adapters present in the system. In this example (Figure 3-6), all four adapters are present, which can be seen by the four new tabs at the top, labeled **Adapter 0**, **Adapter 1**, **Adapter 2**, and **Adapter 3**. This would be the result of having a

single server with four adapters, or a dual server machine with two adapters per server. If the server was a dual server with only one adapter per server, only two adapter tabs would show.




Figure 3-6 Zones 1.4.2 – Top Menu Bar and Adapter Tabs

In a dual server configuration, Zones is unable to determine which adapter belongs to which server. The user can determine which adapters belong to which server by powering on only one server and then running the zoning tool. The adapter tabs that show up will belong to the server that is currently powered on.

1. Power on the first server.
2. Open Zones.
3. Open a session. (There's no need to save this session.)
4. Make note of which adapter in Zones belongs to that server.
5. Exit Zones.
6. Power on the second server
7. Start Zones again.
8. The user may then continue to zone.

Next, enable all the drives by clicking the **Show All** button at the bottom of the home screen (Figure 3-2). This will enable all the slots, and you can now select which drives you would like associated with which adapter.

144 Drives Setup

If the machine is populated with 9mm profile SSD drives, this allows two drives per slot in the machine. It splits the slot, creating primary and secondary drives. By default, only primary drives are seen. To zone the drives on the secondary slots, you will first need to go to the **Setup** menu option (Figure 3-7) on the top menu bar and select **144 Drives**.




Figure 3-7 Zones 1.4.2 – Setup Menu Options

Before being allowed to continue, the following warning will appear (Figure 3-8). If you do *not* have a 144-drive system, you can select the **Cancel** button. If you *do* have a 144-drive system, click **OK**.

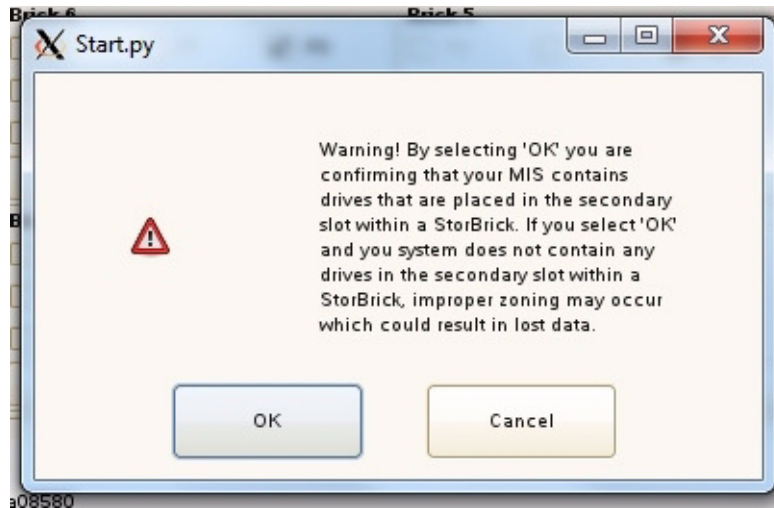


Figure 3-8 Zones 1.4.2 – Warning 144 Drive Configuration

The following diagram (Figure 3-9) shows the layout of the primary and secondary drives in the StorBrick.

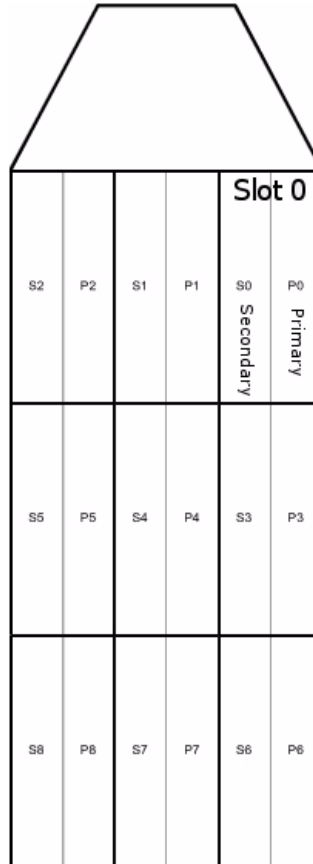


Figure 3-9 StorBrick Diagram for 144-Drive Systems

Saving Zoning Configuration

Once you have allocated the drives to the desired adapters, you can save the configuration by clicking on the button, **Step 2: Save Session**. This saves the configuration in the session folder, and enables the **Step 3: Download Session** button (Figure 3-2).

Clicking on the **Step 3: Download Session** button brings up the **Brick Selection** window (Figure 3-10). This allows you to select which StorBricks will be updated with the new configuration information. Once selected, clicking the OK button downloads the new configuration to the StorBricks. When complete, the **Brick Selection** window will disappear.

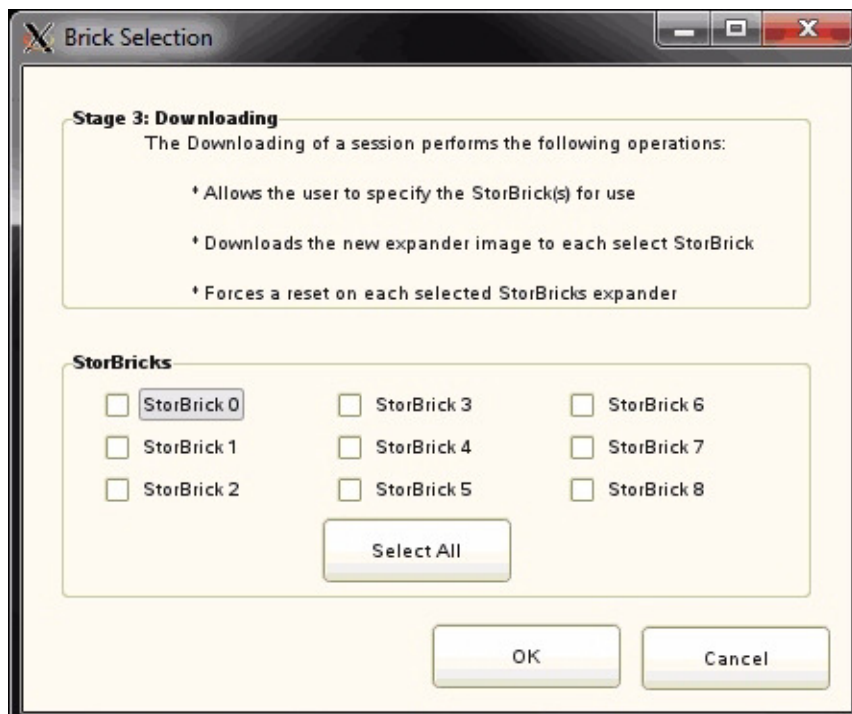


Figure 3-10 Zones 1.4.2 – Brick Selection Window

To complete the zoning process, the machine must be power-cycled (turned off and turned back on again) for the new configuration to re-flash the firmware files and accept the new settings for each StorBrick.

Zoning MIS JBOD Platforms

There are two scenarios for zoning JBODs: the first assumes the JBOD is networked to an MIS Server platform, the second is when a JBOD is networked to a non-MIS server. If the JBOD is connected to an MIS Server, the JBOD can be zoned by clicking on **Setup** from the menu bar (Figure 3-7) and selecting **Zone JBOD**.

To zone a JBOD attached to a non-MIS server, change to the `/opt/Zones/JBOD` directory and type `python JGUI.py`, or through Windows by double-clicking the `JGUI.py` icon in the

C:\Program Files (x86)\Zones\JBOD\ directory. That will bring up the **JBOD Zoning Tool interface** (Figure 3-11).

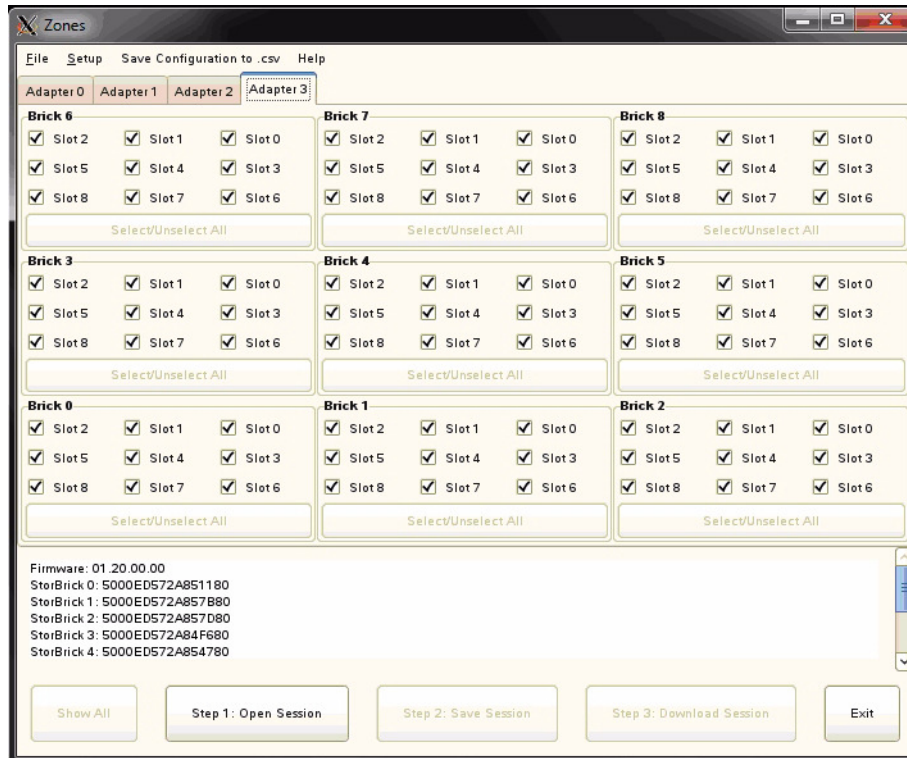


Figure 3-11 Zones 1.4.2 – JBOD Zoning Interface

To begin, click on the button, **Step 1: Open Session** (Figure 3-2).

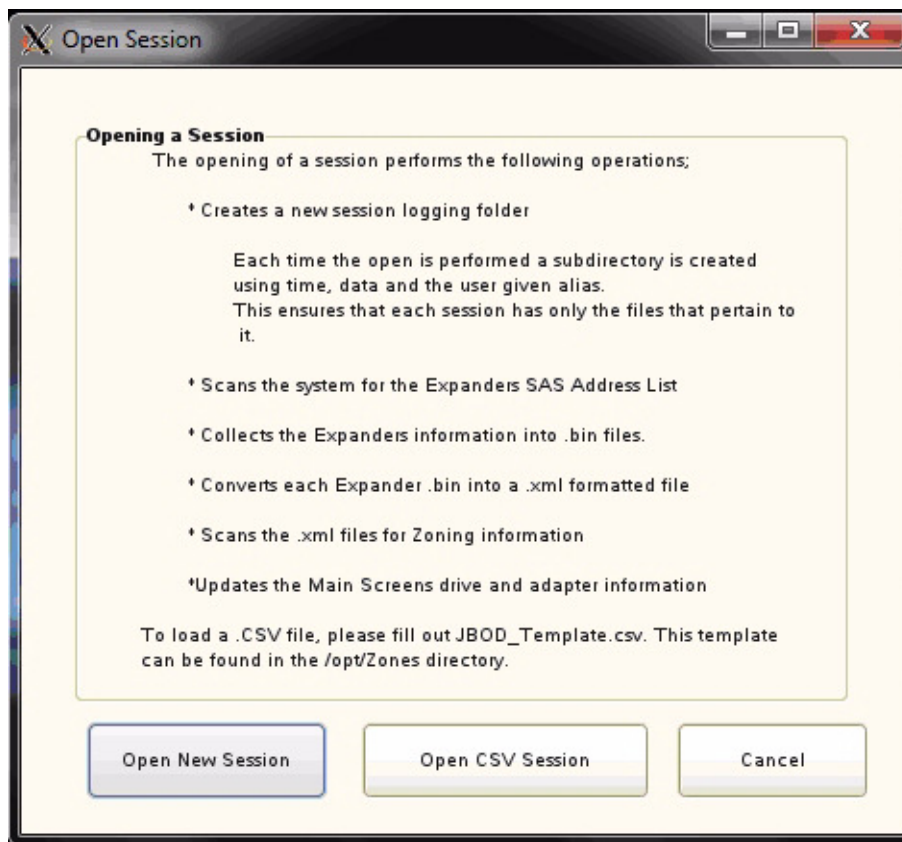


Figure 3-12 Zones 1.4.2 – JBOD Open Session Window

This displays the Open Session window (Figure 3-12). Here you have the option to **Open New Session**, **Open CSV Session**, or **Cancel** the operation. If you have a CSV with the desired zoning configuration, you can choose **Open CSV Session**. (See “Loading a CSV File to Zones,” on page 81 for more information.) Otherwise, click on **Open a New Session**.

Note: Figure 3-12 states: **To load a .CSV file, please fill out JBOD_Template.csv. This template can be found in the /opt/Zones directory.** If the JBOD_Template.csv template file is unavailable, another template can be found in the /opt/Zones/JBOD/ directory for CentOS & Linux, and in the C:\Program Files (x86)\Zones\JBOD directory for Windows. It is called First.csv.

This will bring up the **New Session** prompt, asking for a **Session Alias** to be associated with the zoning configuration. This alias name will be used to create a time-stamped folder with that alias, necessary to store the files used to zone the machine. After entering a session alias, click **OK** (Figure 3-13).

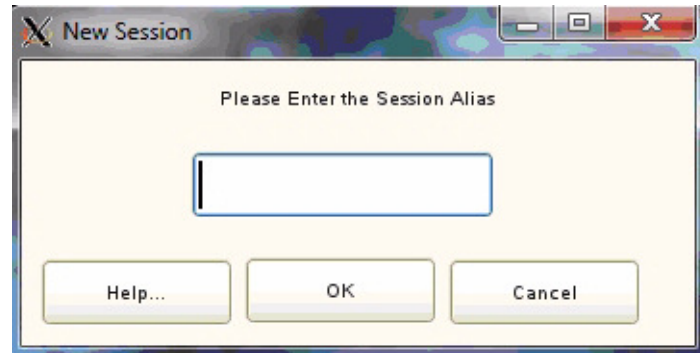


Figure 3-13 Zones 1.4.2 – JBOD New Session Alias

Phy-based Zoning

Phy-based Zoning is automatically enabled as the only option. **Phy-based zoning** allows you to split the drives **between two I/O CARDS** (Figure 1-13 on page 10), though it also may be used when only one I/O card is present.

With **Phy-based zoning**, the main JBOD zoning window will appear with Adapter 0 and Adapter 2 available for zoning (Figure 3-14).

The two adapter tabs correspond to the two JBOD I/O cards the drives can be split between. Which adapter tab (adapter 0 or adapter 2) belongs to which JBOD IO-card is displayed in the text window at the bottom of the interface.

Next, enable all the drives by clicking the **Show All** button at the bottom of the home screen. This will enable all the slots, and you can now select which drives you would like associated with which adapter.

Once you have allocated the drives to the desired adapters, you can save the configuration by clicking on the button, **Step 2: Save Session**. This saves the configuration in the session folder, and enables the **Step 3: Download Session** button.

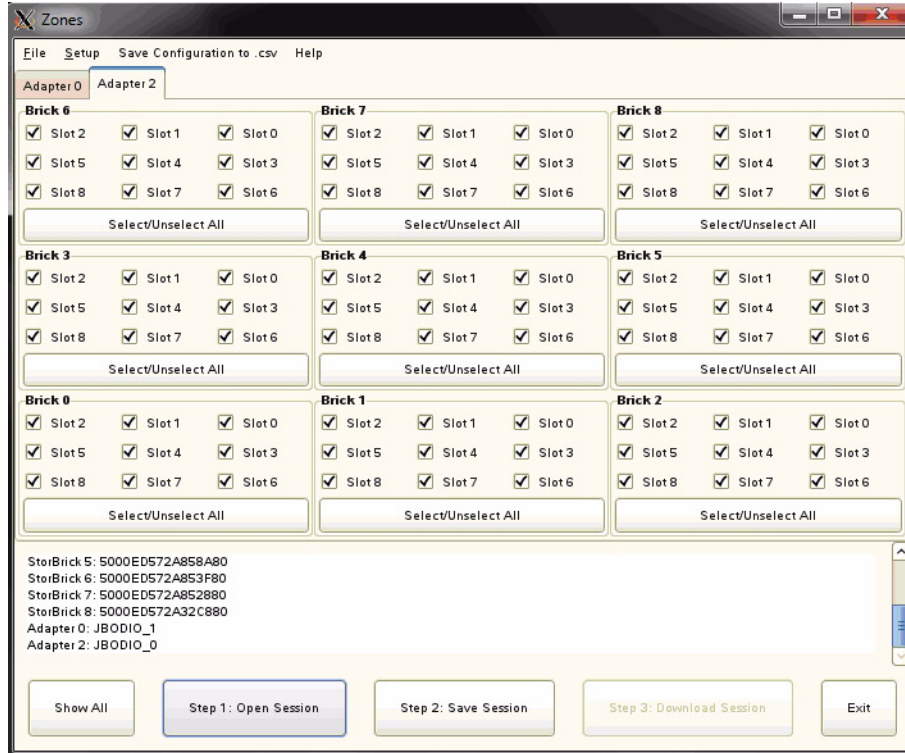


Figure 3-14 Zones 1.4.2 – JBOD Phy-based Zoning Interface

Clicking on the **Step 3: Download Session** button brings up the **Brick Selection** window (Figure 3-15). This allows you to select which StorBricks will be updated with the new configuration information. Once selected, clicking the **OK** button downloads the new configuration to the StorBricks. When complete, the **Brick Selection** window will disappear.

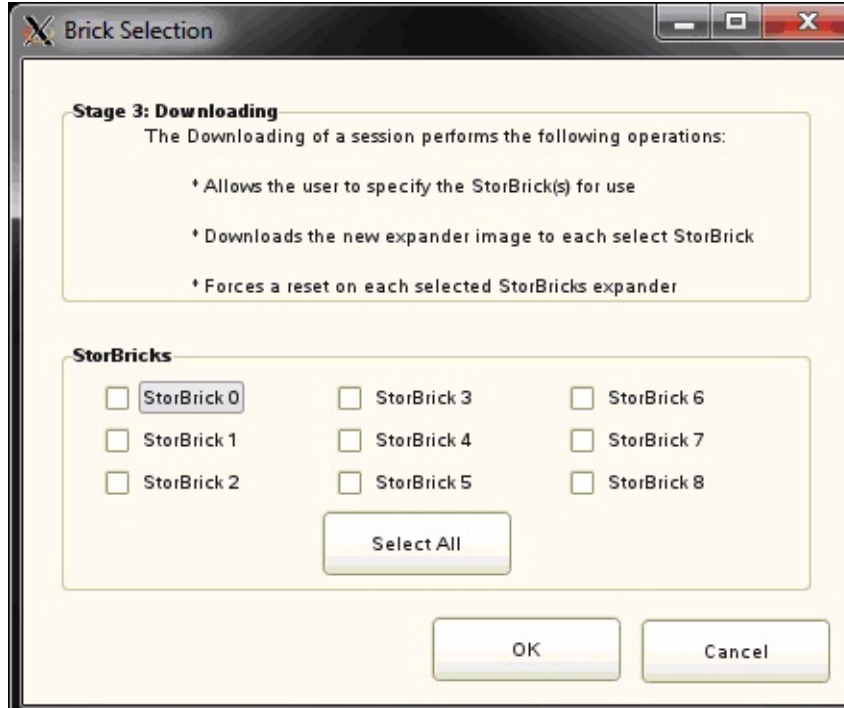


Figure 3-15 Zones 1.4.2 – JBOD Brick Selection Window

To complete the zoning process, the machine in question must then be power-cycled (turned off and turned back on again) for the new configuration to re-flash the firmware files and accept the new settings for each StorBrick.

Loading a CSV File to Zones

The following refers to both MIS Server Platforms and MIS JBOD³ Platforms. Many times it is economical to use a CSV file that contains the zoning configuration and load that into the machine, rather than creating a zoning configuration each time from scratch. When this is the case, open the tool appropriate to the machine on which you are working (either Server or JBOD zoning tools), and click **Step 1: Open Session** (Figure 3-2 on page 70).

³ Loading a CSV file always results in Phy-based zoning.

In the **Open Session** window (Figure 3-3 on page 71) select **Open CSV Session**. This will bring up the file selection window (Figure 3-16).

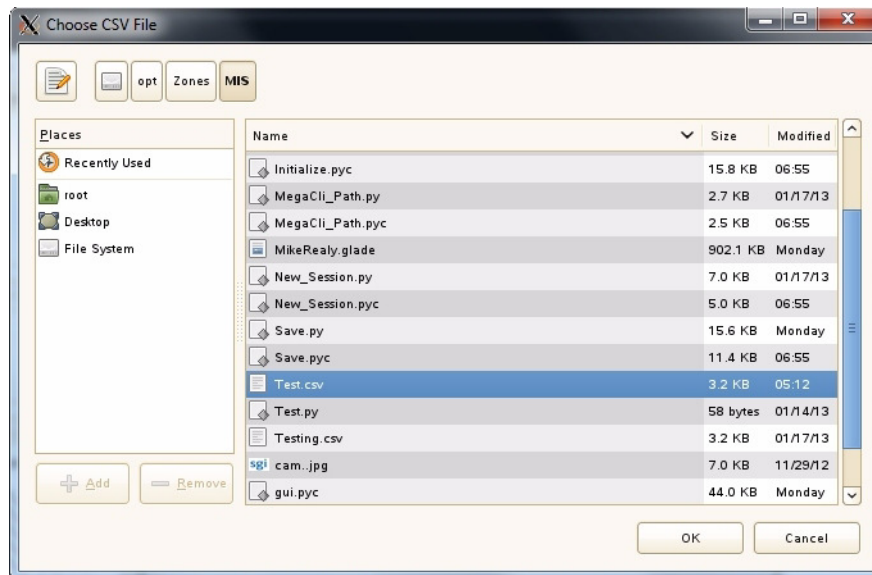


Figure 3-16 Zones 1.4.2 – Choose CSV File Window

Browse to the location to where the desired CSV file is stored, select it, and click **OK**. This will prompt you with a selection for **Primary Only** or **Primary and Secondary** (Figure 3-17).

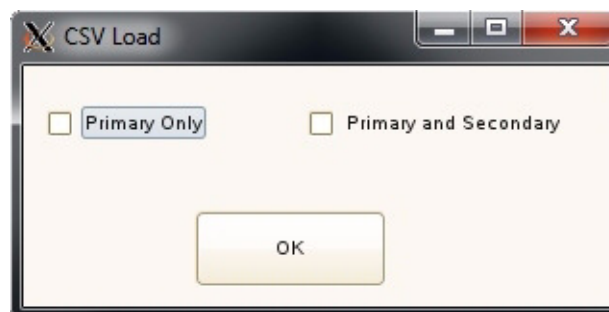


Figure 3-17 Zones 1.4.2 – CSV Load: Primary Only or Primary and Secondary

Place a check next to the box **Primary Only**, unless you are running a platform with a 144 drives (smaller profile 9mm drives that fit two drives—primary and secondary—per slot), in which case, choose **Primary and Secondary**. After making your selection, click **OK**.

A **Warning** box, asking you if you are sure you want to continue (Figure 3-18).




Figure 3-18 Zones 1.4.2 – CSV Warning Window

Clicking **OK** loads the configuration file to the StorBricks. Once complete, the **Warning** box will disappear.

The machine must now be power-cycled (turned off and back on again) for the changes to re-flash the firmware files and take effect the new settings for each StorBrick.

Creating the Drive Groups in MegaRAID (CentOS, Linux & Windows)

Power on the machine to complete power-cycle. Open MegaRAID. Right-click on the expander and select **Create a Virtual Drive** (Figure 3-19).




Figure 3-19 MegaRAID – Create a Virtual Drive

A screen will pop up asking you to choose **Simple** or **Advanced** (Figure 3-20). In Simple mode, the drives are chosen for you. In Advanced mode, you get to choose the drives, and are given additional selections in RAID levels, allowing for spanned (00, 10, 50, 60) drives.

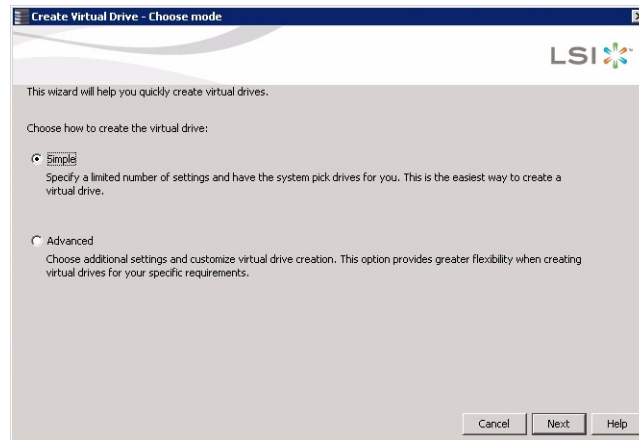


Figure 3-20 MegaRAID – Create Virtual Drive Mode

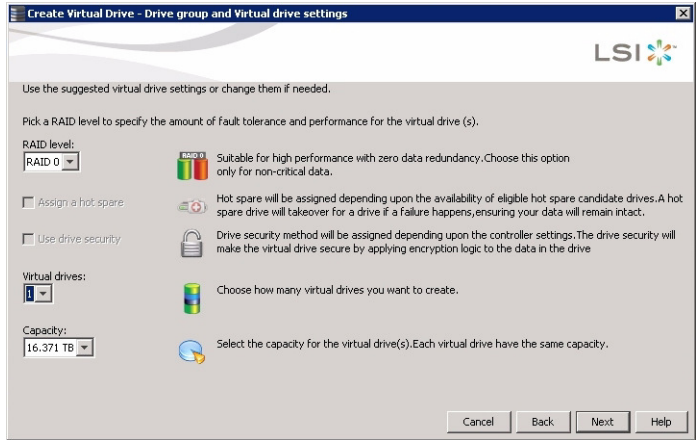


Figure 3-21 Create Virtual Drive – Drive Group Settings

Choose **Write Back BBU** (battery back-up unit). This mode is the safest and the fastest, and will automatically switch from caching mode to write-straight-to-disk whenever battery power has reached low. **Write Through** writes straight to disk. **Write Back** is a cached data flow.

Warning: If you select **Write Back** and power to the system is lost, data is lost.

Click **Next**, and a summary screen verifying settings will appear (Figure 3-22).

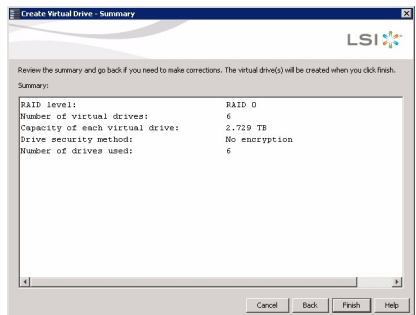


Figure 3-22 Create Virtual Drive – Summary

If the settings are correct, click **Finish**, and click **Ok**.

Formatting the Drives Using Command Line in CentOS

CentOS does not use YaST2 like other Linux-based systems. Instead, drives may be formatted and mounted using command line.

You will need to be at the `root` or add `sudo` before each command.

Verify Drives Seen

To view the drives already recognized by the system, issue the following command:

```
[root@server]# ls /dev/sd*
```

Adding a New Drive

If adding a new drive, you will need to make a partition and format it. If not, skip the following steps. These steps will create a single partition that will take up the entire drive.

```
[root@server]# fdisk /dev/sdb
```

```
WARNING: DOS-compatible mode is deprecated. It's strongly recommended  
to
```

```
    switch off the mode (command 'c') and change display units to  
    sectors (command 'u').
```

```
Command (m for help):
```

Switch off DOS compatible mode and change the units to sectors:

```
Command (m for help): c  
DOS Compatibility flag is not set  
Command (m for help): u  
Changing display/entry units to sectors
```

Create a new partition by typing `n`. We will then select this as a primary partition by entering `p` and select this as the first partition by entering `1`.

```
Command (m for help): n  
Command action  
    e   extended
```

```

    p   primary partition (1-4)
p
Partition number (1-4): 1
First cylinder (1-182401, default 1):
Using default value 1
Last cylinder, +cylinders or +size{K,M,G} (1-182401, default 182401):
Using default value 182401

```

We can review our results by entering p:

```

Command (m for help): p

Disk /dev/sdb: 1500.3 GB, 1500301910016 bytes
255 heads, 63 sectors/track, 182401 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x00000000

```

Device	Boot	Start	End	Blocks	Id	System
/dev/sdb1		1	182401	1465136001	83	Linux

Write (save) the part ion by entering w.

```

Command (m for help): w
The partition table has been altered!

```

Formatting Drives

To format the drives, issue the command `mkfs.ext4 /dev/sdb1`

```

mke2fs 1.41.12 (17-May-2010)
Filesystem label=
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
91578368 inodes, 366284000 blocks
18314200 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=0
11179 block groups
32768 blocks per group, 32768 fragments per group

```

```
8192 inodes per group
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632,
2654208,
    4096000, 7962624, 11239424, 20480000, 23887872, 71663616, 78675968,
    102400000, 214990848
```

This will take some time, and `Writing inode tables: with numbers will detail progress.`

Mount the New Drive

If you don't already have a mount point, make one. This is the directory you will be able to access the newly mounted drive.

```
[root@server]# mkdir /mnt/backup
```

We can now mount the partition we created earlier to the new mount point:

```
[root@server] mount /dev/sdb1 /mnt/backup
```

Formatting the Drives Using YaST2 in Linux

Drives may be formatted using YaST2 Partitioner. In Linux, the folders that the drives will be mounted to need to be created first. Each mount will need a new folder. Some Linux customers will have the ability to issue the YaST2 command, bringing up a GUI to partition drives. Otherwise, drives are formatted and mounted using command line.




Figure 3-23 YaST2 – Server Manager GUI

1. Issue YaST2 command (i.e., # `yast2`) to launch the YaST2 Server Manager GUI (Figure 3-23).
2. Double-click **Partitioner** to launch.
3. A warning message will appear (Figure 3-24). Click **Yes**.



Figure 3-24 YaST2 – Warning Message

Verify Drives Seen

4. Verify that all of your disks have appeared under **Hard Disks** (Figure 3-25).

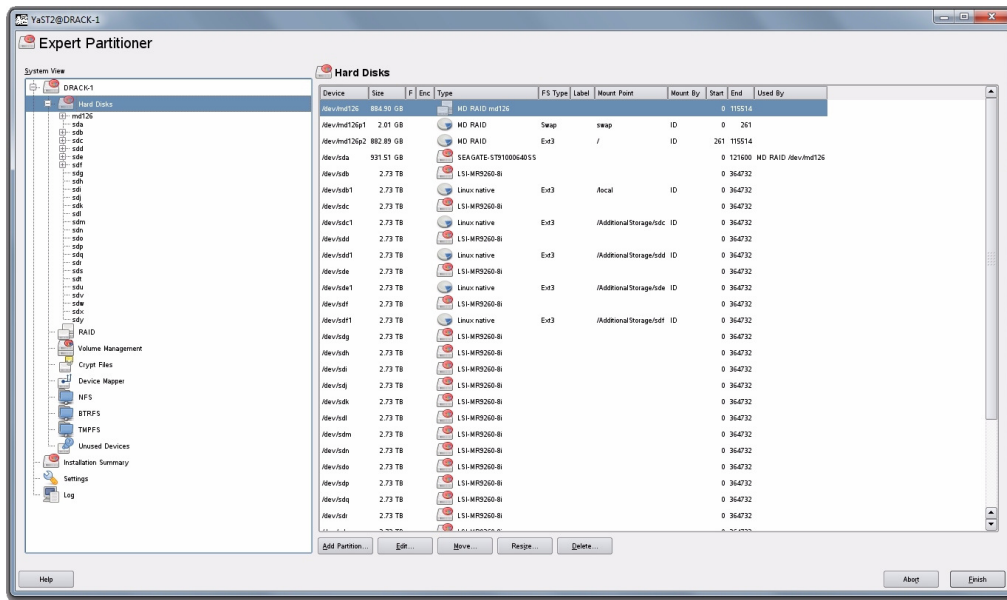


Figure 3-25 YaST2 – Drives Have Appeared

Partitioning Drives

5. Under **Hard Disks**, select the disk you would like to partition and click **Add** at the bottom of the screen (Figure 3-26).




Figure 3-26 YaST2 – Add Button

6. Select the partition size (Figure 3-27) and click **Next**.




Figure 3-27 YaST2 – Select Partition Size

7. Format the partition using `ext3`, mount the disk to your desired folder, and click **Finish** (Figure 3-28).



Figure 3-28 YaST2 – Format & Mount the Drive

8. Verify the partition shows up (Figure 3-29) and click **Next**.

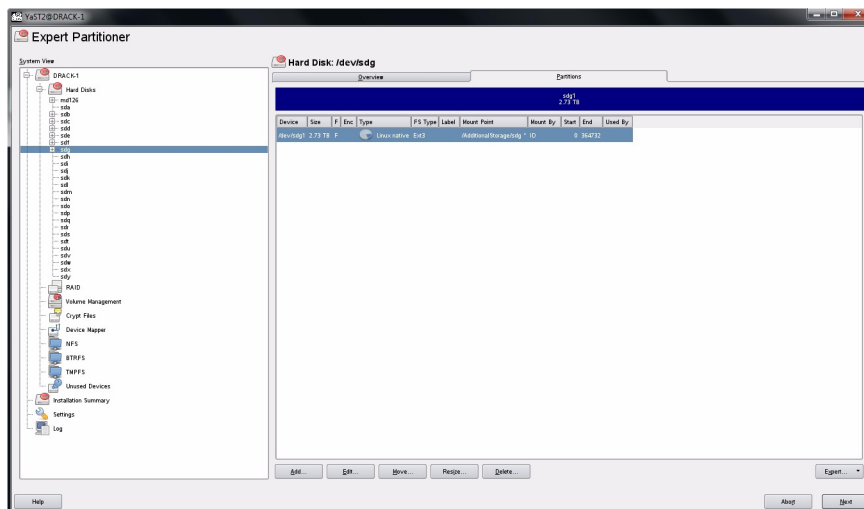


Figure 3-29 YaST2 – Check for Partition

9. Click **Finish** (Figure 3-30).

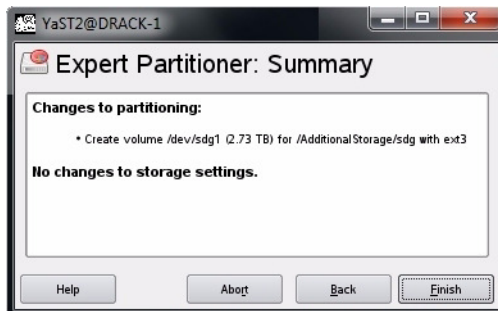


Figure 3-30 YaST2 – Click Finish

It may take several minutes for the system to mount the disk.

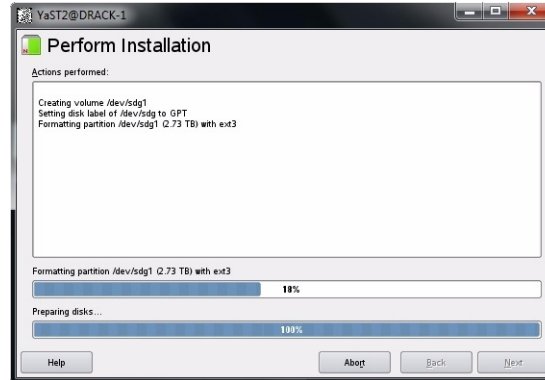


Figure 3-31 YaST2 – Disk Mounting (in process)

Once the disk is mounted (Figure 3-31), the system will return you to the beginning YaST2 GUI.

Windows Server Manager

In Windows, open Server Manager (Figure 3-32).

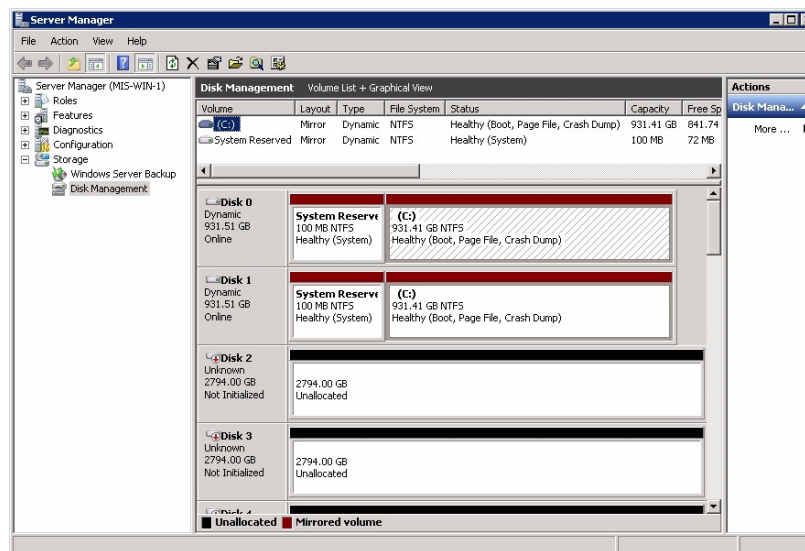


Figure 3-32 Windows Server Manager – Disk Management

Verify Drives Seen

Verify that the only disks the system sees are system drives, and that they are labeled C : \.

Note: Unconfiguring drives removes them from the system.

Formatting the Drives

Drives are formatted using Windows Server Manager (Figure 3-33). Open **Server Manager**—the screen should start at **Disk Management** with the drives showing. If not, click **Storage** in the system tree, and click **Disk Management**. The collection of disks/raidsets will now show.

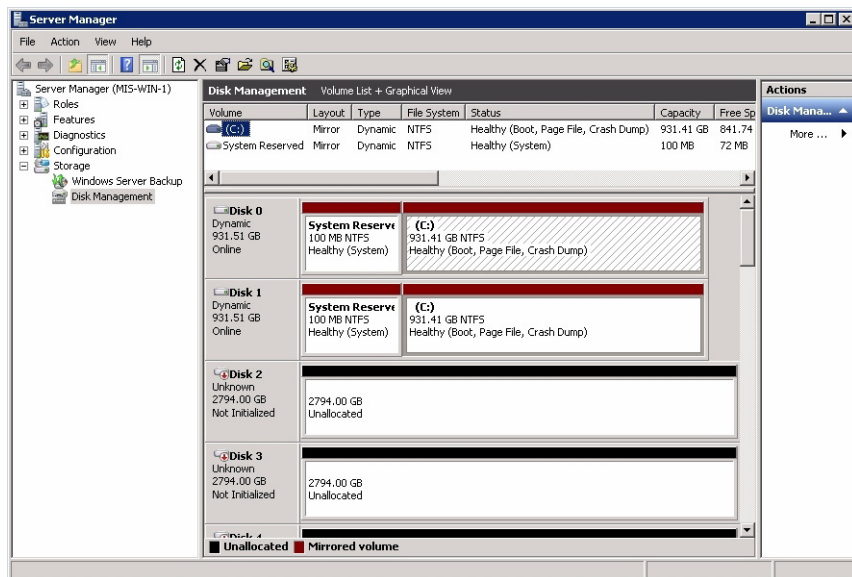


Figure 3-33 Server Manager – Disk Management

Right click in the grey area of the first non-system disk. In the menu that appears choose **Initialize Disk** (Figure 3-34).




Figure 3-34 Server Manager – Initialize Disks

A pop-up window will appear, showing all the uninitialized disks (Figure 3-35).

Warning: Be sure to select **GPT** (GUID Partition Table).



Figure 3-35 Server Manager – Select GPT (GUID Partition Table)

Click **OK**. All the disks should now show as **Online** (Figure 3-36).




Figure 3-36 Server Manager – Disks Initialized and Online

Right click the first non-system disk. Select **New Simple Volume** (Figure 3-37).




Figure 3-37 Server Manager – New Simple Volume

Click **Next** at the **New Simple Volume Wizard** welcome screen (Figure 3-38). Select the size of the volume in MB and click **Next** (Figure 3-39).




Figure 3-38 Server Manager – New Simple Volume Wizard

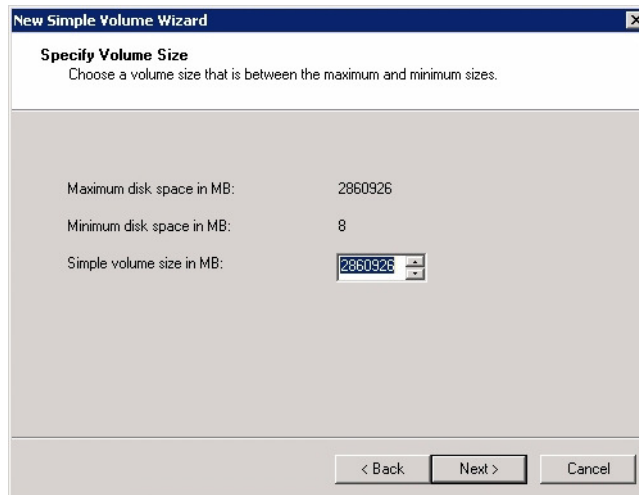


Figure 3-39 New Simple Volume Wizard – Volume Size

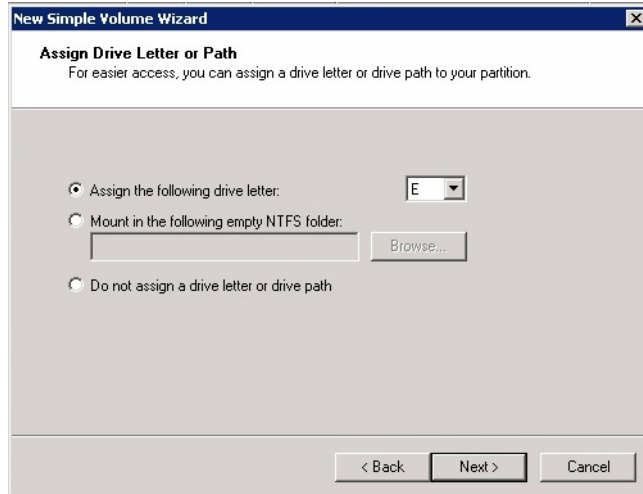


Figure 3-40 New Simple Volume Wizard – Assign Drive Letter or Path

Choose the drive letter to be assigned or click **Next** for the next drive letter available (Figure 3-40).

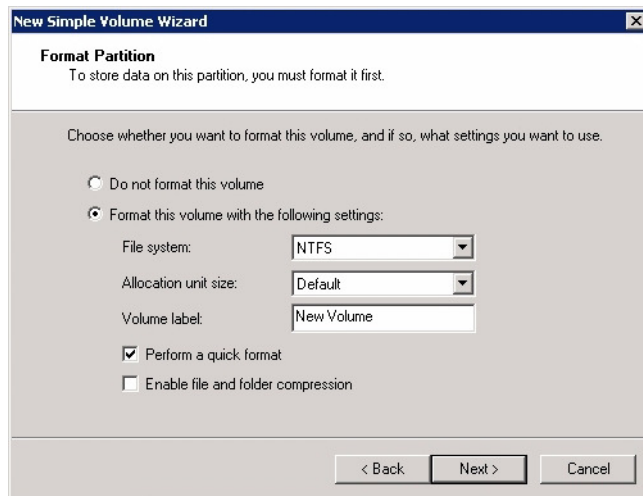


Figure 3-41 New Simple Volume Wizard – Format Partition

Select the format settings to be used, and click **Next** (Figure 3-41).

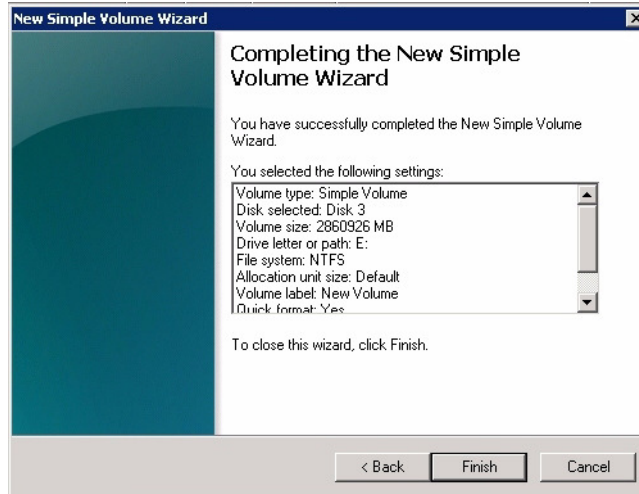


Figure 3-42 New Simple Volume Wizard – Settings Confirmation

Click **Finish** to format the disks (Figure 3-42). New volumes will show in the Disk Management window below the volumes (Figure 3-43).




Figure 3-43 Server Manager – New Simple Volume

CLI Zoning Tool, version 1.4

Note: T10-based zoning is not available for JBODs at this time. Future releases will support T10-based zoning for JBODs

This updated release of the CLI Zoning tool introduces proprietary SGI software, used to zone drives on a MIS Server Platform and MIS JBOD units. The tool also supports diagnostic functions useful to Field Service. The CLI Zoning Tool contains the following features:

- Allows for Zoning JBODs with a single JBOD-IO PCBA in PHY-Based Zoning.

- Includes the 'Expert Mode' commands to aid in diagnosing issues with MIS Enclosures.
- Real-time PHY configuration information available
- Connection to JBOD-IO or fan base for command execution
- Prints JBOD-IO status and fan base status
- Ability to erase Persistent Flash regions programmed by ShackCLI
- New 'hints' in Initialization file for JBOD

Installing CLI Zoning Tool on CentOS or Linux Host

The CLI Zoning Tool is proprietary SGI software, used to zone drives on the MIS Server Platform and MIS JBOD units. To install the software:

1. Create a directory for the application: `mkdir /opt/ShackCLI`
2. Go to <http://support.sgi.com> and log in with your credentials.
3. Download the latest version of **ShackCLI_release_xxx.zip** into the above directory
4. Extract the files: `unzip ShackCLI_release_xxx.zip`

Installing CLI Zoning Tool on Windows Host

The CLI Zoning Tool on a Windows host machine **requires the presence of Python**, version 2.6 or 2.7 installed on the host. Follow the instructions given in the subsequent section, “Installing Python for Windows,” prior to running the CLI Zoning Tool.

To install the CLI Zoning Tool software:

1. Go to <http://support.sgi.com> and log in with your credentials.
2. Download the CLI Zoning Tool for Windows, version 1.4.
3. Unzip `ShackCLI_release_xxx`. The program will ask if you want to create a directory name for the `.zip` file, `c:\python##\ShackCLI\`, where `##` is the version number of Python installed.
4. Click **Yes**.

Basics of CLI Zoning

Note: T10 Zoning is unsupported at this time. Future versions of the CLI Zoning Tool will support T10-based zoning.

The basic idea of Zoning is fairly simple, if you don't need to know the mechanics of T10 or PHY based zoning. The goal is to allow Initiators to see specific disk drives or not see specific disk drives based on your zoning policy. Typically zoning should be configured by the customer to spread across the StorBricks as much as possible. When using RAID controllers, it is **not** recommended to have multiple Initiators seeing the same drive but is supported if the user so chooses.

The ShackCLI uses Comma Separated Values (CSV) files that the user creates to apply Zoning changes. The format of the CSV file is meant to simplify the Zoning configuration down to a simple Initiator to drive connection matrix. The ShackCLI expects a specific format for all CSV files used for Zoning:

Table 3-1 Zoning Configuration

Drive	Initiator 1	Initiator 2	Initiator 3	Initiator4	Initiator 5	Initiator 6	Initiator 7	Initiator 8
0	1	0	0	0	0	0	0	0
1	0	1	0	0	0	0	0	0
2	0	0	1	0	0	0	0	0
3	0	0	0	1	0	0	0	0
4	1	0	0	0	0	0	0	0
5	0	1	0	0	0	0	0	0
6	0	0	1	0	0	0	0	0
7	0	0	0	1	0	0	0	0
8	1	0	0	0	0	0	0	0
9	1	0	0	0	0	0	0	0
10	0	1	0	0	0	0	0	0
11	1	0	0	0	0	0	0	0

Drive	Initiator 1	Initiator 2	Initiator 3	Initiator4	Initiator 5	Initiator 6	Initiator 7	Initiator 8
12	0	0	0	0	0	0	0	0
13	0	0	0	0	0	0	0	0
14	0	1	0	0	0	0	0	0
15	0	0	0	0	0	0	0	0
16	0	0	0	0	0	0	0	0
17	0	0	0	0	0	0	0	0

Column 0 is the drive and columns 1-8 are the Initiators. Each StorBrick may contain up to 18 disk drives when using the SGI 2.5" drive adapter or up to 9 disk drives when using 3.5" disk drives. There is an implied relationship between the Drive Numbers and the StorBrick numbers where Drive 0-8 is StorBrick 0, drive 9-17 is StorBrick 1, etc. Blank rows are supported to help in formatting. If you are not using 2.5" drives, designate the drives 72-143 for Servers, and 82-163 for JBODs in the CSV file, and fill the cells with zeroes so the ShackCLI will know to not map any Initiators to those drives.

A '1' in a cell means the Initiator in that column will be connected to the drive in the row. A '0' in a cell means the Initiator in the column will not be connected to the drive in the row.

All Initiators may be connected to any or all drives or no Initiators may be connected to any drive. Any combination of Initiator to Drive connectivity possible in the CSV file will be supported in the Enclosure, therefore there is no combination of Initiator to drive Zoning that is not supported. It is up to the user to determine how they want to allocate drives to Initiators.

For servers, Initiator 0 is mapped to PHY 7 on the StorBricks and connects to HBA 0, Initiator 1 is mapped to PHY 6 on the StorBricks and connects to HBA 1, etc.

For JBODs, since there are up to two JBOD-IOs PCBAs, JBOD-IO 0 connects to PHY 7/6 on the StorBricks and JBOD-IO 1 connects to PHY 5/4 on the StorBricks. For the StorBrick 8 the PHY mapping is JBOD-IO 0 connects to PHY 5/4 and JBOD-IO 1 connects to PHY 7/6 on the StorBrick. The ShackCLI automatically programs JBODs with the appropriate Initiator to PHY Zoning so you do not have to handle StorBrick 8 any different than any other StorBrick.

If the user wants to read the current zoning on a StorBrick and selects menu option 2, the ShackCLI will print zoning information on a PHY-basis where Initiators are connected to specific PHYs on the StorBrick. In the following example, the ShackCLI is printing zoning information for StorBrick 0 in a JBOD Enclosure, and automatically allows for the wide ports previously

described. If you enter a 1 in the cell corresponding to StorBrick 0 drive 0, the display will show Initiator 0/1 connected to drive 0. If you enter a 1 in the cell corresponding to StorBrick 0 drive 1, the display will show Initiator 2/3 connected to drive 1. This is due to everything in T10- and PHY-based Zoning are connected with PHYs.

Table 3-2 Phy Based Zoning Table for StorBrick 0

Drive	Initiator 0	Initiator 1	Initiator 2	Initiator 3
00	X	X	0	0
01	0	0	X	X
02	X	X	0	0
03	0	0	0	0
04	0	0	0	0
05	0	0	X	X
06	0	0	0	0
07	0	0	0	0
08	0	0	0	0
09	0	0	0	0
10	0	0	0	0
11	0	0	0	0
12	0	0	0	0
13	0	0	0	0
14	0	0	0	0
15	0	0	0	0
16	0	0	0	0
17	0	0	0	0

In the T10 implementation, SAS zoning access control is implemented by linked switch and expander devices, with zoning enabled. These devices define a Zoned Portion of a Service Delivery System (ZPSDS). No host device intervention is required. Each zoning expander device

within the MIS enclosure maintains an identical zone permission table, so zone access control is maintained across the entire ZPSDS. The difference between the expanders lies in the definition of the PHY Zone groups that defines to which Zone Group (ZG) each of the 36 PHYs belongs. The Permission Table then maps the Initiator Zone Groups to the Target Zone groups.

The CLI Zoning Tool performs changes through the use of Comma Separated Values (CSV) files. A Session is defined as the act of querying, editing, saving, and downloading the expander's binary zoning information.

The CLI Zoning Tool runs on any CentOS/Linux/Windows host that has Python 2.6 or Python 2.7 installed. Most Linux-based systems come with Python. The target MIS JBOD/Server does not require an OS since the CLI Zoning Tool uses the Fan Base Ethernet connection to access the StorBricks. Table 3-3 is how the Zone Groups (ZG) are implemented in the MIS Enclosure.

Table 3-3 Zone Group Implementation

Zone Group	Description
ZG0	The Dead Zone that only talks to ZG 127
ZG0 – ZG1	Always enabled
ZG2–3	Enabled for initiators
ZG2	For Initiators that have SMP Zoning Access
ZG3	Initiators that have access to Broadcast
ZG4–7	Reserved per SAS Specification
ZG8–15	The eight possible initiators
ZG16–96	For drives 0–80 for 81 possible drives
ZG97–127	Reserved in the MIS implementation
ZG128-208	For drives 81 - 161. If using 9mm 2.5" drives

A configuration file is used by the CLI application to zone the StorBricks. **A set of standard configuration files are included with the CLI Zoning Tool software package.** A custom file can be created, using a spreadsheet application and then saving it as a `.csv` file (see “Editing the `.csv` File for the CLI Zoning Tool,” on page 111).

Preparing to Zone using the CLI Zoning Tool

CLI Zoning uses the MIS-S9D proprietary network interface. This interface is to be used **ONLY** when zoning using the CLI Zoning Tool. It is located at the front of the chassis at the upper right corner (Figure 3-44). The chassis must be slid out forward at least one inch in order to connect a network cable. (See “Sliding the Chassis Forward/Backward,” on page 133.) Ensure the MIS system is powered on.



Figure 3-44 MIS-S9D Proprietary Network Interface

Use an Ethernet cable to connect a server/laptop running either a CentOS, Linux, or Windows operating system and the CLI Zoning application software. The network port connected to the server/laptop must be set to 192.168.0.xxx (10 will do). The static IP address of the fan base is set to 192.168.0.3. Verify connectivity to the fan base with a ping command to 192.168.0.3 from the server/laptop. If there is no response, or the ping times out, it will be necessary to power cycle the MIS Server or JBOD.

Editing the ShackCLI.ini file for CentOS & Linux

In the /opt/ShackCLI directory, open up the ShackCLI.ini file with a text editor or vi.

The following is a sample file:

```
[main]
# Input filename. This must be either a pathname or a simple
# dash (-), which signifies we'll use standard in.
input_source = cli
```

```
target = 192.168.0.3
[maxsize]
# When we hit this threshold, we'll alert for maximum
# file size.
threshold = 100
[display]
show_footer = yes
# Fill up all SB information before going to Menu
auto_fill = no
[default]
#MIS_Variant = JBOD
MIS_Variant = SERVER
NUM_JBODIO = 2
#ETH_Variant = JBODIO
ETH_Variant = FANBASE
storbrick = 0 1 2 3 4 5 6 7
#storbrick = 0 1 2 3 4 5 6 7 8
cmd = menu
pcsv = /opt/ShackCLI/MIS-Server_4HBA_zoning_PCSV.csv
zcsv = /opt/ShackCLI/ZonePhy_Default.csv
pbcsv = /opt/ShackCLI/MIS-JBOD_1-IOMOD_zoning_PBCSV.csv
#pcsv = C:\Python##\ShackCLI\MIS-Server_4HBA_zoning_PCSV.csv
#zcsv = C:\Python##\hackCLI\ZonePhy_Default.csv
#pbcsv = C:\Python##\ShackCLI\MIS-JBOD_1-IOMOD_zoning_PBCSV.csv
```

First, verify that the target IP address is 192.168.0.3 and is not commented out (i.e., there is no # at the beginning of the line). Next, make the needed changes to the file, as follows.

1. Set the MIS Variant to the type of system to be zoned, JBOD or SERVER. (Be sure the other is commented out.)
2. Change the StorBrick count to be either 0–7 for a server, or 0–8 for a JBOD.
3. Select the type of zoning file to be used `pcsv`, `zcsv`, or `pbcsv`.
4. Add the path to where the configuration file is located.
5. Unless issues develop, leave the remaining selections at default.
6. Save the `ShackCLI.ini` file and close.
7. Execute the CLI command: `python ShackCLI.py --ini ShackCLI.ini --cmd menu`.

This will set the StorBricks to debug mode, and display a menu (“CLI Zoning Tool Main Menu,” on page 108).

Editing the ShackCLI.ini file for Windows

In the `C:\python##\` directory, open up the `ShackCLI.ini` file with an editor or vi.

The following is a sample file:

```
[main]
# Input filename. This must be either a pathname or a simple
# dash (-), which signifies we'll use standard in.
input_source = cli
target = 192.168.0.3
[maxsize]
# When we hit this threshold, we'll alert for maximum
# file size.
threshold = 100
[display]
show_footer = yes
# Fill up all SB information before going to Menu
auto_fill = no
[default]
#MIS_Variant = JBOD
MIS_Variant = SERVER
NUM_JBODIO = 2
#ETH_Variant = JBODIO
ETH_Variant = FANBASE
storbrick = 0 1 2 3 4 5 6 7
#storbrick = 0 1 2 3 4 5 6 7 8
cmd = menu
pcsv = /opt/ShackCLI/MIS-Server_4HBA_zoning_PCSV.csv
zcsv = /opt/ShackCLI/ZonePhy_Default.csv
pbcsv = /opt/ShackCLI/MIS-JBOD_1-IOMOD_zoning_PBCSV.csv
#pcsv = C:\Python##\ShackCLI\MIS-Server_4HBA_zoning_PCSV.csv
#zcsv = C:\Python##\hackCLI\ZonePhy_Default.csv
#pbcsv = C:\Python##\ShackCLI\MIS-JBOD_1-IOMOD_zoning_PBCSV.csv
```

First, verify that the target IP address is `192.168.0.3` and is not commented out (i.e., there is no `#` at the beginning of the line). Next, make the needed changes to the file, as follows.

1. Set the MIS Variant to the type of system to be zoned, JBOD or SERVER. (Be sure the other is commented out.)
2. Change the StorBrick count to be either 0-7 for a server, or 0-8 for a JBOD.
3. Add the path to where the configuration file is located.

4. Unless issues develop, leave the remaining selections at default.
5. Save the `ShackCLI.ini` file and close.
6. Execute the CLI command: `python ShackCLI.py --ini ..\ShackCLI.ini --cmd menu`.

This will set the StorBricks to debug mode, and display a menu (“CLI Zoning Tool Main Menu”).

CLI Zoning Tool Main Menu

Version 1.4.0 Default Menu

```
#### Zoning Commands ####
1) Set Active StorBrick(s)
2) Display Current StorBrick(s) Zoning
3) Update StorBrick(s) Permissions Table From CSV
4) Update StorBrick(s) Phy Zones From CSV
5) Update StorBrick(s) Phy Based Zoning from CSV
6) Change Zoning type (Server <-> JBOD)
7) Save current StorBrick(s) Zoning to CSV
8) Toggle Expert Mode
0) Exit CLI - back to command prompt
Please enter your selection:
```

When the ShackCLI is used in menu mode the simplified menu is displayed. This menu only includes commands associated with the Zoning features of the ShackCLI. Operation of the ShackCLI in terms of Zoning has not changed from version 1.0, however the ShackCLI can now zone JBODs that have only one JBOD-IO in either T10 or PHY Based Zoning configurations. With T10 Zoning in a single JBOD-IO configuration it is now possible to zone up to four Initiators to access the disk drive array.

Warning: If you have opted to use T10 Zoning in a single JBOD-IO configuration, you *must* revert back to PHY Based zoning *before* installing a second JBOD-IO.

Table 3-4 CLI Zoning Tool Menu Options and Descriptions

Menu Option	Description
1) Set Active Storbrick(s)	This menu option allows user to select the StorBrick(s) to act upon. The StorBricks may be entered in any order: 0 1 2 3 4 5 6 7 or 7 6 5 4 3 2 1 0, or in subsets: 0 or 0 1, etc. Storbrick numbers must be less than or equal 7 for MIS Server and less than or equal 8 for MIS JBOD.
2) Display Current StorBrick(s) Zoning	Displays the Zoning configuration that is currently stored in the StorBricks.
3) Update StorBrick(s) Permissions Table From CSV	This menu option uses the <code>csv</code> file described in the <code>ini</code> file under the heading 'p <code>csv</code> ' to modify the T10 Zoning Permission Tables for the selected StorBricks. If no <code>csv</code> file has been specified in the <code>ini</code> file the CLI Zoning Tool will prompt you for the name of the <code>csv</code> file to use.
4) Update StorBrick(s) Phy Zones From CSV	This menu option uses the <code>csv</code> file described in the <code>ini</code> file under the heading 'z <code>csv</code> ' to modify the T10 PHY Zone Groups for the selected StorBricks. If no <code>csv</code> file has been specified in the <code>ini</code> file the CLI Zoning Tool will prompt you for the name of the <code>csv</code> file to use.
5) Update StorBrick(s) Phy Based Zoning from CSV	This menu option uses the <code>csv</code> file described in the <code>ini</code> file under the heading 'pb <code>csv</code> ' to modify the Phy-based Zoning Tables for the selected StorBricks. This is the only supported Zoning configuration for MIS JBOD and is an optional configuration for MIS Server and MIS DC Server. Only one of Phy-based and T10 Zoning should be implemented within an MIS Server (although it is technically possible to mix the Zoning types) and only one of Phy-based Zoning may be supported in an MIS JBOD.

Table 3-4 CLI Zoning Tool Menu Options and Descriptions (**continued**)

Menu Option	Description
6) Change Zoning type (Server <-> JBOD)	This menu option allows you to change the zoning type from SERVER/JBOD to JBOD/SERVER. This command will cause the Selected STORBRICKS T10 Supported flag to be set/unset depending what the current Zoning type is. For example, if the current Zoning type is Phy based and you select this option, then the T10 Zoning Supported flag will be set enabling T10 Zoning to be implemented instead.
7) Save current StorBrick(s) Zoning to CSV	This menu option allows you to save the current configuration of the Enclosures Zoning in a file. This file is compatible with the CLI commands that require a csv file to update Zoning. The csv file format for Phy-based and T10 Zoning are identical therefore one use of this command is to dump a MIS system's T10 Zoning configuration and then rewrite the same file as a Phy-based configuration.
8) Toggle expert mode	This menu option displays all the menu options available.
0) Exit CLI - back to command prompt	This menu option will exit the CLI Zoning Tool.

When ready to zone, complete the following instructions.

1. Execute option 7 to make a copy of the current configuration (be sure to add `.csv` as the file extension). Example: `MIS-System1-zoning-092012-121103.csv`
2. Edit the `.csv` file to the desired zone configuration. Once satisfied, save the file with a different name so as not to over write the saved one.
3. Change the name in the `ShackCLI.ini` file to point it to the new file.
4. Select the update option that fits the configuration (option 3 for T-10, or 5 for PHY).
5. When finished, select option 8 to show the expanded menu, and then option 13 to reset the fan base.
6. Choose option 6, when it asks if you really want to proceed, select Yes. Next, select the type of zoning that was set in the configuration (0 for Phy-based, or 1 for T-10).
7. Select option 13 again to reset the fan base.
8. Verify that the zoning is correct by executing option 2 and reviewing the configuration file.

9. Exit the CLI.

10. Power-cycle the MIS server and reboot any head-of-string controller of a JBOD to refresh the information in the servers.

Editing the .csv File for the CLI Zoning Tool

A configuration file made of comma separated values (a “csv” file) is used by the CLI application to zone the StorBricks. A set of standard configuration files are included with the software package. A custom file can be created, using a spreadsheet application and then saving it as a .csv file.

Figure 3-45 on page 112 is a block diagram of MIS JBOD StorBrick SB0. The other eight StorBricks for the MIS JBOD repeat this, but are offset to other lanes from the I/O modules.

Table 3-5 gives a portion of a csv file. A 1 in the spreadsheet indicates that the drive in question is accessible by the above HBA. The 0 indicates that the path is disabled.

Table 3-5 Zone Group Implementation

SB0 Drive	HBA-A Indicator 1	HBA-B Indicator 2	HBA-C Indicator 3	HBA-D Indicator 4	N/U Indicator 5	N/U Indicator 6	N/U Indicator 7	N/U Indicator 8
0	1	0	0	0	0	0	0	0
1	0	1	0	0	0	0	0	0
2	0	0	1	0	0	0	0	0
3	0	0	0	1	0	0	0	0
4	1	0	0	0	0	0	0	0
5	0	1	0	0	0	0	0	0
6	0	0	1	0	0	0	0	0
7	0	0	0	1	0	0	0	0
8	1	0	0	0	0	0	0	0




Figure 3-45 Block diagram of MIS Server StorBrick SB0

Version 1.4.0 Expanded Menu

```

#### Zoning Commands ####
 1) Set Active StorBrick(s)
 2) Display Current StorBrick(s) Zoning
 3) Update StorBrick(s) Permissions Table From CSV
 4) Update StorBrick(s) Phy Zones From CSV
 5) Update StorBrick(s) Phy Based Zoning from CSV
 6) Change Zoning type (Server <-> JBOD)
 7) Save current StorBrick(s) Zoning to CSV
 8) Toggle Expert Mode
   #### System Testing Commands ####
   #### Expert Mode Commands ####
 9) Display StorBrick(s) Settings
10) Display CLI Settings
11) Enter StorBrick CLI (Must select a single StorBrick)
12) Reboot Fan Base - will not reset StorBricks
13) Reset Fan Base - Danger. This will reset StorBricks
14) Force StorBrick(s) into Debug Mode
15) Exit StorBrick(s) Debug Mode
16) Display PHY Error Counters for selected StorBricks
17) Display PHY Information for selected StorBricks
18) Display StorBrick UUID for selected StorBricks
19) Display StorBrick Firmware Revision Levels for Selected StorBricks
20) Change state of JBODIO ID LED - off, on blink
21) Clear Persistent Flash
22) Display Fan Base Status
23) Display JBODIO Status
24) Change Polling of Expanders by Fan Base or JBODIO
25) Display Disk Drive Status
26) Change Power State of a Disk Drive
27) Power Cycle test
28) Update StorBrick UUID
29) Write/Read Buffer Test
   0) Exit CLI - back to command prompt
Please enter your selection:

```

Select option 8 from the simplified menu to enter expert mode. When in Expert Mode additional features of the ShackCLI are made available to the user.

9) Display StorBrick(s) Settings

This option displays the following information about the StorBricks that have been selected in the Configuration file:

1. Total Errors – total communication errors logged by the ShackCLI during this session. This value should always be zero unless there are hardware or firmware errors in the Enclosure.
2. EXP SAS Address – SAS address of the expander on the StorBrick.
3. SXP SAS Address – SAS address of the target SMP for the expander on the StorBrick.
4. STP SAS Address – SAS address of the STP for the expander on the StorBrick.
5. Permission Table Loaded – When this value is Yes then the ShackCLI has read the current configuration of the Zoning Permission Table into the local memory.
6. Debug Mode Enabled – When this value is Yes then the selected StorBrick is in debug mode and ready to accept commands from the SEP interface.

Example Output of Menu option (9):

```
StorBrick 0:  
  Total Errors: 0  
  EXP SAS Address: 0x5000ED572A851180  
  SXP SAS Address: 0x5000ED572A8511BD  
  STP SAS Address 0x0000000000000003E  
  Permission Table Loaded: No  
  Debug Mode Enabled: Yes
```

10) Display CLI Settings

This menu option prints the current settings of ShackCLI. Generally these settings come from the Initialization file passed to the CLI at run time but may be changed by actions of the user during operation of the ShackCLI.

Example Output of Menu option (10):

```
MIS CLI current settings:  
MIS Variant: JBOD  
StorBricks:  
StorBrick 0 is active  
StorBrick 1 is active  
StorBrick 2 is active  
StorBrick 3 is active  
StorBrick 4 is active  
StorBrick 5 is active  
StorBrick 6 is active  
StorBrick 7 is active
```

```
StorBrick 8 is active
Using Fan Base at IP 10.4.3.229
Permission Table range: 0 through 255
Using CSV file for Permissions Table Update:
C:\Shackleford\ShackZzone\ShackZone\One_Initiators_72_drives.csv
Using CSV file for Phy Zone Update None
Using CSV file for Phy Based Zoning Update
C:\Shackleford\ShackZzone\ShackZone\Two_Initiators_36_drives_jbod.csv
CLI Version: Shack CLI: Version 1.4.0.0 Date: 01/24/2013
```

11) Enter StorBrick CLI (Must select a single StorBrick)

Menu option 11 allows Engineering to debug the StorBrick hardware and firmware and generally would not be used by customers or field service. If you wish to use the StorBrick CLI you must select a single StorBrick using the Initialization file or Menu option 1.

12) Reboot Fan Base - will not reset StorBricks

Menu option 12 will reboot the Fan Base without resetting the StorBricks.

13) Reset Fan Base - Danger. This will reset StorBricks

Menu option 13 will reset the Fan Base and the StorBricks. All the StorBricks in the Enclosure shall be reset regardless of the CLI settings.

14) Force StorBrick(s) into Debug Mode

Menu option 14 allows the user to place the selected StorBricks into Debug Mode which allows the ShackCLI to communicate with the StorBricks internal CLIs. At the start of ShackCLI the selected StorBricks shall be placed into Debug mode automatically. If the selected StorBricks changes during the session this menu option allows the user to place the newly selected StorBricks into Debug Mode.

15) Exit StorBrick(s) Debug Mode

Menu option 15 allows the user to place the selected StorBricks in default mode. This action will result in the selected StorBricks being reset.

Warning: Use this option at your own risk if you are using the ShackCLI on an enclosure that is performing I/O.

16) Display PHY Error Counters for selected StorBricks

Menu option 16 prints the PHY Error counters for the selected StorBricks. PHYs with an abnormal error count may be experiencing issues requiring a service action. Note that the current release of the Fan Base Firmware will truncate the PHY Event counters section of the output due to overflow of the receive buffer. The PHY Events are used for Internal StorBrick Firmware functions and may change in subsequent StorBrick Firmware releases.

- > InvWrdCnt - Invalid work count error count.
- > DispErrCnt - Running Disparity error count.
- > LossSyncCnt - Loss of Sync error count.
- > RstSeqFailCnt - Reset sequence failed count.

PHY mapping to drive slots:

PHY	Drive Slot
0	11
1	8
2	25
3	13
4	17
5	22
6	15
7	19
8	20
9	10
10	9
11	24
12	12
13	16
14	23
15	14
16	18
17	21

Note: The PHY error counters for the drive slots are typically zero. If you see large values in any of the PHY errors for drive slots you probably have an issue with that drive slot that needs attention.

```
INFO: PHY Error Counters for StorBrick 0
```

```
=====
Phy Layer Error Counters
=====
```

```
InvWrldCnt DispErrCnt LossSyncCnt RstSeqFailCnt
=====
```

```
Phy 00 0x00000000 0x00000000 0x00000000 0x00000000
Phy 01 0x00000000 0x00000000 0x00000000 0x00000000
Phy 02 0x00000000 0x00000000 0x00000000 0x00000000
Phy 03 0x00000000 0x00000000 0x00000000 0x00000000
Phy 04 0x00000000 0x00000000 0x00000000 0x00000000
Phy 05 0x00000000 0x00000000 0x00000000 0x00000000
Phy 06 0x000001e0 0x0000019a 0x0000001e 0x00000000
Phy 07 0x0000007a 0x00000065 0x0000001e 0x00000000
Phy 08 0x00000000 0x00000000 0x00000000 0x00000000
Phy 09 0x00000000 0x00000000 0x00000000 0x00000000
Phy 10 0x00000000 0x00000000 0x00000000 0x00000000
Phy 11 0x00000000 0x00000000 0x00000000 0x00000000
Phy 12 0x00000000 0x00000000 0x00000000 0x00000000
Phy 13 0x00000000 0x00000000 0x00000000 0x00000000
Phy 14 0x00000000 0x00000000 0x00000000 0x00000000
Phy 15 0x00000000 0x00000000 0x00000000 0x00000000
Phy 16 0x00000000 0x00000000 0x00000000 0x00000000
Phy 17 0x00000000 0x00000000 0x00000000 0x00000000
Phy 18 0x00000000 0x00000000 0x00000000 0x00000000
Phy 19 0x00000000 0x00000000 0x00000000 0x00000000
Phy 20 0x00000000 0x00000000 0x00000000 0x00000000
Phy 21 0x00000000 0x00000000 0x00000000 0x00000000
Phy 22 0x00000000 0x00000000 0x00000000 0x00000000
Phy 23 0x00000000 0x00000000 0x00000000 0x00000000
Phy 24 0x00000000 0x00000000 0x00000000 0x00000000
Phy 25 0x00000000 0x00000000 0x00000000 0x00000000
Phy 26 0x00000000 0x00000000 0x00000000 0x00000000
Phy 27 0x00000000 0x00000000 0x00000000 0x00000000
Phy 28 0x00000000 0x00000000 0x00000000 0x00000000
Phy 29 0x00000000 0x00000000 0x00000000 0x00000000
Phy 30 0x00000000 0x00000000 0x00000000 0x00000000
Phy 31 0x00000000 0x00000000 0x00000000 0x00000000
Phy 32 0x00000000 0x00000000 0x00000000 0x00000000
```

```
Phy 33 0x00000000 0x00000000 0x00000000 0x00000000
Phy 34 0x00000000 0x00000000 0x00000000 0x00000000
Phy 35 0x00000000 0x00000000 0x00000000 0x00000000
```

```
=====
Link Layer Event Counters
=====
```

```
Event # Event Type Event Threshold
=====
```

```
12ah 0x00000000
240h 0x00000000
341h 0x00000000
42dh 0x00000000
=====
```

```
Phy # EventData1 EventData2 EventData3 EventData4
=====
```

```
Phy 00 0x00000000 0x00000000 0x00000000 0x00000000
Phy 01 0x00000000 0x00000000 0x00000000 0x00000000
Phy 02 0x00000000 0x00000000 0x00000000 0x00000000
Phy 03 0x00000000 0x00000000 0x00000000 0x00000000
Phy 04 0x01606392 0x00000000 0x014904b8 0x000003bc
Phy 05 0x06b084c1 0xec636528 0x025a2c23 0x000003ab
Phy 06 0x0053e9df 0x00000000 0x004bd5a9 0x00000356
Phy 07 0x016b6a91 0x3237dda3 0x0088112f 0x00000359
Phy 08 0x03cf2b5b 0x01d18ca5 0x76319c04 0x00000001
Phy 09 0x00000000 0x00000000 0x00000000 0x00000000
Phy 10 0x00000000 0x00000000 0x00000000 0x00000000
Phy 11 0x00e45221 0x00717733 0x1afe520c 0x00000001
Phy 12 0x00000000 0x00000000 0x00000000 0x00000000
Phy 13 0x00000000 0x00000000 0x00000000 0x00000000
Phy 14 0x00000000 0x00000000 0x00000000 0x00000000
```

17) Display PHY Information for selected StorBricks

Menu option 17 displays details on all 36 PHYs for the selected StorBricks expanders. In the following example the PHY information for StorBrick 0 in a JBOD is displayed. Note that PHYs 7, 6 and 5, 4 are configured as wide ports and have the same SAS address associated.

In a Server configuration there are allowances for up to four HBAs with a single SAS lane attached to each of the HBAs on expander PHYs 7-4. The ShackCLI uses the information in the current configuration file to determine if the enclosure is a JBOD or a Server and configures zoning appropriately, without user intervention.

INFO: Retrieving PHY Information for StorBricks: [0]
 INFO: PHY Information for StorBrick 0

PHY	TYPE	NLR	PHY CNT	ZONE GRP	ZONE BUS	CTRL IIITTTA	STMSTMA PPPPPT	ATTACHED	SAS	ADDR	ROUTE TYPE	CONN TYPE INDX	CONN ELEM LINK	CONN PHY ID	MAP PHY	DFE ADAPT VAL
00		0x0	0x00	0x00	0x00	-----					D	0x00	0x00	0x00	000	
01		0x0	0x00	0x00	0x00	-----					D	0x00	0x00	0x00	001	
02		0x0	0x00	0x00	0x00	-----					D	0x00	0x00	0x00	002	
03		0x0	0x00	0x00	0x00	-----					D	0x00	0x00	0x00	003	
04	EXP	6.0	0x01	0x20	0x00	-----1-	5000ED57_2ED64300				T	0x00	0x00	0x00	004	0x1
05	EXP	6.0	0x01	0x07	0x00	-----1-	5000ED57_2ED64300				T	0x00	0x00	0x00	005	0x0
06	EXP	6.0	0x3f	0xcc	0x00	-----1-	5000ED57_DEAD203F				T	0x00	0x00	0x00	006	0x0
07	EXP	6.0	0x3f	0xee	0x00	-----1-	5000ED57_DEAD203F				T	0x00	0x00	0x00	007	0x1
08	END	6.0	0x01	0xe0	0x00	----1----	5000C500_40AB40C1				D	0x20	0x01	0x00	008	0x0
09	END	6.0	0x01	0x07	0x00	---1---	5000C500_40AB40C2				D	0x00	0x00	0x00	009	0x0
10	END	6.0	0x01	0x40	0x00	---1---	5000C500_4093FBD2				D	0x20	0x09	0x00	010	0x0
11	END	6.0	0x01	0x00	0x00	---1---	5000C500_4093FBD1				D	0x20	0x00	0x00	011	0x0
12		0x0	0x00	0x00	0x00	-----					D	0x20	0x0C	0x00	012	
13		0x0	0x00	0x00	0x00	-----					D	0x20	0x03	0x00	013	
14		0x0	0x00	0x00	0x00	-----					D	0x20	0x0F	0x00	014	
15		0x0	0x00	0x00	0x00	-----					D	0x20	0x06	0x00	015	
16		0x0	0x00	0x00	0x00	-----					D	0x20	0x0D	0x00	016	
17		0x0	0x00	0x00	0x00	-----					D	0x20	0x04	0x00	017	
18		0x0	0x00	0x00	0x00	-----					D	0x20	0x10	0x00	018	
19		0x0	0x00	0x00	0x00	-----					D	0x20	0x07	0x00	019	
20		0x0	0x00	0x00	0x00	-----					D	0x20	0x08	0x00	020	
21		0x0	0x00	0x00	0x00	-----					D	0x20	0x11	0x00	021	
22	END	6.0	0x01	0xdd	0x00	---1---	5000C500_40AAF695				D	0x20	0x05	0x00	022	0x0
23	END	6.0	0x01	0xee	0x00	---1---	5000C500_40AAF696				D	0x20	0x0E	0x00	023	0x0
24	END	6.0	0x01	0x60	0x00	---1---	5000C500_40AB1102				D	0x20	0x0A	0x00	024	0x0
25	END	6.0	0x01	0x07	0x00	---1---	5000C500_40AB1101				D	0x20	0x02	0x00	025	0x0
26		0x0	0x00	0x00	0x00	-----					D	0x00	0x00	0x00	026	
27		0x0	0x00	0x00	0x00	-----					D	0x00	0x00	0x00	027	
28		0x0	0x00	0x00	0x00	-----					D	0x00	0x00	0x00	028	
29		0x0	0x00	0x00	0x00	-----					D	0x00	0x00	0x00	029	
30		0x0	0x00	0x00	0x00	-----					D	0x00	0x00	0x00	030	
31		0x0	0x00	0x00	0x00	-----					D	0x00	0x00	0x00	031	
32		0x0	0x00	0x00	0x00	-----					D	0x00	0x00	0x00	032	
33		0x0	0x00	0x00	0x00	-----					D	0x00	0x00	0x00	033	
34		0x0	0x00	0x00	0x00	-----					D	0x00	0x00	0x00	034	
35		0x0	0x00	0x00	0x00	-----					D	0x00	0x00	0x00	035	
SXP	END	6.0	0x01	0x00	0x00	1-11---	5000ED57_2A8511BD				D	0x2F	0x00	0x00	036	
STP		0x0	0x00	0x00	0x00	----1--					D	0x2F	0x00	0x00	037	
EXP	END	6.0	0x01	0x00	0x00	---1-1-	5000ED57_2A851180				D	0x00	0x00	0x00	038	

18) Display StorBrick UUID for selected StorBricks

Menu option 18 displays the current UUID programmed into each of the selected StorBricks.

19) Display StorBrick Firmware Revision Levels for Selected StorBricks

Menu option 19 displays all the StorBrick firmware revisions of the selected StorBricks. The STORBRICK-x (where x = 0 in the example) should match the slot number of the StorBrick.

```
INFO: Retrieving Firmware Revision Information for StorBricks: [0]
INFO: Firmware Revision Information for StorBrick 0
```

```
=====
Hardware Revision Information:-
=====
```

```
Vendor ID:SGI CORP
Product ID:STORBRICK-0
Product Revision Level:Minor = 0x16, Unit = 0x00
Component ID:0x0223 (Bond Option :36)
Component Revision Level:5(B3)
```

```
=====
Firmware Revision Information:-
=====
```

```
Active Firmware: Active Image
```

```
Boot Image:
Revision: 1.22.0.0
Platform Name:   SGI - Shackleford STORBRICK
Version Name:    STORBRICK-01.22.00.00 01/21/13
Firmware Family: 1  OemFamily: 0
Fast Boot: No   Image Address: 0x14000000
```

```
Active Image:
Revision: 1.22.0.0
Platform Name:   SGI - Shackleford STORBRICK
Version Name:    STORBRICK-01.22.00.00 01/21/13
Firmware Family: 1  OemFamily: 0
Fast Boot: No   Image Address: 0x14080000
```

```
Backup Image:
Revision: 1.20.0.0
```

```
Platform Name:   SGI - Shackleford STORBRICK
Version Name:   STORBRICK-01.20.00.00 07/20/12
Firmware Family: 1  OemFamily: 0
Fast Boot: No  Image Address: 0x14100000
```

```
=====
SDK Revision:-
=====
```

```
HAL Revision: 1.0.0.6
SES Revision: 1.0.0.7
SCE Revision: 1.0.0.1
```

20) Change state of JBOD-IO ID LED - off, on blink

Menu option 20 changes the state of the JBOD-IO ID LED. This is the blue LED on the JBOD-IO PCBA and is generally used to locate the enclosure in the rack:

```
INFO: ID LED State is currently Off
Enter '0' to turn off, '-1' to turn on or '1' to blink:
```

21) Clear Persistent Flash

Menu option 21 will clear the persistent flash for all selected StorBricks. This option will clear all changes made by the ShackCLI over this session and all previous ShackCLI sessions and when the selected StorBricks are reset whatever has been loaded in the primary manufacturing image regions will become the current active image. Use this command if you have made changes with the ShackCLI from which you want to back out, or if there was an error reported by the ShackCLI during configuration.

22) Display Fan Base Status

Menu option 22 will display fan base status, including all the voltage and temperature sensor information the fan base has gathered from the StorBricks:

Please enter your selection: 22

FAN ID	FAN Name	Power	On	LED Blue	LED Green	FAN ID	FAN TACH
0	0	100	1	0	1	0	11160
1	1	100	1	0	1	2	11430
2	2	100	1	0	1	4	11520
3	3	100	1	0	1	6	11250
4	4	100	1	0	1	8	11160
5	5	100	1	0	1	10	11610

StorBrick	ID	Temp Sensor Name	Status_txt	Age	Reading
0	0	SBExpander Temp0	Ok	18054	29.0
0	1	Intake Temp0	Ok	8062	22.0
0	2	SBMid Temp0	Ok	28069	23.0
1	0	SBExpander Temp1	Ok	17120	30.0
1	1	Intake Temp1	Ok	7127	22.0
1	2	SBMid Temp1	Ok	27134	22.0
2	0	SBExpander Temp2	Ok	16187	28.0
2	1	Intake Temp2	Ok	6193	23.0
2	2	SBMid Temp2	Ok	26201	23.0
3	0	SBExpander Temp3	Ok	15253	29.0
3	1	Intake Temp3	Ok	5260	23.0
3	2	SBMid Temp3	Ok	25266	22.0
4	0	SBExpander Temp4	Ok	14319	30.0
4	1	Intake Temp4	Ok	4326	23.0
4	2	SBMid Temp4	Ok	24333	23.0
5	0	SBExpander Temp5	Ok	13386	38.0
5	1	Exhaust Temp5	Ok	3393	32.0
5	2	SBMid Temp5	Ok	23400	32.0
6	0	SBExpander Temp6	Ok	12451	39.0
6	1	Exhaust Temp6	Ok	2459	32.0
6	2	SBMid Temp6	Ok	22465	32.0
7	0	SBExpander Temp7	Ok	11518	39.0
7	1	Exhaust Temp7	Ok	1525	31.0
7	2	SBMid Temp7	Ok	21532	33.0
8	0	SBExpander Temp8	Ok	10583	39.0
8	1	Exhaust Temp8	Ok	590	32.0
8	2	SBMid Temp8	Ok	20598	31.0

StorBrick	ID	Volt Sensor Name	Status_txt	Age	Reading
0	0	1V	Ok	28044	0.998V
0	1	1.8V	Ok	18051	1.805V
0	2	5V	Ok	8059	5.006V
0	3	12V	Ok	48067	11.952V
0	4	3.3V	Ok	38074	3.260V
1	0	1V	Ok	27110	1.000V
1	1	1.8V	Ok	17118	1.802V
1	2	5V	Ok	7125	4.998V
1	3	12V	Ok	47132	11.968V
1	4	3.3V	Ok	37140	3.275V
2	0	1V	Ok	26176	1.001V
2	1	1.8V	Ok	16184	1.792V
2	2	5V	Ok	6192	4.992V
2	3	12V	Ok	46198	11.980V
2	4	3.3V	Ok	36207	3.265V
3	0	1V	Ok	25242	1.000V
3	1	1.8V	Ok	15250	1.793V
3	2	5V	Ok	5257	5.008V
3	3	12V	Ok	45265	11.964V
3	4	3.3V	Ok	35273	3.270V
4	0	1V	Ok	24308	1.000V
4	1	1.8V	Ok	14316	1.803V
4	2	5V	Ok	4324	5.008V
4	3	12V	Ok	44331	11.980V
4	4	3.3V	Ok	34339	3.253V
5	0	1V	Ok	23375	0.995V
5	1	1.8V	Ok	13382	1.801V
5	2	5V	Ok	3390	5.002V
5	3	12V	Ok	43397	11.948V
5	4	3.3V	Ok	33406	3.248V
6	0	1V	Ok	22441	0.996V
6	1	1.8V	Ok	12448	1.806V
6	2	5V	Ok	2456	4.990V
6	3	12V	Ok	42464	11.988V
6	4	3.3V	Ok	32471	3.263V
7	0	1V	Ok	21507	0.999V
7	1	1.8V	Ok	11515	1.803V
7	2	5V	Ok	1523	4.982V
7	3	12V	Ok	41531	11.996V
7	4	3.3V	Ok	31538	3.247V
8	0	1V	Ok	20574	0.997V
8	1	1.8V	Ok	10580	1.795V
8	2	5V	Ok	586	5.012V
8	3	12V	Ok	40596	11.968V
8	4	3.3V	Ok	30604	3.251V

23) Display JBOD-IO Status

Menu option 23 displays information regarding any JBOD-IO PCBA's that are in the enclosure, if the enclosure is in fact a JBOD. In the following example the ShackCLI is retrieving JBOD-IO status for a JBOD populated with two JBOD-IO PCBA's, and status for both JBOD-IOs is displayed:

JBODIO PCBA is in slot # 0:

Other JBODIO PCBA present: Yes

Temp	Name	Status	Status_txt	Timestamp	Age	Reading
Sensor						
0	Expander Temp	0	Ok	104021	1137	35
1	JBODIO Outlet Temp	0	Ok	100021	5195	27
2	LTC2991 Temp	0	Ok	102021	3252	32

Expander Voltage Sensors:

Volt	Sensor	Status	Status	Time	Age	Voltage
Sensor	Name	Status	Text	Stamp	Age	Reading
0	1V	0	Ok	98010	7322	0.998
1	1.8V	0	Ok	100010	5385	1.797
2	5V	0	Ok	102010	3448	4.986
3	12V	0	Ok	104010	1516	11.952
4	3.3V	0	Ok	96010	9583	3.274

PCBA Voltage Sensors:

Sensor	Status	Voltage	
Name	Status	Text	Reading
1.0V	1	normal	0.99
1.8V	1	normal	1.79
5V Standby	1	normal	4.95
12V	1	normal	11.89
Reference	1	normal	1.50

LED Name	State
power	on
status_green	on
status_amber	off
id	off
lan_activity	off
hdd_activity	off

JBODIO PCBA is in slot # 1:

Other JBODIO PCBA present: Yes

Temp Sensor	Name	Status	Status_txt	Timestamp	Age	Reading
0	Expander Temp	0	Ok	685440021	4502	39
1	JBODIO Outlet Temp	0	Ok	685442021	2562	28
2	LTC2991 Temp	0	Ok	685444021	621	30

Expander Voltage Sensors:

Volt Sensor	Sensor Name	Status	Status Text	Time Stamp	Age	Voltage Reading
0	1V	0	Ok	685438010	6690	0.997
1	1.8V	0	Ok	685440010	4754	1.794
2	5V	0	Ok	685442010	2822	4.982
3	12V	0	Ok	685444011	888	11.992
4	3.3V	0	Ok	685436010	8957	3.261

PCBA Voltage Sensors:

Sensor Name	Status	Status Text	Voltage Reading
1.0V	1	normal	1.00
1.8V	1	normal	1.79
5V Standby	1	normal	4.94
12V	1	normal	11.94
Reference	1	normal	1.49

LED Name	State
power	on
status_green	on
status_amber	off
id	off
lan_activity	off
hdd_activity	off

24) Change Polling of Expanders by Fan Base or JBOD-IO

Menu option 24 toggles the state of the fan base or JBOD-IO options in terms of polling StorBricks for voltage and temperature data. In versions of StorBrick Firmware prior to 1.21.0.0, the temperature and voltage polling was not supported while in debug mode, and the JBOD-IO and fan base would report issues with the SMBUS communications used for enclosure management. Use this command if you want to suspend polling of voltage and temperature data while using the ShackCLI.

25) Display Disk Drive Status

Menu option 25 displays information about the drives populated in the selected StorBricks. Drive slots that are not populated will show no connected SAS address:

StorBrick: 0

Drive	SAS Address	Power State	Interface Type	Link Rate
0	0x5000c5004093fbd1	On	SAS	6.0
1	0x5000c50040ab40c1	On	SAS	6.0
2	0x5000c50040ab1101	On	SAS	6.0
3	0x0	On	SAS	0x0
4	0x0	On	SAS	0x0
5	0x5000c50040aaf695	On	SAS	6.0
6	0x0	On	SAS	0x0
7	0x0	On	SAS	0x0
8	0x0	On	SAS	0x0

26) Change Power State of a Disk Drive

Menu option 26 allows the user to turn off/on power to the selected disk drives. The ShackCLI will ask for the StorBrick and drive slot to act upon.

27) Power Cycle test

Menu option 27 runs a drive power cycle test on the selected disk drive slots. This test should not be run while the server/JBOD is connected to a host.

28) Update StorBrick UUID

Menu option 28 allows the user to program the UUID into the StorBrick. This menu option would not be used by a customer since the UUID is programmed at the factory.

29) Write/Read Buffer Test

Menu option 29 runs a Write/Read buffer test on the selected StorBricks drive slots. This test utilizes the StorBricks Embedded SSP Initiator to issue Write/Read buffer commands to SAS drives in the enclosure. The user should take note of the PHY Error Counters prior to running this

test and then take another reading on the PHY Error Counters after running the test to determine if there are slot issues with the selected drive slots.

Disk RAID Support

The MIS Platform supports both software and hardware RAID, standard and nested. Disk performance is improved because more than one disk can be accessed simultaneously. Fault tolerance is improved because data loss caused by a hard drive failure can be recovered by rebuilding missing data from the remaining data or parity drives.

The MIS Platform supports the following RAID levels:

- RAID 0 (striping without parity or mirroring)
- RAID 1 (mirrored without parity or striping)
- RAID 5 (striping with parity)
- RAID 6 (striping with dual parity)
- RAID 00 (spanned drive group striped set from a series of RAID 0 drive groups)
- RAID 10 (mirrored stripes across spanned drive groups)
- RAID 50 (distributed parity and striping across spanned drive groups)
- RAID 60 (distributed parity, with two independent parity blocks per stripe in each RAID set, and disk striping across spanned drive groups)

The onboard MIS server zoning application is supported in both Windows and Linux operating systems. There is also an external fan base CLI, where a system running either CentOS, Linux, or Windows can be used to zone the JBOD in addition to MIS Servers. When LSI MegaRAID HBAs are used, they support RAID 0, RAID 1, RAID 5 and RAID 6, along with their variants. Where ever possible, the zoning and RAID selection for a MIS server should provide for the maximum amount of availability in the event of a StorBrick failure. The StorBrick has four SAS lanes into it from up to four SAS RAID HBAs. These SAS lanes are inputs to a SAS expander that connects to the drives installed in the StorBrick. Since there are eight StorBrick in the MIS-Server, each with nine drives, the highest availability for the data would be RAID groupings that span the StorBricks.

Important: Unless specified, all systems ship as **RAID 6**.

RAID Configuration Notes

To get the best availability, treat each drive brick as an enclosure. For a RAID 1 with only one drive per StorBrick in the group, the loss of a StorBrick does not affect the availability of the data (Figure 3-46).




Figure 3-46 RAID 1 with One Drive per StorBrick

However, with two drives spanning a StorBrick, the loss of a StorBrick will cause the data to be unavailable until the failed storbrick is replaced (Figure 3-47).




Figure 3-47 RAID 1 with Two Drives Spanning a StorBrick

For redundancy, choose one drive from each drive brick as you build the LUNs. Configure the drives in a RAID group to span the StorBricks. For instance, configuring 8 drives from StorBrick 0 as a RAID 5, 7+1 group will work. However, if that StorBrick (StorBrick 0) fails, all 8 drives will become inaccessible, making that RAID group's data unavailable until the StorBrick is replaced.

If, however, drive 0 from each StorBrick (SB0-0, SB1-0, SB2-0, SB3-0, SB4-0, SB5-0, SB6-0 and SB7-0) is used to make up the RAID 5, 7+1 group, and any StorBrick were to fail, only one drive of that group would be affected. The RAID 5 algorithm would be able to recover the data from the remaining drives and the RAID group would remain available. Configurations varies based on needs for capacity versus protection.

There is also the option to assign both a RAID group dedicated spare drive or a global spare drive. For example, a RAID 5 group could be a 6+1 with a dedicated spare drive. In a full up system you would have 9 sets of these RAID groups, all spread across the StorBricks. One method for greater capacity is available by configuring 8 groups of 7+1, one group of 6+1 and a global spare. In a RAID 6 6+2, a spare may not be desirable at all, as one is already part of the group automatically. For configurations that maximize storage, can allow unavailability of data for a time and/or predict a high-success rate for StorBricks, very large RAID groups may be desirable. For typical RAID usage, RAID 5 7+1 or RAID 6 6+2 are most likely.

Warning: Do not configure multiple drives in a RAID group to be on the same StorBrick.

Since there are eight StorBrick in the MIS-Server, each with nine drives, the highest availability for the data would be RAID groupings that span the StorBricks. For example, a RAID 6 configuration should have eight drives, one from each StorBrick. In this configuration, if a StorBrick failed, the data would still be available. For RAID 5 or RAID 6, with one drive per StorBrick in a group, the loss of a StorBrick (SB3) does not affect the availability of the data (Figure 3-48).




Figure 3-48 RAID 5 or 6 with One Drive per StorBrick

If a larger RAID group is desired, it would need to have multiple drives on a StorBrick. Then if a StorBrick were to fail, two drives would be unavailable. If this were a RAID 6 implementation, the data would still be available, though another StorBrick failure or even a drive failure in that group would cause a loss of data availability. If this were a RAID 5 implementation, the data would become unavailable until the failed StorBrick is replaced (Figure 3-49).




Figure 3-49 Loss of a Drive with Multiple Drives on a StorBrick

For RAID 6 with three drives of the group spanning a StorBrick, the data is unavailable until the failed StorBrick is replaced (Figure 3-50).




Figure 3-50 Three Drive Loss in RAID 6 Requires StorBrick Replacement

Spare Drives

The LSI MegaRAID HBA can configure drives to be used as spares. There are two ways of assigning spares. One method is local to a RAID set, be it RAID 1, RAID 3, RAID 5, RAID 6, or any of their variations, in that the assigned spare will only be used within that specific VLUN. The other method is called a global spare, where it can be used by any VLUN that requires one.

The usage of spare drives is an option the customer can select when they configure and create their VLUNs from the physical drives.

For more on RAID configuration, see the MegaRAID guides on the Technical Publications Library (<http://docs.sgi.com>). Operating System software RAID support appears in: Windows Dynamic Disks, or Linux mdadm.

System Maintenance

For warranty and safety considerations, SGI designates the following chassis components as customer-replaceable units (CRUs):

- Power supplies
- Fans
- Disk drives (capacity and boot drives)
- Exterior air ducts
- Front Bezel

These components are all hot-swappable; that is, you can replace them without powering down the platform. Replacement components are issued through SGI Customer Support (see “Product Support” on page xxvi). A trained service technician should install and replace all other components. This chapter describes the steps necessary to replace these CRUs.

- “Detecting Component Failures” on page 132
- “Sliding the Chassis Forward/Backward” on page 133
- “Removing the Front or Rear Chassis Cover” on page 134
- “Replacing a Power Supply” on page 135
- “Replacing a Fan Module” on page 136
- “Replacing a Disk Drive” on page 138
- “Installing the Rack Air Ducts” on page 140

The following tools are required for these replacement procedures:

- #1 and #2 Phillips screwdrivers
- Short #2 Phillips screwdriver
- Flat-head screwdriver

Note: Phillips screwdrivers with magnetic tips are recommended.

Videos of these procedures are available through the SGI Supportfolio web site (<http://support.sgi.com>).



Warning: Review the warnings and precautions listed in, “Important Information” on page xix, before setting up or servicing this chassis.

Detecting Component Failures

In general, when a system component fails, the BMC sends an alert (when configured to do so; see “Alert Email” on page 51 for more information). Drive failures are recognized and reported through MegaRAID Storage Manager (MSM). Either the BMC or MSM generates an alert to the monitoring application for your storage server. The alerts include the system serial number, the suspect component, and a summary of the fault. Drives, power supplies and fans can be replaced using the following procedures listed in this chapter. For all other components, you should inform SGI service of the fault and forward the information from the alerts.

Some component failures are also registered on the Fan Base Service Page. The Fan Base Service Page can be accessed by connecting through the MIS-S9D Proprietary Ethernet port (Figure 4-1) and pointing a browser to the fan base IP: `10.4.3.196/service.xml` for servers (Figure 4-5) and `10.4.3.214/service.xml` for JBODs. Using the Fan Base Service Page, failed fans can be set in a “Safe to Service” mode, a step necessary before changing out a fan module.



Figure 4-1 MIS-S9D Proprietary Network Interface

Sliding the Chassis Forward/Backward

The newly designed rail kit of the MIS chassis allows it to be slid forward or backwards (20"). The new rail design is such that one person can now move the rack. You will need to slide the chassis out to service some of its components. To slide the chassis out in either direction, follow these steps:

1. Push the two release latches in, at the left and right sides and in the center of the rail mounts, towards the center of the chassis.
2. Pull the chassis out using the handles. The chassis will latch at the 20-inch limit.
3. To slide the chassis back in, depress the two release latches near the rail and slide it back in.

Removing the Front or Rear Chassis Cover




Figure 4-2 Front & Rear Chassis Covers

Important: When a chassis cover is removed, an intrusion sensor monitored by the BMC will detect its removal. If the cover is off for more than 15 minutes or any system temperature sensor exceeds its threshold limit, the server will perform an orderly shutdown and power-off.

As shown in Figure 4-2, the top of the chassis is bifurcated; that is, there is a front and rear chassis cover. Except for power supply maintenance, all service actions require that you remove the front or rear chassis cover, detailed below.

1. To remove a chassis cover, first follow the instructions in “Sliding the Chassis Forward/Backward” on page 133.
2. Remove the single security screw from the cover.
3. Push the detent, and slide the cover out and up from the chassis.

Replacing a Power Supply




Figure 4-3 Power Supply Module

1. Using the BMC Web Interface, verify the fault (failed unit) and turn on the locator LED (blue) for that chassis.
2. Locate the failed unit in the specified chassis: its amber service LED should be lit, indicating a fault (see Figure 4-3).
3. Unplug the power supply that will be replaced and remove its power cord.
4. Push the release tab on the back of the power supply.
5. Pull the power supply out using the handle.
6. Replace the failed power module with another of the same model.
7. Push the new power supply module into the power bay until it clicks into the locked position.
8. Plug the AC power cord back into the module, and plug the module in.
9. If the power supply is being hot-swapped on a system that is already powered on, the power supply will power up automatically. If the power supply is being installed in a system that is not powered on, the system should be powered on at this time.
10. You will have to wait a few moments for the power supply to respond to AC power, and complete its internal processes, before its status LEDs are illuminated. Once lit, ensure that it is also recognized as **Good** in the BMC.

Replacing a Fan Module




Figure 4-4 Fan Module

Warning: If the fan is NOT set in a “**Safe to Service**” mode, there are only **10 seconds** available to exchange a fan. Without setting “**Safe to Service**,” the platform will assume a fan has experienced *critical failure*, requiring an emergency shutdown for thermal requirements. Setting “**Safe to Service**” allows more time for replacement procedures (15 minutes) before initiating thermal shutdown.

1. Using the BMC Web Interface, verify the fault (failed module) and turn on the locator LED (blue) for that chassis.
2. Locate that chassis and use an Ethernet cable to connect a server/laptop to the MIS-S9D Proprietary Network Interface (Figure 4-1).
3. Access the Fan Base Service Page by pointing a browser to <http://192.168.0.196/service.xml> (for servers) or <http://192.168.0.214/service.xml> (for JBODs) from the connected server/laptop. This will bring up the Fan Base Service Page (Figure 4-5).

MIS Server Chassis												
Cover is on, uptime 16868006(ms mod 2 ³²)						All FANS <input type="button" value="On"/>			<input type="button" value="Shutdown"/>			
R E A R	D8	D5	D2	StorBrick 5 Good	Boot Drive CM0-1	FAN 5 Good <input type="button" value="On"/> <input type="button" value="Off"/> RPM : 12273 : 11650	StorBrick 4 Good	D0	D3	D6	F R O N T	
	D7	D4	D1		Boot Drive CM0-0			D1	D4	D7		
	D6	D3	D0		Boot Drive CM1-1			D2	D5	D8		
	D8	D5	D2	StorBrick 6 Good	Boot Drive CM1-0	FAN 4 Good <input type="button" value="On"/> <input type="button" value="Off"/> RPM : 11977 : 10704	StorBrick 3 Good	D0	D3	D6		
	D7	D4	D1		D1			D4	D7			
	D6	D3	D0		D2			D5	D8			
	D8	D5	D2	StorBrick 7 Good		FAN 3 Good <input type="button" value="On"/> <input type="button" value="Off"/> RPM : 10836 : 12334	StorBrick 2 Good	D0	D3	D6		
	D7	D4	D1		D1			D4	D7			
	D6	D3	D0		D2			D5	D8			
	Compute Module 1 -						FAN 2 Good <input type="button" value="On"/> <input type="button" value="Off"/> RPM : 11753 : 11021	StorBrick 1 Good	D0	D3		D6
	Compute Module 0 CM0 P1 59.0°C Ok - CM0 P2 55.0°C Ok								D1	D4		D7
	Power Supply 3 Good						FAN 1 Good <input type="button" value="On"/> <input type="button" value="Off"/> RPM : 11384 : 11014	StorBrick 0 Good	D2	D5		D8
Power Supply 2 Good						D0			D3	D6		
Power Supply 1 o						FAN 0 Good <input type="button" value="On"/> <input type="button" value="Off"/> RPM : 10987 : 10763	StorBrick 0 Good	D1	D4	D7		
Power Supply 0 o								D2	D5	D8		

Figure 4-5 Fan Base Service Page

- From the Fan Base Service Page, locate the faulted fan, and click the Off button.
- At the prompt, **Are you sure you want to turn OFF FAN #?** (where # is the number of the fan), click **Ok**. This put the fan in a “**Safe to Service**” mode and will illuminate its blue locator LED.
- Remove the front chassis cover (see “Removing the Front or Rear Chassis Cover” on page 134) and locate the fan module with the illuminated blue LED (Figure 4-4).
- Loosen the thumbscrew, pull out the faulted fan by pulling upward on both the front and rear flanges, and replace it.
- Once the fan module is replaced, seat the fan by pushing between the two LEDs until it seats.
- Re-install the chassis cover and security screw.

10. Using the fan base Service page, return the fan to On/Full Power. At the prompt, **Are you sure you want to turn ON FAN #?** (where # is the number of the fan), click **Ok**. If there are no faults, the fan will power up, report status of **Good**, and its green LED will be lit. This will also turn off the blue locator LED.
11. Unlock the chassis from the extended position and push it back until it locks into the stowed position.
12. Using the BMC, clear the chassis locator LED.

Replacing a Disk Drive



Important: Empty drive carriers cannot be inserted into the storbricks, so slots without HDDs will not have carriers.

To replace a failed disk drive:

1. Using the MegaRAID Storage Manager for your system, verify the fault (failed unit).
2. Using the MegaRAID Storage Manager, set the system to a service state for the removal of the faulted drive. The MegaRAID Storage Manager will turn off the drive. It will then turn on the fault LED (yellow) for that drive.
3. Remove the chassis cover. (See “Removing the Front or Rear Chassis Cover” on page 134.)
4. Locate the faulted drive with the illuminated yellow LED, and remove it from its StorBrick (or boot drive bay). (See “Removing the Drive” on page 139.)
5. Replace the faulted drive. (See “Re-installing the Drive” on page 139.)
6. Once the drive is replaced, re-install the chassis cover and security screw.
7. Unlock the chassis from the extended position and push it back until it locks into the stowed position.
8. Using the MegaRAID Storage Manager, return the system to a normal state. The new drive will be powered on.
9. If there are no faults at this time, the rebuild or mirroring of the data to the new drive will begin.

Removing the Drive

As shown in Figure 4-6, the drives are mounted in drive carriers to simplify their installation and removal from the drive bricks or boot drive bays in the chassis.

To remove the drive, perform the following steps:

1. Ensure that the drive LEDs are off (except the blue locator LED), indicating that the drive is not in use and can be removed.
2. Unlatch the drive carrier by sliding the grey latch toward the drive and pull the drive carrier out of the StorBrick or boot drive bay.
3. Remove the four screws that secure the drive to the drive carrier.
4. Remove the drive from the carrier.




Figure 4-6 Hard Drive Carrier

Re-installing the Drive

To re-install a hard drive into the hard drive carrier, perform the following steps:

1. Place the hard drive carrier on a flat, stable surface such as a desk, table, or work bench.
2. Slide the hard drive into the carrier with the printed circuit board side facing down.
3. Carefully align the mounting holes in the hard drive and the carrier.
Make sure the bottom of the hard drive and bottom of the hard drive carrier are flush.
4. Secure the hard drive using the four screws (see Figure 4-6).
5. Replace the drive carrier into the chassis.
6. Push the drive carrier down to lock it place.
7. Replace the chassis cover and slide the chassis back into place in the rack.
8. Using the MegaRAID Storage Manager (see Chapter 3, “System Software”), add the drive to the system zoning configuration. Once properly zoned, the drive will begin the rebuild process.

Installing the Rack Air Ducts

There are two air ducts that install in the rack above either outer rail for the MIS chassis. In order to install these air ducts, you will need a #2 Phillips screw driver. (Some air duct kits do not attach with a interior screw, but rather with Velcro.)

1. Push the chassis backwards into the server rack, until it locks into place. This will give you access to the inside of the rack, where the air ducts will be installed.
2. Place the air duct above the outer rail, and secure it to the outer rail and rack using the two screws provided.
3. There is a smaller flat-head screw that goes on the inside in the counter-sunk hole. Make sure the screw to the outer rail is sunk deep enough so that it does not catch or scratch the chassis as it returns to its normal position in the rack.
4. The larger truss head screw attaches to the back, where it is large enough to cover the wide, square hole.

The second air duct installs on the other outer rail, in mirror image to the first.




Figure 4-7 Rack Air Duct Placement

Replacing the Front Bezel Grille

The MIS front Bezel Grille can be detached from the front of the chassis by gently tugging on the right side. There are two or three cables (one or two for control panel functionality, and one for the S9D network interface) at the front left corner of the chassis (Figure 4-9). These cables must be unattached from their corresponding sockets on the front bezel grille, before removing the grille from the chassis. Once detached, the new front grille can be installed. To do so, follow these instructions.

1. Orient the front bezel grille so that the left inner edge containing the sockets (Figure 4-8) lines up with the front left edge of the chassis where the cables protrude.
2. Use the two latches on the bezel grille to help attach it to the chassis on the left side, and work from the right side (swinging it open like a door, rather than top or bottom, for ease of cabling).




Figure 4-8 Front Bezel Grille Control Panel Sockets and S9D Network Interface



Figure 4-9 Chassis Front Cables

3. Plug the ribbon cable(s) into their corresponding control panel socket (Figure 4-10).



Figure 4-10 Control Panel Ribbon Cables Attached to Front Bezel

4. Plug the network cable into the underside of the S9D network interface (Figure 4-11s).



Figure 4-11 S9D Network Interface Cable

5. Once all the wires have been successfully connected, swing the grille shut such that it can be snapped into place on the front of the chassis using the attached tabs (Figure 4-12). Push firmly but gently to attach the front grille.




Figure 4-12 Front Bezel Grille Attach Tabs

Troubleshooting

This chapter describes troubleshooting for the problems listed below, that customers are able to perform without the assistance of SGI Field Technicians. Chapter 2, “System Interfaces,” describes use of the control panel to monitor the overall system status and the status of specific components. Chapter 4, “System Maintenance,” describes how to replace defective components.

- “No Video” on page 146
- “Losing the System’s Setup Configuration” on page 146
- “Safe Power-Off” on page 146
- ““Missing” Firmware Files” on page 147
- “Cannot Receive Email Alerts Using BMC” on page 148
- “Path not found errors after installing Python for Windows” on page 150
- “Intel BIOS POST error messages and handling description” on page 150
- “Beep Codes” on page 154
- “Not all beep codes signal an error” on page 155
- “Fans don’t spin when power button pressed” on page 155
- “Fans spin when power button is pressed, but no video display is seen” on page 155
- “BIOS or board logo appear, but OS load screen never appears” on page 155
- “OS load screen appears, but nothing further” on page 155
- “OS fully loads, but errors are seen” on page 156
- “Cannot boot after new software or drivers were installed” on page 156
- “Cannot boot after settings were changed in BIOS” on page 156
- “System Gets Stuck in a Reboot Cycle” on page 156
- “Shortening the Cable Management Arm” on page 156

Other troubleshooting information is contained in “Related Publications” on page xxv (such as, [*Intel® Remote Management Module 4 and Integrated BMC Web Console User Guide*](#) (pdf)).

Important: If you require software to run or monitor the MIS platform, need technical assistance, or to contact SGI Customer Support, please see “Product Support” on page xxvi.

No Video

If the power is on but there is no video, remove all add-on cables (besides power & video). Use a different monitor to determine if the fault lies with it. Use the speaker to determine if any beep codes exist. Refer to “Beep Codes” on page 154, and Appendix B, “BIOS Error Codes” for details.

Losing the System’s Setup Configuration

Make sure there is no fault with the power supplies. A fault with the power supply may cause the system to lose the CMOS setup information. If this does not fix the Setup Configuration problem, contact SGI for repairs.

Safe Power-Off

There are different ways to shut down an MIS machine, each more extreme than the next. The most polite way is to go into the machine’s operating system and select **Shut Down**. This will prompt the user to enter a password before allowing the shut-down process. They include,

- Using the OS GUI power-off button at the console screen, if a keyboard/mouse/video monitor is connected.
- When logged in via an ssh session and executing a “shutdown” or “poweroff” command.
- When logged in to the BMC and using the power control page to power off the sever.
- Using the remote console screen GUI power-off button, if a KVM RMM4Lite session is established through the BMC.

- Connecting to the fan base using the MIS-S9D proprietary network interface and executing a “shutdown” or “poweroff” command.

If the platform is an MIS dual-server and **both** servers are powered up, performing the above steps **only** powers off the server with which you are working. The fans, drives and second server will remain powered on until the second server is powered off, then all power (but standby) will be turned off.

For a JBOD Unit, the power button on the front panel will turn off the power to that I/O module. If a second module is installed and powered on, it, the fans and the drives will remain on until it, too, is powered off.

The next way shut down the platform is to hold down the **Power** button on the front of the unit (Figure 2-1) until the machine powers off. However, this is not a polite way to power off the machine, and will require a recovery process at start-up.

If power is lost from outside the machine (power outage), the machine will recognize the loss of power and execute an emergency shutdown procedure. If there is a battery back-up unit installed, it will protect integrity of cache in the event of power or server failure.

“Missing” Firmware Files

There may come a time when you will need to update the zoning tool in order to support the firmware that is present on the StorBricks in a new MIS Server or JBOD machine. If this is the case, you will receive an error message that states either: **Unable to find directory that contains firmware files** or **Unable to find firmware file**, and below that, the version number of the firmware files necessary (example: **01.30.00.00**).

The firmware folders necessary for download will contain the version number in the folder name. For this example, the desired firmware folders will have a name similar to `storbrick-release-01.30.00.00 <date>`.

You will need to download the necessary firmware files and place them into the `/opt/Zones/Versions` folder for Linux and the `c:\Program Files (x86)\Zones\Versions` folder for Windows.

Cannot Receive Email Alerts Using BMC

In order to receive email alerts, the MIS Server must be networked to a SMTP mail server. Next, using the BMC, go to the **Configuration** page, and click on **Alert Email** (Figure 5-1).

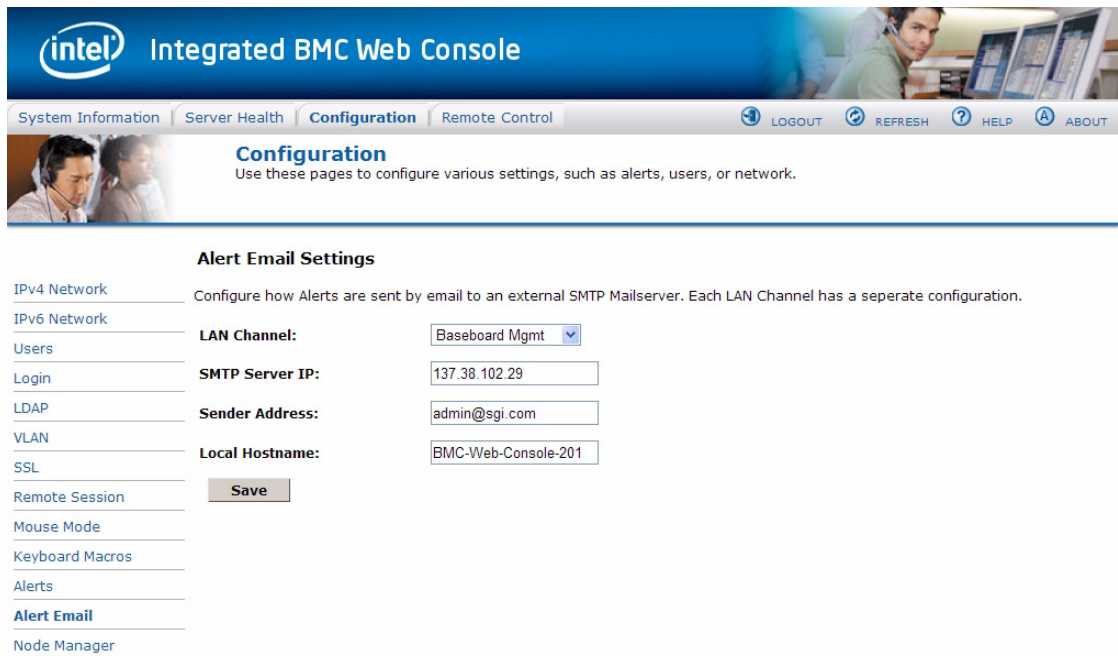


Figure 5-1 BMC Alerts Email Settings Page

The **LAN Channel** should be set to **Baseboard Mgmt**. The **SMTP Server IP** should be the IP address of the mail server to which the MIS Platform is connected. **Sender Address** can be any email address configured on the SMTP server. **Local Hostname** can be any name to identify the machine sending the alert (limit of 18 characters). When finished, click the **Save** button at the bottom of the page. A pop-up will let you know that the changes have been saved (Figure 5-2).

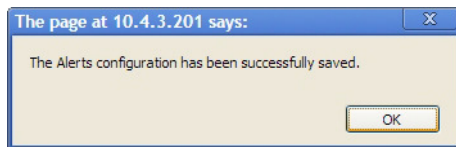


Figure 5-2 BMC Alert Configuration Success Pop-up

Next, click on **Alerts** in the left-hand navigation pane (Figure 5-3).

The screenshot shows the Intel Integrated BMC Web Console interface. The top navigation bar includes 'System Information', 'Server Health', 'Configuration' (selected), and 'Remote Control'. Below the navigation bar, the 'Configuration' section is active, with a sub-header 'Alerts' and a description: 'Use these pages to configure various settings, such as alerts, users, or network.' The left-hand navigation pane lists various system settings, with 'Alerts' highlighted. The main content area for 'Alerts' includes:

- Alerts**: Configure which system events generate Alerts and the external network destinations they should be sent to.
- IPv4 Network**, **IPv6 Network**, **Users**, **Login**, **LDAP**, **VLAN**, **SSL**, **Remote Session**, **Mouse Mode**, **Keyboard Macros**
- Globally Enable Platform Event Filtering:** Enabled Disabled
- Select the events that will trigger alerts:**
 - Temperature Sensor Out of Range
 - System Restart
 - Fan Failure
 - Power Supply Failure
 - BIOS: Post Error Code
 - Node Manager Exception
 - Watchdog Timer
 - Voltage Sensor Out of Range
 - Chassis Intrusion
 - Memory Error
 - FRB Failure
 - Hard Drive Failure
- Check All** **Clear All**
- LAN Channel to Configure:** Baseboard Mgmt
- Alert Destination #1:**
 - SNMP Send SNMP Alerts to IP: 0.0.0.0
 - Email Send Email to: field_technician@sgi.com
- Alert Destination #2:**
 - SNMP Send SNMP Alerts to IP: 0.116.116.112
 - Email Send Email to: [empty field]
- Save** **Send Test Alerts**

Figure 5-3 BMC Alerts Configuration Page

Make sure all the checkboxes for alerts contain a check for each alert to be sent. The **LAN Channel to Configure** should be set to **Baseboard Mgmt**. **Alert Destination #1** should be set to **Email**, and **Send Email to:** should be the email address of the account to receive the alerts. When finished, click the **Save** button at the bottom of the page. A pop-up will let you know that the changes have been saved. To test your configuration, click the **Send Test Alerts** button. You will receive a pop-up letting you know that the email was sent, and you will receive an email notifying you of a successful test.

Path not found errors after installing Python for Windows

Some Windows machines will need the variable path set for Python, and will display “Path not found” errors if not set. The path is set by doing the following:

1. Go to **Control Panel** and double-click the **System** icon.
2. Click on the **Advanced** tab.
3. Click the **Environment Variables** button at the bottom.
4. In the system variables box in at the bottom, find the entry `Path` and select it.
5. Click the **Edit** button and add to the end of the string: `;c:\python##` where `##` is the version number of Python installed on the Windows machine.
6. Click **Apply** and **Okay**.

Python will now run normally.

Intel BIOS POST error messages and handling description

Refer to this table if you receive a code during the POST of booting my Intel® Server Board. What does it mean?

See table for listing of codes for BIOS events relevant to Intel® Server Boards.

Table 5-1 Intel BIOS POST Error Messages

Error Code	Error Message Response	Event Type
004C	Keyboard / interface error	Major
0012	CMIS date / time not set	Major
0048	Password check failed	Fatal
0141	PCI resource conflict	Major
0146	Insufficient memory to shadow PCI ROM	Major
0192	L# cache size mismatch	Fatal
0194	CPUID processor family are different	Fatal

Table 5-1 Intel BIOS POST Error Messages (**continued**)

Error Code	Error Message Response	Event Type
0195	Front side bus mismatch	Major
5220	Configuration cleared by jumper	Minor
5221	Passwords cleared by jumper	Major
8110	Processor 01 internal error (IERR) on last boot	Major
8111	Processor 02 internal error (IERR) on last boot	Major
8120	Processor 01 thermal trip error on last boot	Major
8121	Processor 02 thermal trip error on last boot	Major
8130	Processor 01 disabled	Major
8131	Processor 02 disabled	Major
8160	Processor 01 unable to apply BIOS update	Major
8161	Processor 02 unable to apply BIOS update	Major
8190	Watchdog timer failed on last boot	Major
8198	Operating system boot watchdog timer expired on last boot	Major
8300	Baseboard management controller failed self-test	Major
8305	Hot swap controller failed	Major
84F2	Baseboard Management Controller failed to respond	Major
84F3	Baseboard management controller in update mode	Major
84F4	Sensor data record empty	Major
84FF	System event log full	Minor
8500	Memory component could not be configured in the selected RAS mode	Major
8520	DIMM_A1 failed Self Test (BIST)	Major
8521	DIMM_A2 failed Self Test (BIST)	Major
8522	DIMM_A3 failed Self Test (BIST)	Major

Table 5-1 Intel BIOS POST Error Messages (continued)

Error Code	Error Message Response	Event Type
8523	DIMM_A4 failed Self Test (BIST)	Major
8524	DIMM_B1 failed Self Test (BIST)	Major
8525	DIMM_B2 failed Self Test (BIST)	Major
8526	DIMM_B3 failed Self Test (BIST)	Major
8527	DIMM_B4 failed Self Test (BIST)	Major
8528	DIMM_C1 failed Self Test (BIST)	Major
8529	DIMM_C2 failed Self Test (BIST)	Major
852A	DIMM_C3 failed Self Test (BIST)	Major
852B	DIMM_C4 failed Self Test (BIST)	Major
852C	DIMM_D1 failed Self Test (BIST)	Major
852D	DIMM_D2 failed Self Test (BIST)	Major
852E	DIMM_D3 failed Self Test (BIST)	Major
852F	DIMM_D4 failed Self Test (BIST)	Major
8580	DIMM_A1 Correctable ECC error encountered	Minor/Major after 10 events
8581	DIMM_A2 Correctable ECC error encountered	Minor/Major after 10 events
8582	DIMM_A3 Correctable ECC error encountered	Minor/Major after 10 events
8583	DIMM_A4 Correctable ECC error encountered	Minor/Major after 10 events
8584	DIMM_B1 Correctable ECC error encountered	Minor/Major after 10 events
8585	DIMM_B2 Correctable ECC error encountered	Minor/Major after 10 events
8586	DIMM_B3 Correctable ECC error encountered	Minor/Major after 10 events

Table 5-1 Intel BIOS POST Error Messages (**continued**)

Error Code	Error Message Response	Event Type
8587	DIMM_B4 Correctable ECC error encountered	Minor/Major after 10 events
8588	DIMM_C1 Correctable ECC error encountered	Minor/Major after 10 events
8589	DIMM_C2 Correctable ECC error encountered	Minor/Major after 10 events
858A	DIMM_C3 Correctable ECC error encountered	Minor/Major after 10 events
858B	DIMM_C4 Correctable ECC error encountered	Minor/Major after 10 events
858C	DIMM_D1 Correctable ECC error encountered	Minor/Major after 10 events
858D	DIMM_D2 Correctable ECC error encountered	Minor/Major after 10 events
858E	DIMM_D3 Correctable ECC error encountered	Minor/Major after 10 events
858F	DIMM_D4 Correctable ECC error encountered	Minor/Major after 10 events
8601	Override jumper set to force boot from lower alternate BIOS bank of flash ROM	Minor
8602	WatchDog timer expired (secondary BIOS might be bad)	Minor
8603	Secondary BIOS checksum fail	Minor
92A3	Serial port component was not detected	Major
92A9	Serial port component encountered a resource conflict error	Major
0xA000	TMP device not detected	Minor
0xA001	TMP device missing or not responding	Minor

Table 5-1 Intel BIOS POST Error Messages (**continued**)

Error Code	Error Message Response	Event Type
0xA002	TMP device failure	Minor
0xA003	TMP device failed self test	Minor

Beep Codes

The BMC may generate beep codes upon detection of failure conditions. Beep codes are sounded each time the problem is discovered (for example, on each power-up attempt), but are not sounded continuously. Common supported codes are listed in Appendix B, “BIOS Error Codes” for details.

Another source of beep codes may be the LSI MegaRAID card. Table 5-2 contains a summary of the LSI MegaRAID card beep codes. These beep codes indicate activity and changes from the optimal state of your RAID array.

Table 5-2 Summary of LSI MegaRAID Card Beep Codes

Beep Code	LSI Firmware State	Cause (Depending on RAID Level)
3 seconds on, 1 second off	SPEAKER_OFFLINE_ENTRY	RAID 0: One or more drives offline. RAID 1: Two drives offline. RAID 5: Two or more drives offline. RAID 6: More than two drives offline.
1 second on, 1 second off	SPEAKER_DEGRADED_ENTRY	RAID 1: A mirrored drive failed. RAID 5: One drive failed. RAID 6: One or two drives failed.
1 second on, 3 seconds off	SPEAKER_HOTSPARE_ENTRY	A hot spare drive has completed the rebuild process and has been brought into the array.

Not all beep codes signal an error

Intel boards and systems are designed to indicate USB readiness by a series of beep codes during POST, before video becomes available. These beeps mean USB is powered and initialized.

Device such as a pen drives or USB CD/DVD ROM drives attached to external USB port will generate a beep once the device is recognized, powered and initialized.

These beep codes do not signal any error. They signal USB and external device readiness during POST.

Fans don't spin when power button pressed

Is at least one power supply fan spinning? If yes, there is good power to the modules. Verify all required power cables are correctly plugged into the chassis.

If no, there is potential lack of clean power to the module. Swap power cables, try different wall circuit, or try replacing a power supply module (see “Replacing a Power Supply” on page 135).

Fans spin when power button is pressed, but no video display is seen

Are there any beeps? If **yes**, beep codes are listed in Table 5-2 on page 154. If the answer is **no**, there may be a memory card error or with the processor. For either case, contact SGI Technical Support (see “Product Support” on page xxvi).

BIOS or board logo appear, but OS load screen never appears

This may be a problem with add-in cards in the server system. Contact SGI Technical Support (see “Product Support” on page xxvi).

OS load screen appears, but nothing further

This may be a problem with add-in cards in the server system. Contact SGI Technical Support (see “Product Support” on page xxvi).

OS fully loads, but errors are seen

Use operating system logging utility such as Windows Event Viewer or Linux dmesg to narrow the source of the error, and contact SGI Technical Support (see “Product Support” on page xxvi).

Cannot boot after new software or drivers were installed

If you recently installed new software or new device drivers, try booting into Safe Mode and uninstall the new software or driver. If you can now boot normally, there may be a compatibility issue between the new software or driver and some component in your system. Contact the software manufacturer for assistance.

Cannot boot after settings were changed in BIOS

Certain changes in BIOS settings (such as chipset timing or latency, memory timing or latency, processor clock frequency, etc.) can cause a system to no longer boot. If you are able to enter the BIOS Setup by pressing F2, reset the BIOS to factory defaults by pressing F9. Save and exit the BIOS Setup. If you cannot enter the BIOS Setup, contact SGI Technical Support (see “Product Support” on page xxvi).

System Gets Stuck in a Reboot Cycle

If the platform gets stuck in a cycle where it is continuously rebooting without ever going into full power on, try replacing two of the power supplies, even if they are not reporting faults.

If a power supply has failed and does not correctly issue its fault (i.e., it *falsely* reports that it is still working), it will cause the platform to go into standby power mode, and the platform will try rebooting in order to power up. This can cause the platform to get stuck in a reboot cycle, and never actually successfully powering on. Replacing the faulty power supplies resolves this issue.

Shortening the Cable Management Arm

The cable management arm is an optional addition to the MIS chassis that helps keep cables organized in the rack. It attaches to a plate on the rear of the MIS chassis and to the rack in which

the chassis is mounted. The cable management arm comes in a length designed for 26" deep racks, such as the D-Rack. It can be shortened to accommodate other rack sizes (Figure 5-4).




Figure 5-4 Cable Management Arm

Removing segments from the cable management arm is done using a thin flat-head screwdriver.

- For a 26" deep rack, do not remove any sections.
- For a 30" deep rack, remove 3 sections of the arm.
- For a 34" deep rack, remove 6 sections of the arm.

There are two ways that segments can be removed: either while the cable management arm is not attached to the chassis or rack, and devoid of cables, or while attached to the rack and chassis, with cables within.

For the first method, segments are removed from the *end* of the arm. Use the flat-head screwdriver and slip it between the plastic pieces that form the joint of the arm (Figure 5-5).



Figure 5-5 Remove Segments from the End of an Empty Arm

Pry the joint apart on one side (Figure 5-6).



Figure 5-6 Pry the Joint Apart

Pull the joint open (Figure 5-6).



Figure 5-7 Pull the Segments from the Arm

The segments are now removed from the end of the arm (Figure 5-7), and the arm can now be attached to the chassis and rack (Figure 5-8).



Figure 5-8 Shortened Arm Ready for Installation

The second way to shorten the cable management arm is while it is already installed on the rack, with cables in it. In this case, the segments are removed from the *middle* of the cable management arm (Figure 5-9).




Figure 5-9 Remove Segments from the Middle of a Full Arm

As before, using a thin flat-head screwdriver, pry apart the joint of a segment on one side (Figure 5-10).



Figure 5-10 Pry Open the Segment Joint

Do this to the other side, and then pull apart the arm (Figure 5-11).



Figure 5-11 Pull Apart the Cable Management Arm

Repeat the process on the other end of the segment you wish to remove.



Figure 5-12 Pry Apart the Joint at the Other End

Once complete, separate the segment from the cables by twisting the segment off (Figure 5-13).



Figure 5-13 Removing the Segment from the Cables

Bring the two spliced ends of the cable management arm together, and snap the joint together (Figure 5-14).



Figure 5-14 Join the Spliced Ends of the Cable Management Arm

The cable arm should now be the appropriate length for the depth of the rack.

Technical Specifications

Table A-1 describes the technical specifications for the SGI MIS platform.

Table A-1 Technical Specifications

Attribute	Specification
Overview	
Profile	4U Standard-depth chassis
Product type	SGI MIS Server Platform (single or dual server), or SGI MIS JBOD Unit
Connectivity	Up to four SGI MIS JBOD units per SGI MIS Server Platform
Mount	<ul style="list-style-type: none"> –Standard 19-inch rack-compatible rail kit (weight-dependent) –SGI 19-inch Destination Rack (D-Rack), 42U –Up to 10 chassis per D-Rack
Chassis Dimensions	
Height	6.94" (176 mm)
Width	16.9" (429.2 mm)
Depth	36" (914.4 mm)
Max weight	220 lbs.
Power	
AC Input	100-240 VAC (50-60Hz) single phase
Safety	<ul style="list-style-type: none"> –NRTL certified to UL60950-1 –CB certified to EN60950/IEC60950

Table A-1 Technical Specifications (continued)

Attribute	Specification
EMC	<ul style="list-style-type: none"> –North America FCC Part 15 Class A –Canada ICES-003 –Europe EN55022/EN55024 –Korea KN22, KN24 –Taiwan CNS 13438
Operating Environment	
Operating temperature range	41° to 95° F [5° to 35° C] (up to 5,000 ft. [1,500 m]) <ul style="list-style-type: none"> – Derate max temperature (95° F [35° C]) by 1.8° Fahrenheit [1° Celsius] per 1,000 ft. [305 m] of altitude above 5,000 ft. [1525 m] – Temperature rate of change must not exceed 18° F [10° C] per hour
Operating humidity range	8% to 80% non-condensing <ul style="list-style-type: none"> – Humidity rate of change must not exceed 10% relative humidity per hour
Operating altitude	Up to 10,000 ft. [3,050 m]
Shipping temperature range	minus 40° to 140° F [minus 40° to 60° C]
Shipping Humidity	10% to 95% non-condensing
Shipping Altitude	Up to 40,000 ft. [12,200 m]
Storage Temperature	41° F [5° C] to 113° F [45° C] <ul style="list-style-type: none"> – Temperature rate of change must not exceed 36° F [20° C] per hour
Storage Humidity	8% to 80% non-condensing
Storage Altitude	up to 40,000 ft. [up to 12,200 m]
Sgi MIS Server Specifications	
Servers/System	<ul style="list-style-type: none"> –One or two server modules per system –Dual-socket processors per server
Processor support	<ul style="list-style-type: none"> –Supports Intel® Xeon® E5-2600 series processors –Supports Intel Turbo Boost Technology 2.0

Table A-1 Technical Specifications (continued)

Attribute	Specification
Max cores	16 per server (32 per system)
Memory	<ul style="list-style-type: none"> -Up to 8 DDR3 DIMMs (8, 16, 32 GB) for a single server configuration -Up to 16 DIMMs for a dual server configuration. -Max 256GB per server module
Boot drives	<ul style="list-style-type: none"> -Two per server module, mirrored using RAID 1 -3.5" or 2.5" (15mm or 9.5mm) -SAS or SATA -Rotational or SSD -Up to 300 GB, 6Gb/s, 15,000rpm
Supported operating systems	<ul style="list-style-type: none"> -CentOS 6.3 -RHEL 6.2 -SLES 11 SP1 -VMware ESXi 5.0 -Windows 2008 R2 SP1
Networking	Up to four user-specified PCIe HBAs, full-height (4.25") and half-depth (3.375"), externally or internally facing.
Expansion slots	<ul style="list-style-type: none"> Single server: 6 x PCIe gen 2 Dual server: 4 x PCIe gen 2 per server (8 total)
RAID controllers	<ul style="list-style-type: none"> -8 to 32 SAS ports via 8-port PCIe cards -Support for RAID 0, 1, 5, 6, 00, 10, 50, and 60
External storage attachment	Up to 4 SGI MIS JBOD Units per server module via PCIe cards
Internal storage	<ul style="list-style-type: none"> -Up to 72 SAS or SATA 15mm, 2.5" or 3.5" drives -Up to 144 SAS or SATA 9.5mm, 2.5" drives -Drive size and capacity can be mixed in groups of 8 -Supported internal drives: SAS or SATA, rotational or SSD
SGI MIS JBOD Specifications	
Internal Storage	<ul style="list-style-type: none"> -Up to 81 SAS or SATA 15mm, 2.5" or 3.5" drives -Up to 162 SAS or SATA 9.5mm, 2.5" drives -Drive type and size can be mixed in groups of 9 -Supported internal drives: SAS or SATA, rotational or SSD
Connectivity	8 or optional 16 SAS ports

Zone Permission Groups Rules

The first 8 bits are the Master Group. The first bit (0) has access to bit 1 and nothing else. This is why the Master Group 0 is always set to 0. Master Group 1 is all FF—it “sees” everyone. Master Groups 2 & 3 talk to all the Initiators (bits 8-15), and currently, Master Groups 4-7 point up to Master Group 1 (this may change in the future).

The next 8 bits are the Initiator Groups. Bit 8 through 11 corresponds to Adapter 1, Paths 0 through 3 and bits 12-15 correspond to Adapter 2, Paths 4 through 7.

The next 16 through 24 bits are the Zone Permission Groups, with the 25th bit always set for the unassigned slot. The remaining bits are reserved for future use. See Figure B-1 and Figure B-2 for examples.

