



SGI® Modular InfiniteStorage™ (MIS) Platform
User Guide, version 1.0

007-5818-005

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	June 2012 Original printing.
002	August 2012 Edited and updated for technical and editorial changes. Images updated to reflect changes, added Zones and CLI Zoning Tool software information, available RAID configurations updated.
003	October 2012 Electromagnetic Compatibility (EMC) compliance and safety information included. Section on CPU/Riser/HBA configuration restrictions and options added. Weight safety maximums included. Additional information on RAID options and conditions provided. Includes updates from the stand-alone version of Chapter 3 issued in September 2012. Information on grouping control added per customer feedback (Appendix C).
004	July 2013 Updated to include new Zones 1.4.2 and the CLI Zoning tool version 1.4. Troubleshooting section expanded. Added Environmental Requirements.
005	August 2013 Added information on the latest firmware, which includes the Fan Base Service page. Added server module LED codes. Updated information on power supply, fan, boot and capacity drive monitoring and replacement procedures. Added BMC IP configuration instructions, along with adding user accounts, and email alert configuration instructions. Expanded information on IPMI. Added table to troubleshooting section detailing which side to access for which components.

:

Contents

Introduction	xix
Audience	xix
Important Information	xix
Environmental Requirements	xx
Safety Precautions	xxi
ESD Precautions	xxii
Safety & Emissions	xxii
Electromagnetic Compatibility	xxiii
Safety Certification	xxiii
Chapter Descriptions	xxiv
Related Publications	xxv
Conventions	xxvi
Product Support	xxvi
CRU/FRU	xxvii
Purchasable Support & Maintenance Programs	xxvii
Reader Comments	xxvii
1 System Overview	1
MIS Enclosure	5
Front Grille and Control Panels	6
Rear Panel Components	7
Cable Management Arm	9
MIS Common Modules	10
Power Supply Module	11
Fan Assembly Module	13
StorBrick Module	13
MIS Server Platform or JBOD Unit	16
Server Module	16

	Layout of Server CPUs, and PCIe Risers HBAs 22
	Boot Drives Module 24
	MIS JBOD I/O Module 25
	System Block Diagram 25
	System Layout 27
2	System Interfaces 29
	Control Panel 30
	MIS JBOD vs. Server Control Panel 32
	Powering an MIS Platform 33
	Using IPMITool 34
	Turning Off an MIS Platform 34
	Server Module LEDs 35
	JBOD LEDs 38
	Disk Drive LEDs 38
	Drive Monitoring 39
	Boot Drive Monitoring 39
	Capacity Drive Monitoring 42
	Power Supply LEDs 45
	Power Supply Monitoring 46
	Requirements 47
	Viewing Existing Filters 47
	Setting Filters 48
	Clearing Filters 48
	Fan LEDs 49
	Fan Monitoring 49
	Fan Base Service Page 49
	MIS-S9D Proprietary Network Interface 51
	BMC Integrated Web Console 52
	Setting the BMC IP Address 53
	Setting Up User Accounts for the BMC 54
	Help Files 54
	System Debug Log 54
	Server Health 55

	Sensor Readings	55
	Event Log	57
	Power Statistics	58
	Configuration Tab	59
	Alerts	59
	Alert Email	61
	Server Power Control	64
	Virtual Front Panel	65
	Intelligent Management Platform Interface (IPMI 2.0)	66
	Enabling IPMI on a Linux System	67
	Accessing IPMI on a Linux System	67
3	System Software	69
	Overview	69
	Section Guide	70
	Downloading & Installing Software	72
	Required Downloads	73
	Before Updating from Zones 1.0	73
	Linux	73
	Windows	73
	Installing MegaRAID Storage Manager for Linux	74
	Zones 1.4.2 for Linux	75
	Starting Zones 1.4.2 for Linux	75
	MegaRAID Storage Manager for Windows	75
	Python for Windows	76
	Zones 1.4.2 for Windows	76
	Starting Zones 1.4.2 for Windows	77
	Zones 1.4.2 for Linux & Windows	77
	Zoning MIS Server Platforms	77
	Zones Interface	78
	Opening a Session	79
	144 Drives Setup	81
	Saving Zoning Session	83
	Zoning MIS JBOD Platforms	84

Zones JBOD Interface 85
Opening a Session 85
Adapter Tab Assignment 88
Saving a Session 89
Downloading a Session 89
Loading a CSV File to Zones 90
Creating the Drive Groups in MegaRAID (Linux & Windows) 91
Formatting the Drives using YaST2 in Linux 93
Verify Drives Seen in YaST2 95
Partitioning the Drives in YaST2 96
Windows Server Manager 98
Verify Drives Seen in Windows Server Manager 98
Formatting the Drives in Windows Server Manager 99
CLI Zoning Tool, version 1.4 105
Before Updating from CLI Zoning Tool 1.0 105
Installing CLI Zoning Tool on CentOS or Linux Host 105
Installing CLI Zoning Tool on Windows Host 106
Basics of CLI Zoning 106
Preparing to Zone using the CLI Zoning Tool 110
Editing the ShackCLI.ini file for CentOS & Linux 111
Editing the ShackCLI.ini file for Windows 112
CLI Zoning Tool Main Menu 114
Version 1.4.0 Default Menu 114
Editing the .csv File for the CLI Zoning Tool 116
Version 1.4.0 Expanded Menu 118
9) Display StorBrick(s) Settings 118
10) Display CLI Settings 119
11) Enter StorBrick CLI (Must select a single StorBrick) 120
12) Reboot Fan Base - will not reset StorBricks 120
13) Reset Fan Base - Danger. This will reset StorBricks 120
14) Force StorBrick(s) into Debug Mode 120
15) Exit StorBrick(s) Debug Mode 120
16) Display PHY Error Counters for selected StorBricks 121

	17) Display PHY Information for selected StorBricks123
	18) Display StorBrick UUID for selected StorBricks125
	19) Display StorBrick Firmware Revision Levels for Selected StorBricks125
	20) Change state of JBOD-IO ID LED - off, on blink.126
	21) Clear Persistent Flash126
	22) Display Fan Base Status126
	23) Display JBOD-IO Status129
	24) Change Polling of Expanders by Fan Base or JBOD-IO130
	25) Display Disk Drive Status131
	26) Change Power State of a Disk Drive131
	27) Power Cycle test131
	28) Update StorBrick UUID131
	29) Write/Read Buffer Test131
	Disk RAID Support132
	RAID Configuration Notes133
	Spare Drives135
4	System Maintenance137
	Detecting Component Failures138
	Selecting Which Side to Access.139
	Sliding the Chassis Forward/Backwards140
	Removing the Front or Rear Chassis Cover.140
	Replacing a Power Supply141
	Replacing a Fan Module.143
	Replacing a Boot Drive145
	Replacing a Capacity Drive148
	Removing a Drive149
	Re-installing a Drive149
	Checking the System Air Flow150
	Installing the External Air Ducts151
5	Troubleshooting153
	No Video154
	Losing the System's Setup Configuration154

Safe Power-Off	154
“Missing” Firmware Files	155
Cannot Receive Email Alerts Using BMC	155
Intel BIOS POST error messages and handling description	158
Beep Codes	161
Not all beep codes signal an error	163
Fans don’t spin when power button pressed	163
Fans spin when power button is pressed, but no video display is seen	163
BIOS or board logo appear, but OS load screen never appears	164
OS load screen appears, but nothing further	164
OS fully loads, but errors are seen	164
Cannot boot after new software or drivers were installed	164
Cannot boot after settings were changed in BIOS	164
System Gets Stuck in a Reboot Cycle	165
Shortening the Cable Management Arm	165
Installing the Cable Management Arm	172
A Technical Specifications	173

Figures

Figure 1-1	SGI Destination Rack (D-Rack)	4
Figure 1-2	MIS Chassis and Case	5
Figure 1-3	Bi-directional Rail Mounts	6
Figure 1-4	Single Control Panel	6
Figure 1-5	Dual Control Panel	7
Figure 1-6	Rear View – MIS Single Server Platform	7
Figure 1-7	Rear View – MIS Single Server Platform (four power supplies)	8
Figure 1-8	Rear View – MIS Dual Server Platform	8
Figure 1-9	Rear View – MIS Dual Server Half-Height Platform (single dual-server module)	8
Figure 1-10	Rear View – MIS JBOD Unit (single I/O module)	9
Figure 1-11	Rear View – MIS JBOD Unit (dual I/O modules)	9
Figure 1-12	Cable Management Arm	10
Figure 1-13	Power Supply Module (rated at 1100 Watts)	12
Figure 1-14	Power Supply Numbering	12
Figure 1-15	Fan Assembly Module (each contains two impellers)	13
Figure 1-16	StorBrick Modules for 3.5" or 2.5" 15mm Drives (left) and 2.5" 9.5 mm Drives (right)	14
Figure 1-17	3.5" 15mm Drive and Carrier	15
Figure 1-18	3.5" 15mm Drive Carrier (top view, with thumb latch)	15
Figure 1-19	Two 2.5" 9.5mm Drives and Carrier	15
Figure 1-20	2.5" 9.5mm Drive Carrier (isometric view with dual thumb latches)	16
Figure 1-21	MIS Single Server Platform, version 1.0.	17
Figure 1-22	MIS Dual Server Platform, version 1.0	18
Figure 1-23	MIS JBOD Unit, version 1.0	19
Figure 1-24	MIS Server Module – Single Server	20
Figure 1-25	Single Server Module – Component View	20
Figure 1-26	MIS Dual Server Module – Half Height	21

Figure 1-27	Dual Server Module – Component View 21
Figure 1-28	CPU and PCIe Riser Layout 22
Figure 1-29	HBA Population Layout 23
Figure 1-30	Boot Drive Module 24
Figure 1-31	I/O Module for MIS JBOD Unit 25
Figure 1-32	MIS JBOD Midplane I/O Connector (right & left views) 25
Figure 1-33	System-Level Block Diagram. 26
Figure 2-1	MIS Control Panel 31
Figure 2-2	BIOS POST Code and System LEDs. 36
Figure 2-3	NIC Port LEDs 37
Figure 2-4	JBOD I/O Panel LEDs 38
Figure 2-5	Disk Drive LEDs 38
Figure 2-6	MegaRAID Storage Manager – Alerts Pop-up Window 42
Figure 2-7	MegaRAID Storage Manager – Configure Alerts Window 43
Figure 2-8	MegaRAID Storage Manager – Mail Server Tab. 44
Figure 2-9	MegaRAID Storage Manager – Email Tab 45
Figure 2-10	Power Supply LEDs 46
Figure 2-11	Fan Module LEDs. 49
Figure 2-12	Fan Base Service Page (Server) 50
Figure 2-13	Fan Base Service Page (JBOD) 51
Figure 2-14	MIS-S9D Proprietary Network Interface 52
Figure 2-15	BIOS Server Management Tab 53
Figure 2-16	BMC Web Console – Logout, Refresh, and Help Buttons 54
Figure 2-17	BMC Web Console – System Debug Log 55
Figure 2-18	BMC Web Console – Server Health 56
Figure 2-19	BMC Web Console – Event Log 58
Figure 2-20	BMC Web Console – Power Statistics (not used) 58
Figure 2-21	BMC Web Console – Alerts Page 59
Figure 2-22	BMC Web Console – Alert Email Settings 61
Figure 2-23	BMC Web Console – Alert Email Result. 62
Figure 2-24	Alert Email – Power Cord Pulled from Power Supply Module 63
Figure 2-25	Alert Email – Power Module Pulled From Chassis 63
Figure 2-26	BMC Web Console – Power Control and Status 65

Figure 2-27	BMC Web Console – Virtual Front Panel	66
Figure 3-1	Zones 1.4.2 – Current Zoning Configuration	78
Figure 3-2	Zones 1.4.2 – Bottom Row of Buttons	78
Figure 3-3	Zones 1.4.2 – Step 1: Opening a Session	79
Figure 3-4	Zones 1.4.2 – New Session Alias	79
Figure 3-5	Zones 1.4.2 Warning Box – Alias Rules	80
Figure 3-6	Zones 1.4.2 – Top Menu Bar and Adapter Tabs	80
Figure 3-7	Zones 1.4.2 – Setup Menu Options	81
Figure 3-8	Zones 1.4.2 – Warning 144 Drive Configuration	82
Figure 3-9	StorBrick Diagram for 144-Drive Systems	83
Figure 3-10	Zones 1.4.2 – Brick Selection Window	84
Figure 3-11	Zones 1.4.2 – JBOD Zoning Interface	85
Figure 3-12	Zones 1.4.2 – JBOD Open Session Window	86
Figure 3-13	Zones 1.4.2 – JBOD New Session Alias	87
Figure 3-14	Zones 1.4.2 – JBOD Phy-based Zoning Interface	88
Figure 3-15	Zones 1.4.2 – JBOD Brick Selection Window	89
Figure 3-16	Zones 1.4.2 – Choose CSV File Window	90
Figure 3-17	Zones 1.4.2 – CSV Warning Window	91
Figure 3-18	MegaRAID – Create a Virtual Drive	91
Figure 3-19	MegaRAID – Create Virtual Drive Mode	92
Figure 3-20	Create Virtual Drive – Drive Group Settings	92
Figure 3-21	Create Virtual Drive – Summary	93
Figure 3-22	Yast2 Server Manager GUI	94
Figure 3-23	YaST2 – Warning Message	95
Figure 3-24	YaST2 – Drives Have Appeared	95
Figure 3-25	YaST2 – Add Button	96
Figure 3-26	YaST2 – Select Partition Size	96
Figure 3-27	YaST2 – Format & Mount the Drive	96
Figure 3-28	YaST2 – Check for Partition	97
Figure 3-29	YaST – Click Finish	97
Figure 3-30	YaST2 – Disk Mounting (in process)	98
Figure 3-31	Windows Server Manager – Disk Management	99
Figure 3-32	Server Manager – Disk Management	100

Figure 3-33	Server Manager – Initialize Disks	100
Figure 3-34	Server Manager – Select GPT (GUID Partition Table)	101
Figure 3-35	Server Manager – Disks Initialized and Online	101
Figure 3-36	Server Manager – New Simple Volume	102
Figure 3-37	Server Manager – New Simple Volume Wizard	102
Figure 3-38	New Simple Volume Wizard – Volume Size	103
Figure 3-39	New Simple Volume Wizard – Assign Drive Letter or Path	103
Figure 3-40	New Simple Volume Wizard – Format Partition	104
Figure 3-41	New Simple Volume Wizard – Settings Confirmation	104
Figure 3-42	New Simple Volume in Server Manager	105
Figure 3-43	MIS-S9D Proprietary Network Interface	111
Figure 3-44	Block diagram of MIS Server StorBrick SB0	117
Figure 3-45	RAID 1 with One Drive per StorBrick	133
Figure 3-46	RAID 1 with Two Drives Spanning a StorBrick	133
Figure 3-47	RAID 5 or 6 with One Drive per StorBrick	134
Figure 3-48	Loss of a Drive with Multiple Drives on a StorBrick.	135
Figure 3-49	Three Drive Loss in RAID 6 Require StorBrick Replacement	135
Figure 4-1	MIS-S9D Proprietary Network Interface	138
Figure 4-2	Front & Rear Chassis Covers	140
Figure 4-3	Top Cover Security Screw	141
Figure 4-4	Power Supply Module	141
Figure 4-5	Fan Module	143
Figure 4-6	Fan Base Service Page	144
Figure 4-7	Capacity Drive and Carrier	148
Figure 4-8	MIS Chassis Midspan Support Brace	151
Figure 4-9	Rack Air Duct Placement	152
Figure 5-1	BMC Alerts Email Settings Page	156
Figure 5-2	BMC Alert Configuration Success Pop-up	156
Figure 5-3	BMC Alerts Configuration Page	157
Figure 5-4	BMC Alert Configuration Success Pop-up	158
Figure 5-5	Cable Management Arm	166
Figure 5-6	Remove Segments from the End of an Empty Arm	167
Figure 5-7	Pry the Joint Apart	167

Figure 5-8	Pull the Segments from the Arm168
Figure 5-9	Shortened Arm Ready for Installation168
Figure 5-10	Remove Segments from the Middle of a Full Arm169
Figure 5-11	Pry Open the Segment Joint169
Figure 5-12	Pull Apart the Cable Management Arm170
Figure 5-13	Pry Apart the Joint at the Other End170
Figure 5-14	Removing the Segment from the Cables.171
Figure 5-15	Join the Spliced Ends of the Cable Management Arm171
Figure 5-16	Velcro Cable Ties172

Tables

Table -1	MIS Server Platform Region and EMC Compliance References . . .	xxiii
Table -2	MIS Server Platform Region and Functional Safety Listing Mark . . .	xxiv
Table 2-1	MIS Server Platform Control Panel Buttons and LEDs	32
Table 2-2	System Status LED Meanings	36
Table 2-3	NIC Port LED meanings	37
Table 2-4	Disk Drive LEDs	39
Table 2-5	Events	41
Table 2-6	Power Supply LEDs	46
Table 2-7	Server Power Control Actions	64
Table 3-1	Zoning Configuration107
Table 3-2	Phy Based Zoning Table for StorBrick 0108
Table 3-3	Zone Group Implementation109
Table 3-4	CLI Zoning Tool Menu Options and Descriptions114
Table 3-5	Zone Group Implementation116
Table 4-1	Access Points139
Table 5-1	Intel BIOS POST Error Messages158
Table 5-2	Summary of BMC Beep Codes162
Table 5-3	Summary of LSI MegaRAID Card Beep Codes162
Table A-1	Technical Specifications173

Introduction


This guide describes the features and components of the SGI® Modular InfiniteStorage™ (MIS) platform, version 1.0. With two main configurations possible for the enclosure (server and storage, or JBOD—Just Bunch Of Disks) this guide covers the different configurations, their respective components, interface panels, indicator lights and meanings, software, maintenance, and troubleshooting.


Audience

This guide is written for owners/users of the MIS platform, version 1.0. 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, 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 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 rail replacement. 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). More information on powering off the MIS Platform is available in Chapter 4, “System Maintenance.”

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. It includes information on Electromagnetic Compatibility and Safety Certification.

Electromagnetic Compatibility

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

Table -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 Certification
Taiwan	BSMI CNS 13438 Certification
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 -2 lists the region and compliance reference for safety compliance.

Table -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,” 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. Additional information includes the operating systems supported, and RAID configurations possible with the MIS enclosures.

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 for monitoring hardware status and configuring alerts.

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 (Zones 1.4 and CLI Zoning Tool 1.3). There are certain prerequisite programs necessary, depending on the operating system, and this chapter gives instructions for download and installation of these programs. In this chapter are step-by-step instructions for Zones tool, its features and their function, plus warnings and error codes. Screen shots are given for a walk-through of the tool, plus step-by-step instructions for zoning. This Chapter also contains information on using the CLI Zoning tool (both programs can zone JBODs at this time).

Chapter 4, “System Maintenance,” describes how to use Sensor Data Records 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 storage drive in a StorBrick (15mm and 9mm profiles), how to replace a boot drive, how to replace a fan module, and additional air flow precautions.

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.

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, see Figure 2-17 on page 61)
- 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) for translation of the hex codes used in Alert Emails (see Figure 2-24 on page 70)
- Intel’s *Intel Server Board S2600JF, Technical Product Specification*:
- IPMItool comes with Linux. Information on IPMItool can be found through <http://support.sgi.com>
- IPMIutil is for Windows and is available for download through <http://ipmiutil.sourceforge.net/>. Information on this tools can be found at <http://ipmiutil.sourceforge.net/docs/UserGuide>. (SGI does not provide support for Windows, or any Windows utilities.)
- 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
- 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.

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 is a high-density, integrated storage server platform. The MIS Platform uses a 4U rackmount system, and can be either a compute and storage server, or a “Just Bunch Of Disks” expansion storage unit (MIS JBOD unit). The MIS Server Platform can be single or dual server. Up to 5 MIS enclosures (server & JBODs) or 6 JBODs can be mounted into a SGI Destination rack (D-Rack), as shown in Figure 1-1. (Other 3rd-party 19" racks are also supported.) A single D-Rack has space for up to 10 enclosures, however, due to floor weight regulations, only 5-6 units may be installed in a D-Rack (see “Environmental Requirements” on page xx).

Features of the modular design of the MIS 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 can be attached to one MIS Server Platform.



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

The MIS Server Platform features:

- Up to 2 server modules per platform.

- One or two Intel® Xeon® E5-2600 series processors per server motherboard.
- 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).
- Up to 8 DDR3 DIMMs (4 GB, 8 GB, or 16 GB) for a single-server motherboard configuration, and up to 16 DIMMs for a dual-server motherboard configuration.
- Up to 4 HBAs for a single server, full-height (4.25") and half-depth (3.375"), externally or internally facing. Up to 4 HBAs (half-height, half-depth; 2 per server module) for a dual server. There are an additional two internally facing, half-height and half-depth HBAs per server module, used by the system. The MIS Single Server Platform can have a total of six HBAs, where a Dual Server Platform can have a maximum of eight (including those used by the system).
- Up to 3 PCIe riser cards for single server systems (dual-servers have a mandatory 3 PCIe risers, regardless of card count).
- 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 total maximum of six. (Unique BBU PCIe technology allows the inclusion of BBUs without the consumption of any of the available PCIe slots.)
- Two boot drives per server: SAS or SATA, rotational or SSD, up to 300GB, mirrored using RAID 1.
- Dual GbE networking onboard, with an optional 2 port 10GbE, 2 port GbE, or 4 port¹ 8Gb FC PCIe cards (4 optional networking PCIe cards maximum, external facing only; risers 1 and 2).

The System Overview will explore:

- “MIS Enclosure” on page 5 including
 - “Front Grille and Control Panels” on page 6
 - “Rear Panel Components” on page 7 and
 - “Cable Management Arm” on page 9
- “MIS Common Modules” on page 10 presents

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 card. The four port can be used in the single server.

-
- “Power Supply Module” on page 11,
 - “Fan Assembly Module” on page 13, and
 - “StorBrick Module” on page 13.
 - “MIS Server Platform or JBOD Unit” on page 16, discusses the presence of the
 - “Server Module” on page 16 its available features and associated
 - “Layout of Server CPUs, and PCIe Risers HBAs” on page 22 and
 - “Boot Drives Module” on page 24
 - or the presence of a ninth StorBrick and “MIS JBOD I/O Module” on page 25

On the final pages,

- A “System Block Diagram” on page 25 shows a diagram of the enclosure with a dual-server configuration (optional JBOD components show in grey), signals and power connections,
- While a “System Layout” on page 27 gives a map of the physical layout of the components in the chassis.



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

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 1-2 MIS Chassis and Case

The SGI MIS chassis features a front bezel display with an internal EMI grille (Figure 1-2). The unique bi-directional sliding rail mounts (Figure 1-3) 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 their full length to be serviced.




Figure 1-3 Bi-directional 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 1-4 shows a single control panel and Figure 1-5 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 (to be used by SGI field operatives only). Indicator light meanings and button functions are explained in, “Control Panel” in Chapter 2.




Figure 1-4 Single Control Panel




Figure 1-5 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 (only two are pictured in any of the figures here). 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 1-6, features a single server module with two USB ports, a video port, and two NIC ports. Figure 1-7 show a MIS Server Platform with the optional four power supplies. The MIS Server Platform (dual server) rear panel as shown in Figure 1-8 has a second server module with its own set of USB ports, video port, and NIC ports. Figure 1-9 shows an MIS Server Platform (single dual-server module) which features dual server construction with a single server installed, and the option of upgrade to include a second server later. Figure 1-11 shows the rear panel of the MIS JBOD with two I/O modules.




Figure 1-6 Rear View – MIS Single Server Platform




Figure 1-7 Rear View – MIS Single Server Platform (four power supplies)




Figure 1-8 Rear View – MIS Dual Server Platform




Figure 1-9 Rear View – MIS Dual Server Half-Height Platform (single dual-server module)




Figure 1-10 Rear View – MIS JBOD Unit (single I/O module)




Figure 1-11 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 1-12).




Figure 1-12 Cable Management Arm

Instructions on shortening the cable management arm are in Chapter 5, “Troubleshooting.”

MIS Common Modules

This section describes the internal modules that both the MIS Server and MIS JBOD share in common. Designed to deliver a high level of reliability, scalability and manageability, the MIS platform makes use of modules to contain key components. Whether the unit is an MIS Server or a JBOD, both chassis contain the following hot-swappable modules:

- Up to four power supplies (two redundant) (Figure 1-13);
- Six fan assemblies (Figure 1-15);

- Capacity drives installed in StorBricks (Figure 1-16).

The power supply modules are high-efficiency, hot-swappable power supplies 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. And instead of the conventional disk architecture, the unique StorBrick modules—innovative, highly-dense drive modules used to house drive bays—allows 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.



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

Power Supply Module

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 rest in standby (if not off or in a fault state). The power supplies provide 12V DC main power and 5V DC standby power. The power supplies are hot-swappable and can be replaced under full load. Power supplies are numbered 0-3 from the bottom up, on the rear panel of the enclosure (Figure 1-14).

The power supply is an 1100 Watt AC to DC power-factor-corrected (PFC) power supply that converts standard AC mains power into a main output of 12 VDC for powering intermediate bus architectures (IBA) in high performance and reliability servers, routers, and network switches. The power supply meets international safety standards and displays the CE-Mark for the European Low Voltage Directive (LVD).

The power supplies are hot-swappable as long as the redundant power option is installed. The power supplies can be replaced under full load as long as there are three or more functioning power supplies present.

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




Figure 1-13 Power Supply Module (rated at 1100 Watts)




Figure 1-14 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. The fuses are not accessible from the outside and are therefore not serviceable parts (instead the module is replaced as a whole).

Fan Assembly Module




Figure 1-15 Fan Assembly Module (each contains two impellers)

Six fan assemblies mounted in the middle of the chassis 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.

StorBrick Module

Each StorBrick Module contains up to nine 3.5" or 2.5" 15mm drives (Figure 1-16), or eighteen 2.5" 9mm drives (Figure 1-20 on page 16), mounted in the StorBrick using proprietary drive carriers (Figure 1-17). A sliding thumb latch securely fastens the drive carriers in place (Figure 1-18, 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 and/or SATA drives (rotational or SSDs).




Figure 1-16 StorBrick Modules for 3.5" or 2.5" 15mm Drives (left) and 2.5" 9.5 mm Drives (right)



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




Figure 1-17 3.5" 15mm Drive and Carrier



Figure 1-18 3.5" 15mm Drive Carrier (top view, with thumb latch)




Figure 1-19 Two 2.5" 9.5mm Drives and Carrier




Figure 1-20 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 (Figure 1-21 or Figure 1-22) and the MIS JBOD Unit (Figure 1-23) is the presence of the compute server module (Figure 1-24 or Figure 1-26) and boot drives (Figure 1-30) in the Server Platform, or a ninth StorBrick (Figure 1-16) and I/O modules (Figure 1-31) and associated midplane (Figure 1-32) in the JBOD.

Server Module

The MIS Server Platform can be single or dual server (Figure 1-24 or Figure 1-26) 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 (4 GB, 8 GB, or 16 GB) for a single server configuration. Up to 16 DIMMs for a dual-server board configuration,
- Up to 4 HBAs for a single server, full-height (4.25") and half-depth (3.375"), externally or internally facing. Up to 4 HBAs (half-height, half-depth; 2 per server module) for a dual server. (See Figure 1-28 on page 22)
- 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. Up to three battery back up units per server module for a dual server platform, for a total maximum of six. (Unique BBU PCIe technology allows the inclusion of BBUs without the consumption of any of the available PCIe slots.)




Figure 1-21 MIS Single Server Platform, version 1.0




Figure 1-22 MIS Dual Server Platform, version 1.0




Figure 1-23 MIS JBOD Unit, version 1.0




Figure 1-24 MIS Server Module – Single Server




Figure 1-25 Single Server Module – Component View




Figure 1-26 MIS Dual Server Module – Half Height




Figure 1-27 Dual Server Module – Component View

Layout of Server CPUs, and PCIe Risers HBAs

Figure 1-28 shows the CPU and riser layout. Because of cabling restrictions in a single server, single CPU systems, only two internal SAS HBAs are allowed. The first CPU handles Riser 1 and 2. The optional second CPU would manage Riser 3. If the second CPU is not installed, Riser 3 is non-operational. When a second CPU is installed, HBAs populated on Riser 3 are internal facing SAS HBAs only, which connect to the StorBricks (Figure 1-29).





Figure 1-28 CPU and PCIe Riser Layout

**Dual CPUs:**

RISER2-J1 is reserved for a internal SAS HBA and when populated connects to the StorBricks

RISER2-J2 is reserved for a internal SAS HBA and when populated connects to the StorBricks

RISER1-J1 can be populated with external facing HBAs

RISER1-J2 can be populated with external facing HBAs

If a internal SAS HBA is not populated in RISER2-J2, an external HBA can be populated

If a internal SAS HBA is not populated in RISER2-J1, an external HBA can be populated

Single CPU:

RISER2-J1 is reserved for a internal SAS HBA and when populated connects to the StorBricks

RISER2-J2 is reserved for a internal SAS HBA and when populated connects to the StorBricks

RISER1-J1 can be populated with external facing HBAs

RISER1-J2 can be populated with external facing HBAs

If a internal SAS HBA is not populated in RISER2-J2, an external HBA can be populated

Figure 1-29 HBA Population Layout

Boot Drives 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, used to store server data and the server operating system. Supported operating systems include:

- Red Hat® Enterprise Linux (RHEL) 6.2,
- SUSE LINUX® Enterprise Server 11 SP1, or
- VMware® ESX 5.0.

For Microsoft® Windows® 2008 R2 SP1, the hardware is WHQL certified and sold as such. Customers who wish to run Windows buy only the hardware from SGI. The customer is then responsible for operating system installation and configuration, and all support arrangements with Microsoft. SGI does not provide technical support for Windows drivers, utilities (including IPMIutil), device manager, etc.




Figure 1-30 Boot Drive Module

MIS JBOD I/O Module




Figure 1-31 I/O Module for MIS JBOD Unit

JBOD I/O modules slide into a midplane (Figure 1-32), which connect to the SAS controllers.




Figure 1-32 MIS JBOD Midplane I/O Connector (right & left views)

System Block Diagram

Figure 1-33 shows the system-level block diagram for a fully populated dual-server (the optional JBOD components are shown in grey: “STORBRICK 8” and “JBOD only”).

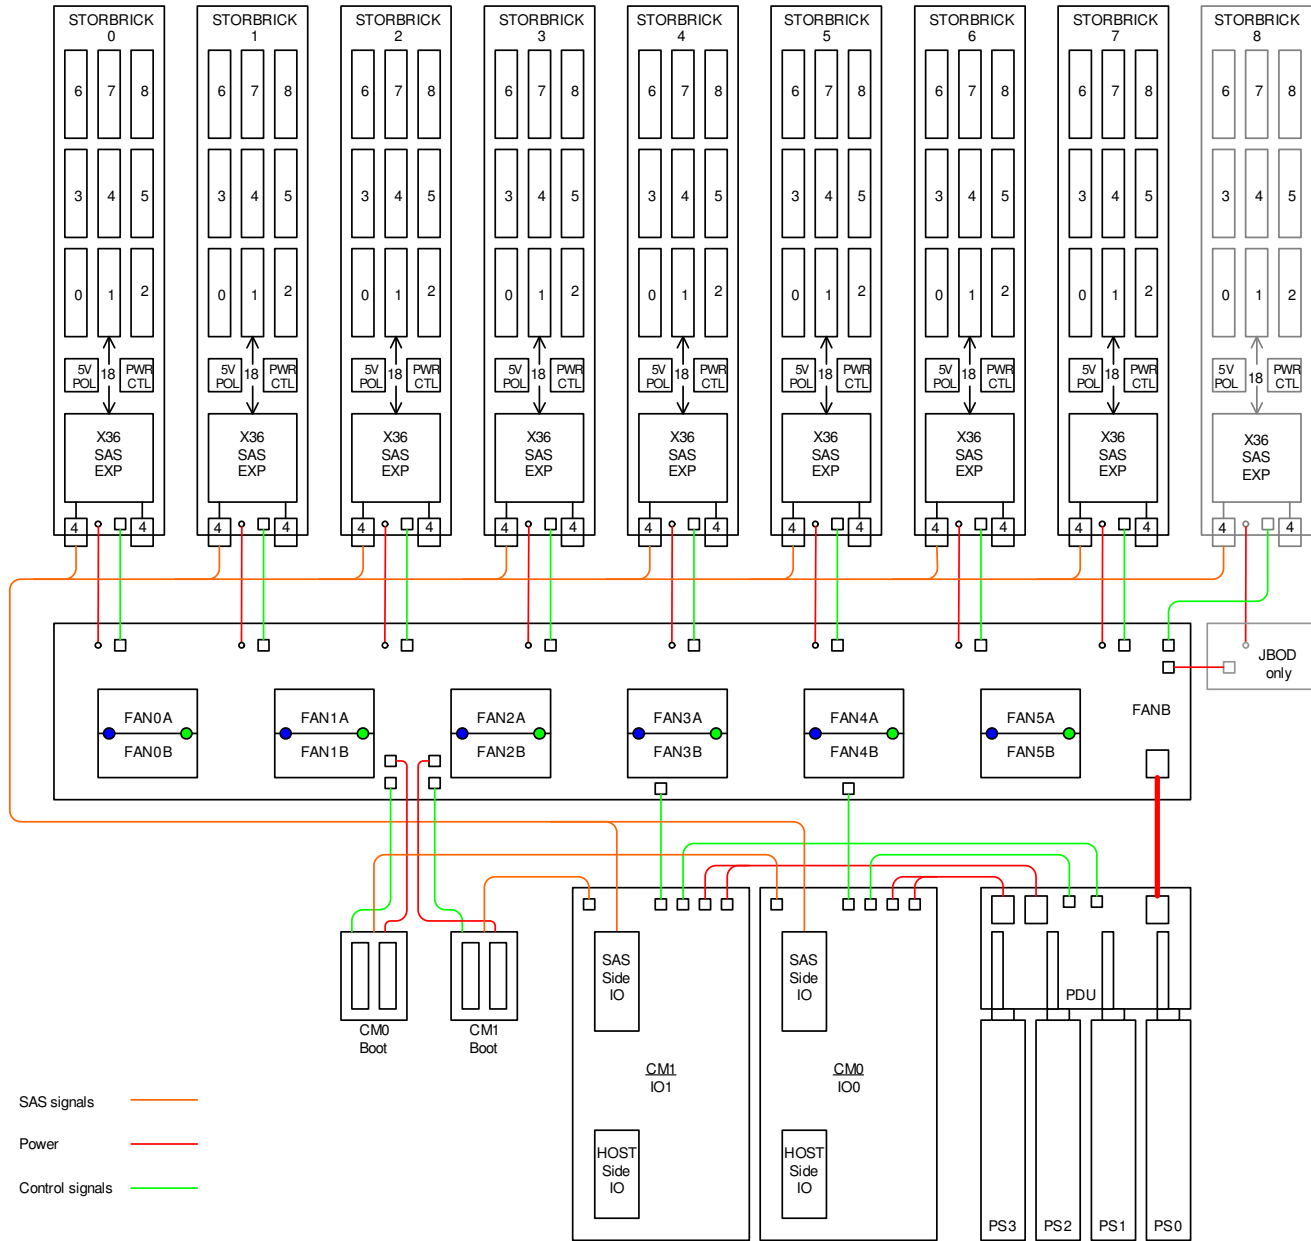
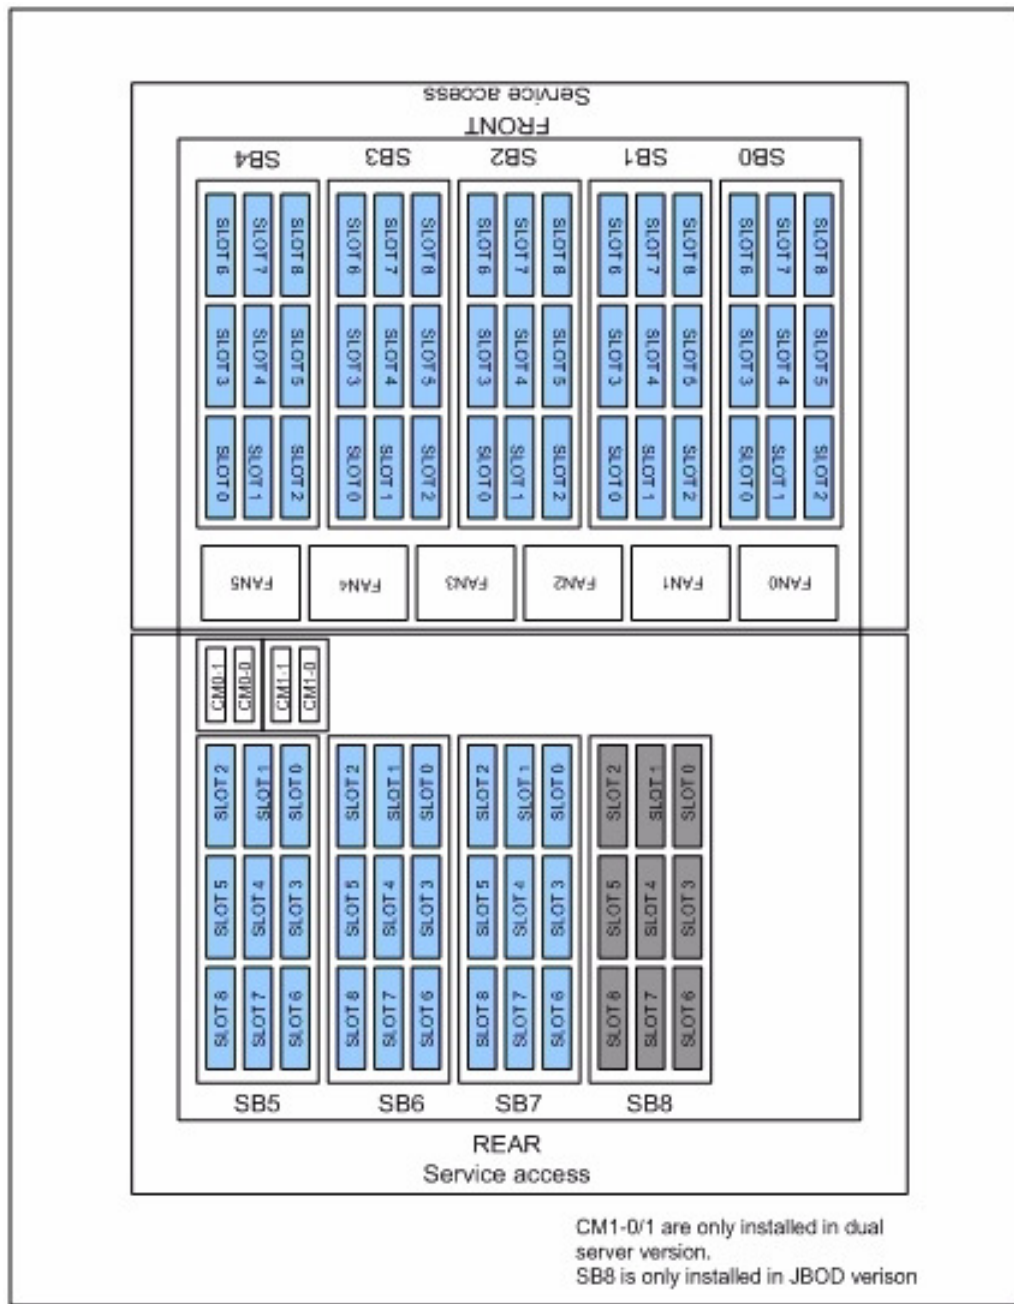


Figure 1-33 System-Level Block Diagram

System Layout

On the following page is the layout of the MIS platform.



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.

- “Control Panel” on page 30
 - “MIS JBOD vs. Server Control Panel” on page 32
 - “Powering an MIS Platform” on page 33
 - “Turning Off an MIS Platform” on page 34
- “Server Module LEDs” on page 35
- “JBOD LEDs” on page 38
- “Disk Drive LEDs” on page 38
- “Drive Monitoring” on page 39
 - “Boot Drive Monitoring” on page 39
 - “Capacity Drive Monitoring” on page 42
- “Power Supply LEDs” on page 45
 - “Requirements” on page 47
 - “Viewing Existing Filters” on page 47

- “Setting Filters” on page 48
- “Clearing Filters” on page 48
- “Fan LEDs” on page 49
- “Fan Monitoring” on page 49
- “Fan Base Service Page” on page 49
- “MIS-S9D Proprietary Network Interface” on page 51
- “BMC Integrated Web Console” on page 52
 - “Setting the BMC IP Address” on page 53
 - “Setting Up User Accounts for the BMC” on page 54
 - “Help Files” on page 54
 - “System Debug Log” on page 54
 - “Server Health” on page 55
 - “Event Log” on page 57
 - “Power Statistics” on page 58
 - “Configuration Tab” on page 59
 - “Server Power Control” on page 64
 - “Virtual Front Panel” on page 65
- “Intelligent Management Platform Interface (IPMI 2.0)” on page 66
 - “Enabling IPMI on a Linux System” on page 67
 - “Accessing IPMI on a Linux System” on page 67

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

The control panel (Figure 2-1) interface consists of five indicator lights and four buttons. The **NMI Reset** button (Non-Maskable Interrupt) is not supported and is non-functional.

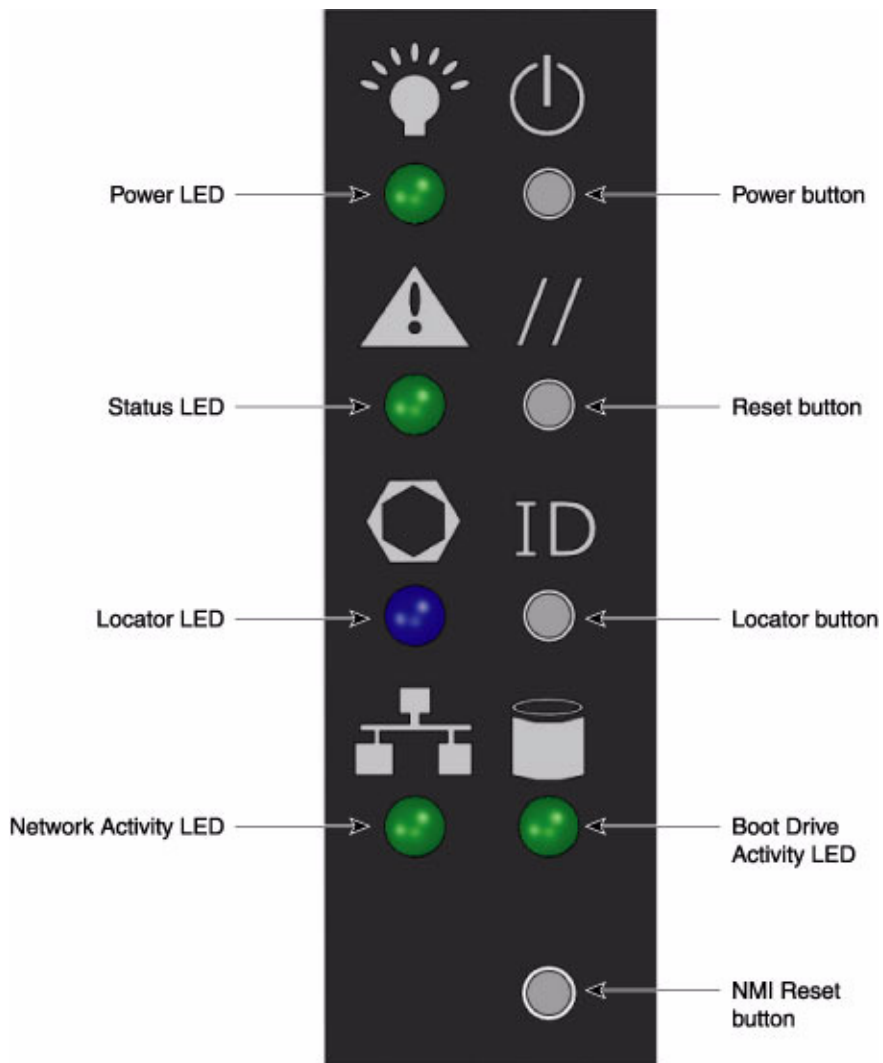


Figure 2-1 MIS Control Panel

MIS JBOD vs. Server Control Panel

The control panel (Figure 2-1) for the MIS JBOD is exactly the same as the MIS Server Platform. However, 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.

When there are *two* I/O modules in a JBOD, the top control panel connects to the bottom I/O module, and the bottom control panel accesses the top I/O module.

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.

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

LED/Button	Description
Locator button	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: <code>-H <ip address> -P <password> -U <user> chassis identify</code>
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	Unused.

Powering an MIS Platform

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 BMC, 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.

The front panel LED blinks for 30-60 seconds upon startup while the baseboard management controller (BMC) is initializing. If not, then the BMC is not functioning, and try to gather system event log (SEL) information for investigation (see "Event Log" on page 57 for more information.). Additional, 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 Power-On Self-Test (POST).

When the MIS Server is first powered on, it gives you the option to press **F2** to enter the BIOS, **F6** to enter the Boot Menu (to change how the system boots), or **F12** for a Network Boot (what's also known as PXE or "pixie boot").

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.

Using IPMItool

Note: The Linux IPMI Service must be enabled and running. See “Intelligent Management Platform Interface (IPMI 2.0)” on page 66.

First, verify that the power cables are properly connected to the system. Use the IPMItool `power on` command to power up the server.

On a local server: # `ipmitool power on`
For a remote server: # `ipmitool -H <host> -U <user> -P <pswd> power on`
For example: # `ipmitool -H 137.8.64.3 -U admin -P Password01 power on`
where the host IP address is 137.8.64.3, the user is admin with the password: Password01.

Turning Off an MIS Platform

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 Operating System GUI power-off button at the console screen.
- When logged in via an ssh session and executing an IPMItool `power off` command.
- Using the remote console screen GUI power-off button, if a keyboard/video monitor/mouse (KVM) RMM4 session is established through the BMC.
- Using the Fan Base Interface Page and pressing the **Shutdown** button.

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.

To shut down using IPMItool, use the following commands:

On a local server: # **ipmitool power off**

For a remote server: # **ipmitool -H <host> -U <user> -P <pswd> power off**

For example: # **ipmitool -H 137.8.64.3 -U admin -P Password01 power off**
where the host IP address is 137.8.64.3, the user is admin with the password: Password01.

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.

Server Module LEDs

On the back panel of the server module are three sets of LED indicator lights, not including any indicator lights that may exist on external HBAs. From the top down, these sets of LEDs include 8 internal Power-On Self-Test (POST) Code Diagnostic LEDs (labeled 1-8, Figure 2-2), bi-colored System Status and blue ID LEDs (A & B, respectively, Figure 2-2), and two NIC ports, each with a green Link LED and bi-colored Activity Rate LED (Figure 2-3).

See [Intel Server Board S2600JF, Technical Product Specification](#), page 170, for more information on what each POST LED pattern indicates.




Figure 2-2 BIOS POST Code and System LEDs

Table 2-2 System Status LED Meanings

LED/Button	Status	Description
Green Solid	Okay	– System ready.
Green Blink	Degraded	– Unable to use all installed memory when more than one DIMM is installed – PCIe correctable link errors – CPU disabled – Non-critical threshold crossed (temperature, voltage, power) – Battery failure
Amber Blink	Non-Fatal	– PCIe uncorrectable link errors – Critical threshold crossed (temperature, voltage, power)

LED/Button	Status	Description
Amber Solid	Fatal	<ul style="list-style-type: none"> – DIMM failure, no good memory is present – Run-time memory uncorrectable error – CPU configuration error (e.g., processor stepping mismatch), overtemperature – Power fault, insufficient power, not enough power supplies present
Off	Not Ready	– Main power off




Figure 2-3 NIC Port LEDs

Table 2-3 NIC Port LED meanings

LED/Button	Description	NIC State
Green LED (A)	On	Active connection
	Blinking	Send/Receive activity
Green/Amber (B)	Off	10 Mbps
	Amber	100 Mbps
	Green	1000 Mbps

JBOD LEDs




Figure 2-4 JBOD I/O Panel LEDs

On the back panel of the JBOD I/O Module are four LEDs. The top is a blue ID LED labeled with a circle in a hexagram, next is a bi-colored Status LED, **ExP** (external SAS ports) and **SxP** (internal SAS ports) activity lights.

Disk Drive LEDs

Figure 2-5 shows the disk drive LEDs, a blue ID LED and bi-colored Status LED.



Figure 2-5 Disk Drive LEDs

Table 2-4 describes the meaning of disk drive LEDs.

Table 2-4 Disk Drive LEDs

Status LED	Blue LED	Description
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.

Drive Monitoring

Capacity drive monitoring can be performed by MegaRAID Storage Manager. Boot drive monitoring, since it is not a part of the StorBricks, is monitored using **Multiple Devices Administration** tool: `mdadm`. By default, if the MIS platform comes with a factory-installed operating system, boot drives are mirrored using RAID 1 and capacity drives are shipped with RAID 6. For those MIS Servers not shipped with an OS, RSTe must be enabled in the BIOS to allow boot drive mirroring and hot-swappable capabilities.

The following sections explains how to use the MegaRAID Storage Manager software to monitor the status of capacity drives and the `mdstat` command to monitor boot drives (Windows 2008 uses the Windows System Resource Manager to enable email alerts, see [KB Article cc732728](#) for more information).

The MegaRAID Storage Manager software does a background check every hour to verify if the controller and the system time are in synch. If the time difference between the controller and the system is more than 90 seconds, the MegaRAID Storage Manager software synchronizes the time so that the controller time and the system time are in synch. `mdadm` polls the md arrays every 60 seconds (default value). This can be adjusted using the `-d` mode and entering in a value of seconds, though it is recommended that you keep the default value and only adjust this time for troubleshooting.

Boot Drive Monitoring

Issue a `#cat /proc/mdstat` command to see current status (`stat`) of of the multiple devices (`md`) attached to the processor (`proc`).

Boot drive status is indicated by either a U for up or _ for down. In the following example, everything is normal ([UU]), and md127 is the container for md125 and md126.

```
# cat /proc/mdstat
Personalities : [raid1] [raid0] [raid10] [raid6] [raid5] [raid4]
md125 : active raid1 sdb[1] sdc[0]
        111357952 blocks super external:/md127/0 [2/2] [UU]

md126 : active raid1 sdb[1] sdc[0]
        5556224 blocks super external:/md127/1 [2/2] [UU]

md127 : inactive sdc[1](S) sdb[0](S)
        6306 blocks super external:imsm
```

In the next example below, a drive has failed. Notice that sdc[0] is now missing from md125, md126 (their numbers changing from [2/2] to [2/1]), and the container md127. Additionally, a failed drive is indicated in the notation [U_].

```
# cat /proc/mdstat
Personalities : [raid1] [raid0] [raid10] [raid6] [raid5] [raid4]
md125 : active raid1 sdb[1]
        111357952 blocks super external:/md127/0 [2/1] [U_]

md126 : active raid1 sdb[1]
        5556224 blocks super external:/md127/1 [2/1] [U_]

md127 : inactive sdb[0](S)
        3153 blocks super external:imsm
```

To enable monitoring of the boot drives, the following usage of the mode `--monitor` is employed:

```
mdadm --monitor <raiddevice> [options] <devices>
```

This usage causes `mdadm` to periodically poll a number of md arrays and to report on any events noticed. If any devices are listed on the command line, `mdadm` will only monitor those devices, otherwise, all arrays will be monitored. Further, adding `--scan` causes the operation to be applied to all arrays listed in `/proc/mdstat`. If `--scan` is given, then an email address **must** be specified on the command line. If `--scan` is given and an email address is *not* available, then `mdadm` will not monitor anything. Without `--scan`, `mdadm` will continue monitoring as long as something was found to monitor.

The result of monitoring the arrays is the generation of events. These events are passed to a given email address using the option `-m=<email address>` or `--mail=<email address>`. If no email is given, then each event is reported to `stdout`.

The `-t` or `--test` option generates a **TestMessage** alert for every array found at startup. This alert gets mailed and passed to the alert program. This can be used for testing that alert message to get through successfully.

The different events are listed in Table 2-5 below.

Table 2-5 Events

Tag	Description
DeviceDisappeared	An md array which previously was configured appears to no longer be configured.
RebuildStarted	An md array started reconstruction.
Rebuild##	Where ## is 20, 40, 60, or 80, this indicates that rebuild has passed that many percentage of the total.
RebuildFinished	An md array that was rebuilding, isn't any more, either because it finished normally or was aborted.
Fail	An active component device of an array has been marked as faulty.
NewArray	A new md array has been detected in the <code>/proc/mdstat</code> file.
DegradedArray	A newly noticed array appears to be degraded. This message is not generated when <code>mdadm</code> notices a drive failure which causes degradation, but only when <code>mdadm</code> notices that an array is degraded when it first sees the array.
TestMessage	An array was found at startup, and the <code>--test</code> option was given.

Only **Fail**, **DegradedArray**, and **TestMessage** cause email to be sent. Each event has an associated array device (e.g., /dev/md1) and possibly a second device. For **Fail**, the second device is the relevant component device. Use the command `mdadm --help` or `mdadm <option> --help` for more information on `mdadm` or the `mdadm` options, respectively.

Capacity Drive Monitoring

When devices automatically go from an optimal state to a different state, such as a drive failure, the MegaRAID Storage Manager software gets those events from the controller and gives a notification using different alert delivery methods. A pop-up notification is started automatically when you login to the operating system (Figure 2-6).

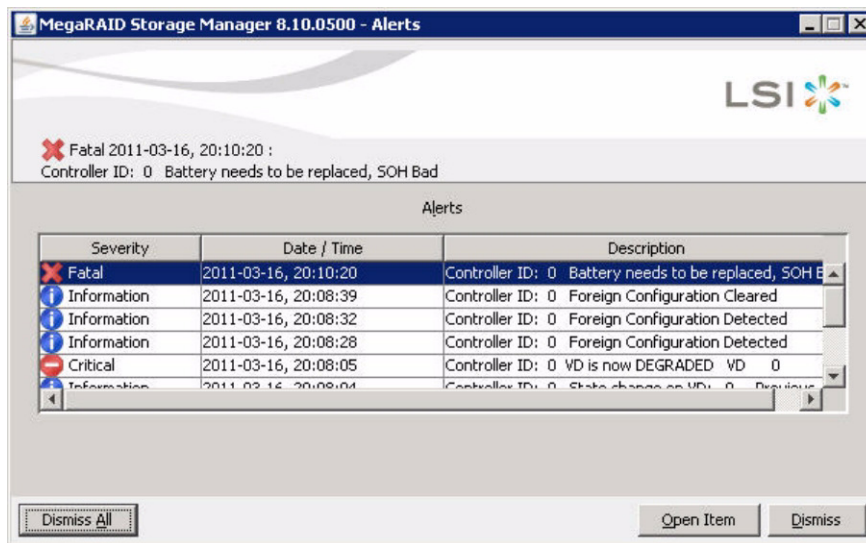


Figure 2-6 MegaRAID Storage Manager – Alerts Pop-up Window

To receive email alerts for capacity drive failures, from the main screen select **Tools** and **Configure Alerts**. In the **Configure Alerts** window, the **Alert Settings** tab allows you edit the alert delivery method for different severity levels by clicking the **Edit** button, or change the method of delivery and/or severity level for each individual event using the **Change Individual Events...** button (Figure 2-7).

Note: The **Save Backup...** button and **Load Backup...** buttons are used to save/load the *entire* alerts configuration to/from an .xml backup file, not just the settings for the tab with which you are working.

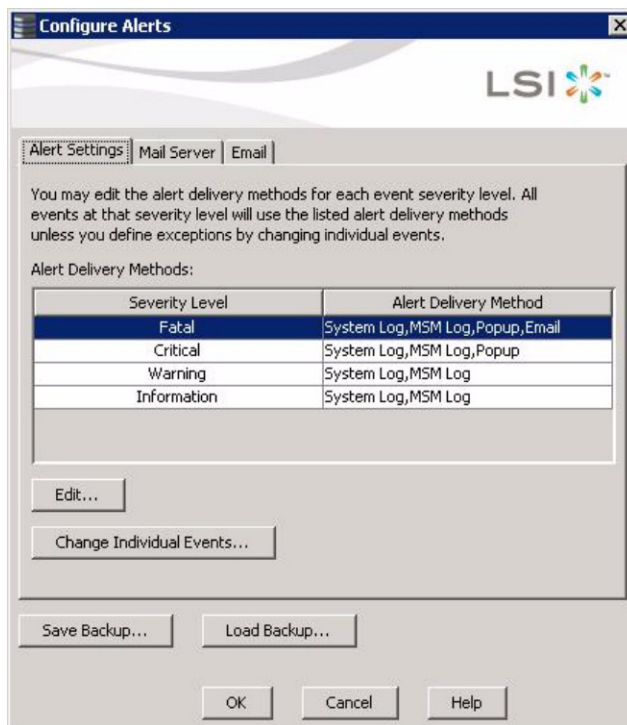


Figure 2-7 MegaRAID Storage Manager – Configure Alerts Window

Important: Whenever you load a saved backup file, all unsaved changes made in the current session will be lost.

The **Mail Server** tab is where to enter the *sender* email address for alerts, the SMTP server name or IP address of the mail server to send the alerts, and user name and password fields (filled only when the SMTP server requires authentication, Figure 2-8).



The screenshot shows the 'Mail Server' configuration tab in the MegaRAID Storage Manager. The LSI logo is in the top right corner. The interface has three tabs: 'Alert Settings', 'Mail Server', and 'Email'. The 'Mail Server' tab is active. The configuration fields are as follows:

- Sender email address: monitor@server.com
- SMTP Server: 127.0.0.1
- Port: 25
- Use Default
- This server requires authentication
- User name: (empty field)
- Password: (empty field)

At the bottom, there are buttons for 'Save Backup...', 'Load Backup...', 'OK', 'Cancel', and 'Help'.

Figure 2-8 MegaRAID Storage Manager – Mail Server Tab

The **Email** tab is where you to add/remove email addresses for recipients of alert notifications (Figure 2-9). To configure an email account to receive alerts, enter the email address you want to add in the **New recipient email address** field and click **Add**. The email address should appear in the **Recipient email addresses** pane.

Once the changes to the Mail Server tab and Email tabs have been made, click the Test button to make sure your configuration is correct. If satisfied, click **OK**.

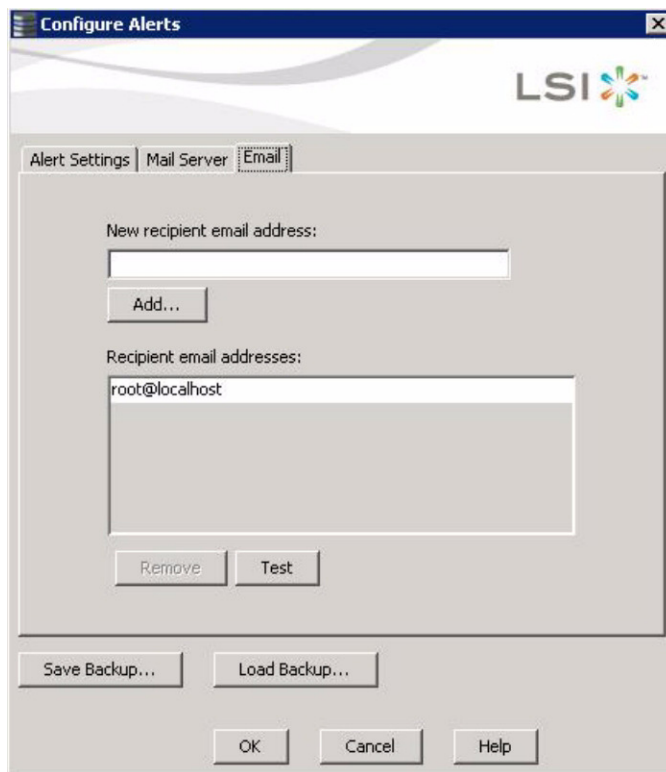


Figure 2-9 MegaRAID Storage Manager – Email Tab

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 2-10). Table 2-6 describes the function of the power supply LEDs




Figure 2-10 Power Supply LEDs

Table 2-6 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 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 LEDs




Figure 2-11 Fan Module LEDs

Like drives, fans also have a green/amber status LED and a blue locator LED. The green LED will be lit when a fan is operating normally. An amber LED will be lit if the fan is in fault. The locator LED is lit automatically when fans are set in a “Safe to Service” mode.

Fan Monitoring

Monitoring of the fan modules is done through the Fan Base Service page. Alert emails for fan failures is done using the preconfigured Platform Event Filters designated by Intel (see “Alerts” on page 60) through the BMC.

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.

Fan Base Service Page

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). The Fan Base Service Page is the customer interface for the fan baseboard, and can be accessed by connecting through the **MIS-S9D Proprietary Ethernet** port (Figure 2-14) and pointing a browser to:

`192.168.0.3/service.xml`. This page can also be reached through the JBOD I/O module ethernet port, using its DHCP address: `<DHCP address>/service.xml`.

MIS Server Chassis											
Cover is on, uptime 16868006(ms mod 2 ³²)				All FANS <input type="button" value="On"/>				<input type="button" value="Shutdown"/>			
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	R E A R	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				RPM : 11384 : 11014					D1	D4	D7
Power Supply 3 Good				FAN 1 Good <input type="button" value="On"/> <input type="button" value="Off"/> RPM : 11384 : 11014					D2	D5	D8
Power Supply 2 Good								StorBrick 0 Good	D0	D3	D6
Power Supply 1 o				FAN 0 Good <input type="button" value="On"/> <input type="button" value="Off"/> RPM : 10987 : 10763					D1	D4	D7
Power Supply 0 o									D2	D5	D8

Figure 2-12 Fan Base Service Page (Server)

The Fan Base Service Page layout is based on a top-down view of a chassis, which is why there are different layouts for the server and JBOD versions of the page. Notably, the server page features boot drives and server modules (Figure 2-12), where the JBOD page has a 9th StorBrick with capacity drives and JBOD I/O modules (Figure 2-13).

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.

MIS JBOD Chassis											
uptime 1478848(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 polling suspended	FAN 5 Good <input type="button" value="On"/> <input type="button" value="Off"/> RPM : 11889 : 10556		StorBrick 4 polling suspended	D0	D3	D6	F R O N T
	D7	D4	D1		D1	D4		D7			
	D6	D3	D0		D2	D5		D8			
	D8	D5	D2	StorBrick 6 polling suspended	FAN 4 Good <input type="button" value="On"/> <input type="button" value="Off"/> RPM : 11403 : 10761		StorBrick 3 polling suspended	D0	D3	D6	
	D7	D4	D1		D1	D4		D7			
	D6	D3	D0		D2	D5		D8			
	D8	D5	D2	StorBrick 7 polling suspended	FAN 3 Good <input type="button" value="On"/> <input type="button" value="Off"/> RPM : 10881 : 10666		StorBrick 2 polling suspended	D0	D3	D6	
	D7	D4	D1		D1	D4		D7			
	D6	D3	D0		D2	D5		D8			
	D8	D5	D2	StorBrick 8 polling suspended	FAN 2 Good <input type="button" value="On"/> <input type="button" value="Off"/> RPM : 11510 : 10687		StorBrick 1 polling suspended	D0	D3	D6	
	D7	D4	D1		D7	D4		D1			
	D6	D3	D0		D2	D5		D8			
	Power Supply 3 Good				FAN 1 Good <input type="button" value="On"/> <input type="button" value="Off"/> RPM : 11325 : 10723		StorBrick 0 polling suspended	D0	D3	D6	
	Power Supply 2 Good							D1	D4	D7	
	Power Supply 1 o				FAN 0 Good <input type="button" value="On"/> <input type="button" value="Off"/> RPM : 11126 : 10616			D2	D5	D8	
	Power Supply 0 o							D0	D3	D6	

Figure 2-13 Fan Base Service Page (JBOD)

MIS-S9D Proprietary Network Interface

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. The MIS-S9D interface is to be used when accessing the fan base for servicing, or when zoning. It is located at the front of the chassis at the upper right corner (Figure 2-14). The chassis must be slid out forward at least one inch in order to connect a network cable.



Figure 2-14 MIS-S9D Proprietary Network Interface

Ensure the MIS system is powered on. Use an Ethernet cable to connect a server/laptop to the MIS-S9D Proprietary Network Interface. 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 baseboard 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 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).

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

Setting the BMC IP Address

The BMC shares NIC 1 in a server module with the operating system. To access the BMC via IPMItool on a remote system or via the Integrated BMC Web Console, you must configure the BMC with its own IP address. The BMC is configured to obtain an IP address through DHCP by default.

1. The BMC IP Address is set during POST, by pressing **F2** to go into the BIOS.



Figure 2-15 BIOS Server Management Tab

2. Next, navigate to the **Server Management** tab (Figure 2-15), scroll down to **BMC LAN Configuration** and press **Enter**.
3. Scroll down to **IP Source** and select **Static**.
4. Configure the **IP Address**, **Subnet Mask**, and **Gateway IP** as required.
5. Press **F10** to save your settings.

Important: Keep a record of the **DHCP Host Name** and assigned **IP address**. These will be needed later when you are trying to connect to the Integrated BMC Web Console.

Setting Up User Accounts for the BMC

To be able to log into the BMC, user accounts must be configured through the BIOS.

1. In the BIOS, scroll down to **User Configuration** and change the **User ID** name as needed.

Note: You cannot change anonymous or root, nor can you log into the Integrated BMC Web Console using an anonymous account. You must either configure the root account or create a new user account.

2. Scroll down to **User Password**. Set the password (must be entered twice). User names and passwords are cAsE sEnSiTiVe.
3. When finished, press **F10** to save, and **Esc** to exit. The system will reboot with the new LAN settings.

Help Files

To view the **Help** files within the BMC, Click the **Help** button in the upper right corner (Figure 2-16).




Figure 2-16 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.*

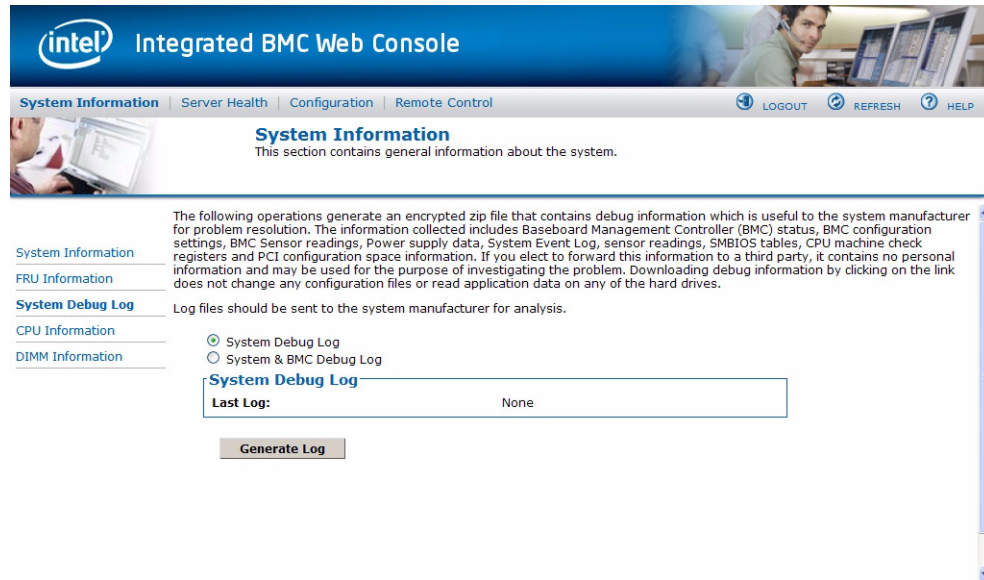


Figure 2-17 BMC Web Console – System Debug Log

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.**

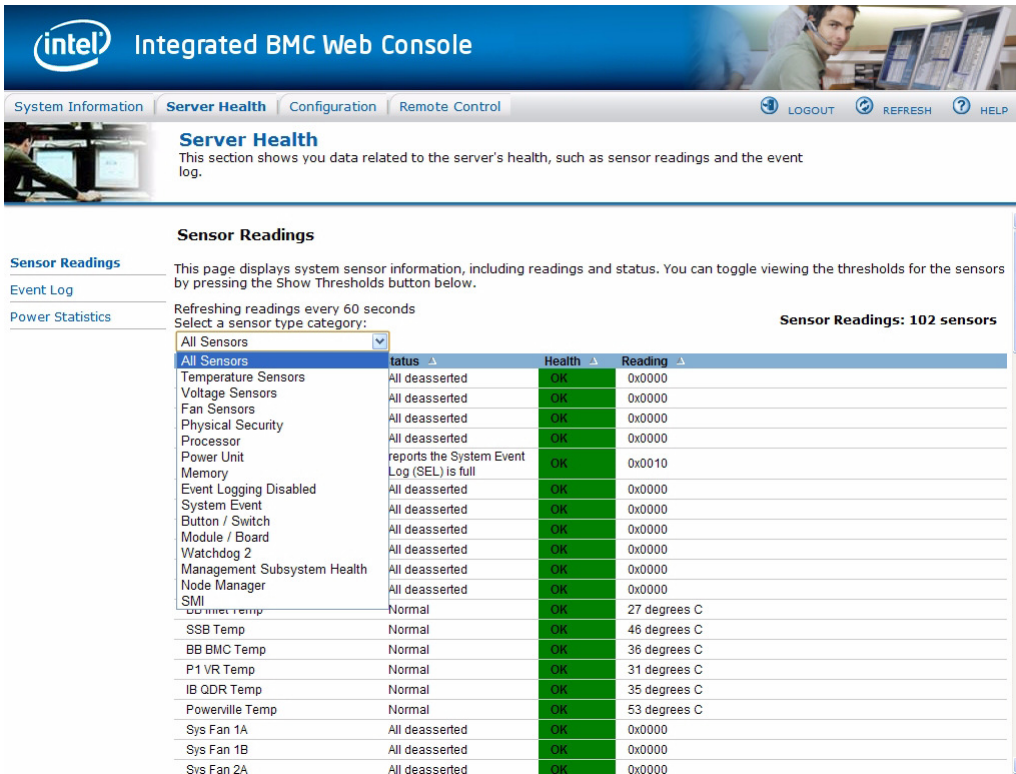


Figure 2-18 BMC Web Console – Server Health

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**. The fan base firmware monitors all Power Supply Units (PSUs) installed in the system, and generates SEL entries regardless of whether or not the alerts are enabled. This functionality cannot be turned off.

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). Items that are monitored using the SEL, such as power supply status, can generate alerts through the BMC Integrated Web Console (see "Alerts" on page 59).

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 IPMI tools for your operating system (IPMItool – Linux-based systems, IPMIutil – Windows)

Note: All BIOS sensors and some BMC and ME sensors are "event-only"; meaning that they are not readable using the IPMI `Get Sensor Reading` command but rather are used just for event logging purposes.

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).

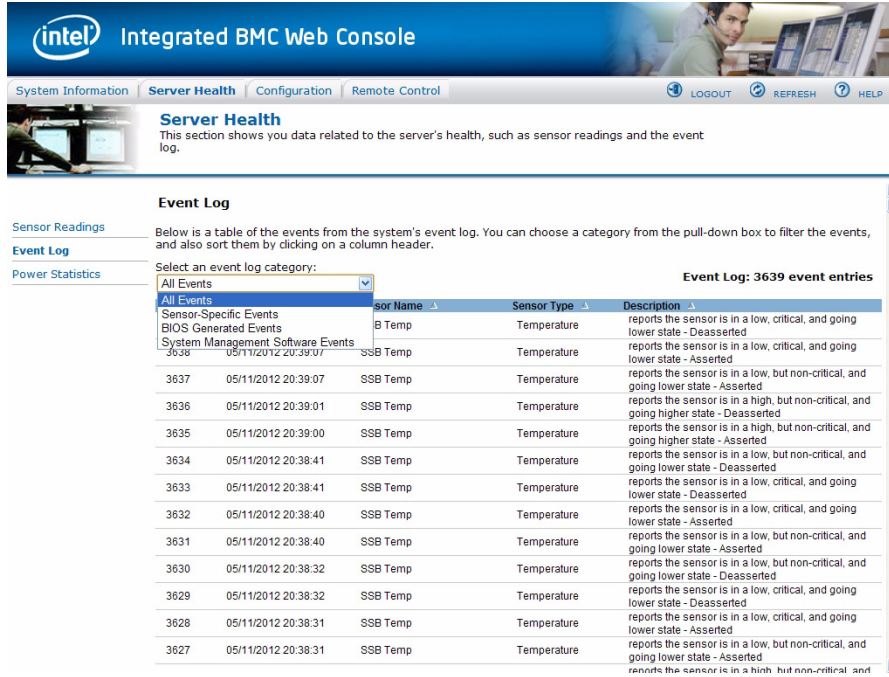


Figure 2-19 BMC Web Console – Event Log

Power Statistics

Not used by MIS platforms. See “Power Supply Monitoring” on page 46 instead.

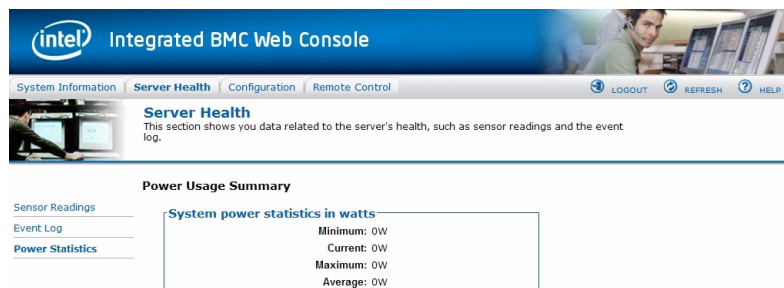


Figure 2-20 BMC Web Console – Power Statistics (not used)

Configuration Tab

The Configuration tab of the BMC Web Console is used to configure various settings, such as alerts, users accounts for the BMC, and network settings. It contains the following menu options in the left navigation pane: **IPv4 Network**, **IPv6 Network** (not used), **Users**, **Login**, **LDAP**, **VLAN**, **SSL**, **Remote Session**, **Mouse Mode**, **Keyboard Macros**, **Alerts**, **Alert Email**, and **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.

The screenshot displays the BMC Web Console interface. At the top, there is a navigation bar with the Intel logo and the text 'Integrated BMC Web Console'. Below this, a secondary navigation bar includes 'System Information', 'Server Health', 'Configuration' (which is highlighted), and 'Remote Control'. To the right of this bar are links for 'LOGOUT', 'REFRESH', and 'HELP'. Below the navigation bar, a sub-header reads 'Configuration' followed by the instruction 'Use these pages to configure various settings, such as alerts, users, or network.'

The main content area is titled 'Alerts' and contains the following configuration options:

- Alerts**: Configure which system events generate Alerts and the external network destinations they should be sent to.
- 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: 100.38.137.30
 - Email Send Email to: mis-ad@sgi.com
- Alert Destination #2:**
 - SNMP Send SNMP Alerts to IP: 0.0.0.0
 - Email Send Email to:
- Save** **Send Test Alerts**

Figure 2-21 BMC Web Console – Alerts Page

To receive alerts for power supplies, see “Power Supply Monitoring” on page 46. 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.

Globally Enable Platform Event Filtering

Global control for enabling or disabling platform event filtering. This can be used to prevent sending alerts until you have fully specified your desired alerting policies.

Warning: Do NOT disable except under explicit instruction from SGI Technical Support. When filtering is globally disabled through this setting, alerts (such as power supply faults) will *not* be sent.

Select Events

Select one or more system events that will trigger an Alert. These events correspond to the 13 IPMI preconfigured Platform Event Filters (PEFs) set by Intel. This area only allows the selection of those 13 PEF filters configured by Intel: **Temperature Sensor Out of Range, System Restart, Fan Failure, Power Supply Failure¹, BIOS: POST Error Code, Node Manager Exception², Watchdog Timer, Voltage Sensor Out of Range, Chassis Intrusion, Memory Error, FRB Failure, Hard Drive Failure³.**

Warning: Clearing all events disables Alerts.

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:

-
- ¹ Not used by SGI – Power Supply failure is reported through the BMC, but Power Supply Alerts are not configured through this interface. Instead, power supply failure is reported by configuring Platform Event Filters 5, 19, and 20 (see “Power Supply Monitoring” on page 46).
 - ² Requires configuration of Node Manager.
 - ³ Boot drives only. Capacity Drives are monitored using MegaRAID Storage Manager software (“Creating the Drive Groups in MegaRAID (CentOS, Linux & Windows)” on page 102).

- 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.

Alert Email

The screenshot shows the BMC Integrated Web Console interface. At the top, there is a blue header with the Intel logo and the text "Integrated BMC Web Console". Below the header is a navigation bar with tabs for "System Information", "Server Health", "Configuration" (which is active), and "Remote Control". To the right of the navigation bar are buttons for "LOGOUT", "REFRESH", and "HELP". Below the navigation bar is a sub-header for "Configuration" with the text "Use these pages to configure various settings, such as alerts, users, or network." Below this is a list of configuration options on the left side, including "IPv4 Network", "IPv6 Network", "Users", "Login", "LDAP", "VLAN", "SSL", "Remote Session", "Mouse Mode", "Keyboard Macros", "Alerts", "Alert Email" (which is highlighted), and "Node Manager". The main content area is titled "Alert Email Settings" and contains the following configuration fields:

- Alert Email Settings**: Configure how Alerts are sent by email to an external SMTP Mailserver. Each LAN Channel has a separate configuration.
- LAN Channel:** Baseboard Mgmt (selected from a drop-down menu)
- SMTP Server IP:** 100.38.137.29
- Sender Address:** curtis@sgi.com
- Local Hostname:** test-toast

At the bottom of the configuration area is a "Save" button.

Figure 2-22 BMC Web Console – Alert Email Settings

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).

In this example, the mail server’s IP address is 100.38.137.29. When the **Send Test Alert** button is pushed, it creates an alert email from the Sender Address (in this example, Figure 2-23, `curtis@sgi.com`), with the heading “**Alert from <Local Hostname>**” (in this example, 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 2-23 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 (**RT**), 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 2-23), 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**).

```
-----Original Message-----
From: AlertEmail [mailto: ]
Sent: Tuesday, July 23, 2013 7:05 PM
To:
Subject: Alert from <dlv1>

Event that generated this alert:
RID:0013 TS:07/24/2013 01:05:14 SN:MIS PSU2 STATUS
ST:Power Supply ED:Input Lost ET:Asserted EC:OK
RID:0013 RT:02 TS:51EF284A GID:002E ER:04 ST:08 S#:72
ET:6F ED:03 FF FF EX:00 FF FF FF FF FF FF FF FF
```

Figure 2-24 Alert Email – Power Cord Pulled from Power Supply Module

```
-----Original Message-----
From: AlertEmail [mailto: ]
Sent: Tuesday, July 23, 2013 6:12 PM
To:
Subject: Alert from <dlv1>

Event that generated this alert:
RID:0011 TS:07/24/2013 00:11:45 SN:MIS PSU1 STATUS ST:Power
Supply ED:Presence detected ET:Deasserted EC:OK
RID:0011 RT:02 TS:51EF1BC1 GID:002E ER:04 ST:08 S#:71 ET:EF
ED:00 FF FF EX:00 FF FF FF FF FF FF FF FF
```

Figure 2-25 Alert Email – Power Module Pulled From Chassis

Server Power Control

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

Table 2-7 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.

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.

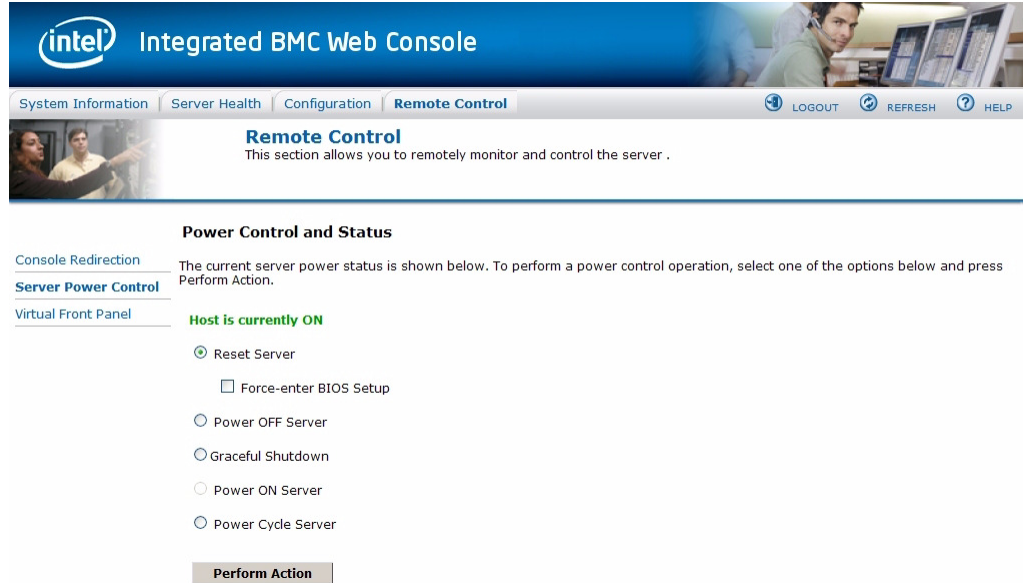


Figure 2-26 BMC Web Console – Power Control and Status

Virtual Front Panel

The Virtual Front Panel is a graphic representation of the front panel, providing remote functionality virtually. From here you can turn on the blue chassis identification LED to locate the chassis in the rack (Figure 2-27).

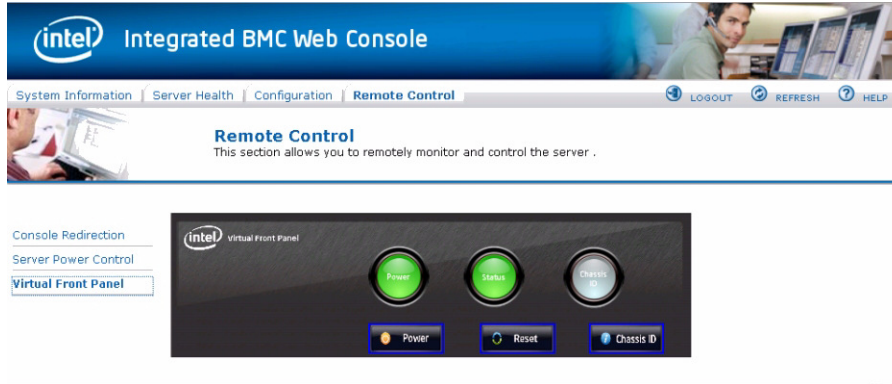


Figure 2-27 BMC Web Console – Virtual Front Panel

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).

The BMC supports LAN Alerting in the form of SNMP Traps that follows the Platform Event Trap format. SNMP Traps are typically sent as unreliable datagrams. However, IPMI includes PET Acknowledge and retry options that allow an IPMI-aware remote application to provide a positive acknowledge that the trap was received. Alert-Over-LAN notifies remote system management applications about PEF selected events, regardless of the state of the server's operating system. LAN alerts can be sent over any of the LAN channels supported by a system.

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://ipmiutil.sourceforge.net/>.

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.

Enabling IPMI on a Linux System

To use IPMI on a Linux system, the `ipmi` service must be running. Use the `# service ipmi status` command to determine if the service is running.

If the service is running, the output will be similar to:

```
# service ipmi status
ipmi_msghandler module loaded.
ipmi_si module loaded.
ipmi_devintf module loaded.
/dev/ipmi0 exists.
```

If the service is not running, use the `# ipmi service start` command to start the service.

```
# service ipmi start
Starting ipmi drivers: [ OK ]
```

Accessing IPMI on a Linux System

On a Linux system, use the `# ipmitool` command to access IPMItool.

On a local system, simply use the `ipmitool` command with the command it should run:

```
# ipmitool <command>
```

For example:

```
# ipmitool power status
Chassis Power is on
```

From a remote system, you must also specify a host, user, and password to access the system:

```
# ipmitool -H <host> -U <user> -P <password> power status
```

For example:

```
# ipmitool -H 137.38.64.3 -U admin -P Password1 power status
Chassis Power is off
```

Note: To access IPMI from a remote system, the BMC on the local system must be configured for network access. (Refer to “Setting the BMC IP Address” on page 53 for the procedure to enable network access.)

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.

The following sections take you through downloading zoning software, zoning, and possible RAID configurations. Both Zones 1.4.2 and the CLI Zoning Tool 1.4 have had significant updates from version 1.0.

- “Downloading & Installing Software,” on page 72,
- “Before Updating from Zones 1.0,” on page 73,
- “Zones 1.4.2 for Linux & Windows,” on page 77,
- “Creating the Drive Groups in MegaRAID (Linux & Windows),” on page 91,
- “Formatting the Drives using YaST2 in Linux,” on page 93,
- “Windows Server Manager,” on page 98,
- “CLI Zoning Tool, version 1.4,” on page 105,
- and “Disk RAID Support,” on page 132.

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 HBA in what server can see which drive sets. When open zoning is enabled, all the SAS connections can see all of the drives. For dual-ported SAS drives, **both** ports 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.

Phy-based Zoning is implemented on the MIS Server. 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 Linux and 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 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 cable connected to the MIS-S9D proprietary network interface (Figure 3-43).

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 (Linux & Windows),” on page 91 for more information).

Section Guide

Inside the first section you will find instructions on downloading and installing the software needed. The next section gives instructions on how to use the software to zone MIS Platforms.

- “Downloading & Installing Software,” on page 72,

- “Required Downloads,” on page 73,
- “CLI Zoning Tool, version 1.4,” on page 105,
- “Installing MegaRAID Storage Manager for Linux,” on page 74,
- “Zones 1.4.2 for Linux,” on page 75,
- “Installing CLI Zoning Tool on Windows Host,” on page 106,
- “MegaRAID Storage Manager for Windows,” on page 75,
- “Python for Windows,” on page 76,
- “Zones 1.4.2 for Windows,” on page 76.

The second section gives step-by-step instructions for using the Zones application.

- “Zones 1.4.2 for Linux & Windows,” on page 77,
 - “Zoning MIS Server Platforms,” on page 77,
 - “Opening a Session,” on page 79,
 - “144 Drives Setup,” on page 81,
 - “Saving Zoning Session,” on page 83,
 - “Zoning MIS JBOD Platforms,” on page 84,
 - “Adapter Tab Assignment,” on page 88
 - “Saving a Session,” on page 89
 - “Downloading a Session,” on page 89
 - “Loading a CSV File to Zones,” on page 90.

The third, fourth and fifth sections deal with running the software, including step-by-step instructions on “Creating the Drive Groups in MegaRAID (Linux & Windows),” on page 91, “Formatting the Drives using YaST2 in Linux,” on page 93, “Windows Server Manager,” on page 98, to finish up the necessary steps for configuring MIS platforms.

- “Creating the Drive Groups in MegaRAID (Linux & Windows),” on page 91,
- “Formatting the Drives using YaST2 in Linux,” on page 93,
 - “Verify Drives Seen in YaST2,” on page 95,
 - “Partitioning the Drives in YaST2,” on page 96
- “Windows Server Manager,” on page 98

- “Verify Drives Seen in Windows Server Manager,” on page 98
- “Formatting the Drives in Windows Server Manager,” on page 99

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 “Python for Windows,” on page 76).

- “CLI Zoning Tool, version 1.4,” on page 105,
 - “Preparing to Zone using the CLI Zoning Tool,” on page 110
 - “CLI Zoning Tool Main Menu,” on page 114,
 - “Version 1.4.0 Default Menu,” on page 115,
 - “Editing the .csv File for the CLI Zoning Tool,” on page 116.

Finally, “Disk RAID Support” explains the different RAID arrays available and the benefits and drawbacks of each. There are special considerations in creating RAID arrays for use on StorBricks, that is, there are RAID configurations, namely 6+2 and 7+2, that ensures there is no single point of failure on the StorBricks within the MIS system. “RAID Configuration Notes” details how to manage those concerns.

- “Disk RAID Support,” on page 132,
 - “RAID Configuration Notes,” on page 133,
 - “Spare Drives,” on page 135.

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 may depend on the operating system on the MIS Server. MIS Servers and JBODs can be 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 3-43) in order to zone the machine in question. Otherwise, the operating system installed on the boot drives in is what decides which software to download and use.

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 networked 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

MIS Linux and Linux-based systems (Red Hat) – “Installing MegaRAID Storage Manager for Linux”, and “Zones 1.4.2 for Linux”

MIS Windows systems – “MegaRAID Storage Manager for Windows”, “Python for Windows”, and “Zones 1.4.2 for Windows”

No MIS Operating System – “CLI Zoning Tool, version 1.4” on host machine¹ with Python 2.6 or 2.7 installed.

Before Updating from Zones 1.0

Linux

If you are updating from Zones 1.0, you should have an `/opt/Zones/` directory containing the program and session files. Simply rename that folder `/opt/Zones_old_1.0/` before installing Zones 1.3. This will keep all your previous session files available for future use, should you need them.

Windows

There are no necessary steps necessary to take, prior to installing Zones 1.4.2 for Windows. Zones 1.0 files appear under the directory `C:\Zones\`.

¹ If the host laptop/server running the CLI Zoning Tool 1.4 is a Windows machine, Python for Windows is also required.

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.
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 to a directory that is easy to navigate to using a command line interface (e.g., /home/<user_name>).
4. Change directory to where you have saved your .tar file (e.g., # cd /home/<user_name>)
5. Untar the .tar file, # tar -zxvf <filename.tar> (e.g., -zxvf 13.04.03.01_Linux(x86)_MSM.tar)
6. You should see an output of files and a new directory called disk. Change to the disk directory (i.e., cd disk)
7. Run the installation by typing ./install.
8. 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
9. Select option 1 - Complete.

10. 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.

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.
2. Download Zones 1.4.2 for Linux.
3. Unzip `Zones.zip` from the command line (e.g., `# unzip Zones.zip`). A `Zones` folder will appear.
4. 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.*

5. Change directory into the new `Zones` folder (e.g., `# cd /opt/Zones`)
6. 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.

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://lsi.com> and search for MegaRAID Storage Manager (on the search page, you can refine the search by choosing **File Type > Management Software and Tools**), or contact SGI Customer Support for the latest working version.
2. Select the latest Windows version for download, and click **Accept** on the license agreement page.
3. Download and install.

MegaRAID Storage Manager GUI for Windows will appear.

Python for Windows

Both Zones 1.4.2 for Windows and the CLI Zoning Tool 1.4 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.
2. Select the version of Python required, or contact SGI Customer Support for the latest working version.
3. Download and start the installation.
4. The installation will ask which directory in which to install python (the default is `c:\Python##\` where ## is the version number).
5. On the **Customize Python** pane, click **Next**.
6. 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.

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>
2. Download Zones 1.4.2 for Windows.
3. Extract the `Zones.zip` files to the destination folder `c:\Program Files (x86)\`. This will create the directory `c:\Program Files (x86)\Zones\`.
4. In the directory `c:\Program Files (x86)\Zones\`, look for the extracted file `pygtk-all-in-one-2.22.6.win32-py2.7.msi`.
5. Run file extraction by double-clicking on `pygtk-all-in-one-2.22.6.win32-py2.7.msi`.

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`.

Zones 1.4.2 for Linux & Windows

This updated release of Zones features the familiar graphic user interface, now streamlined to make zoning faster and easier to execute. Both Zones 1.4.2 for Linux and Zones 1.4.2 for Windows both function the same, now with only a slight difference in appearance. The following instructions work for both Linux 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 Windows,” or “Starting Zones 1.4.2 for Linux,” on page 75.)

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 who can see what. When **open zoning** is enabled, all the SAS connections can see all of the drives. For dual-ported SAS drives, **both** ports will be exposed, so the drives will show up **twice**. This situation will cause conflict between the HBAs.

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. (Zones will not display the 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

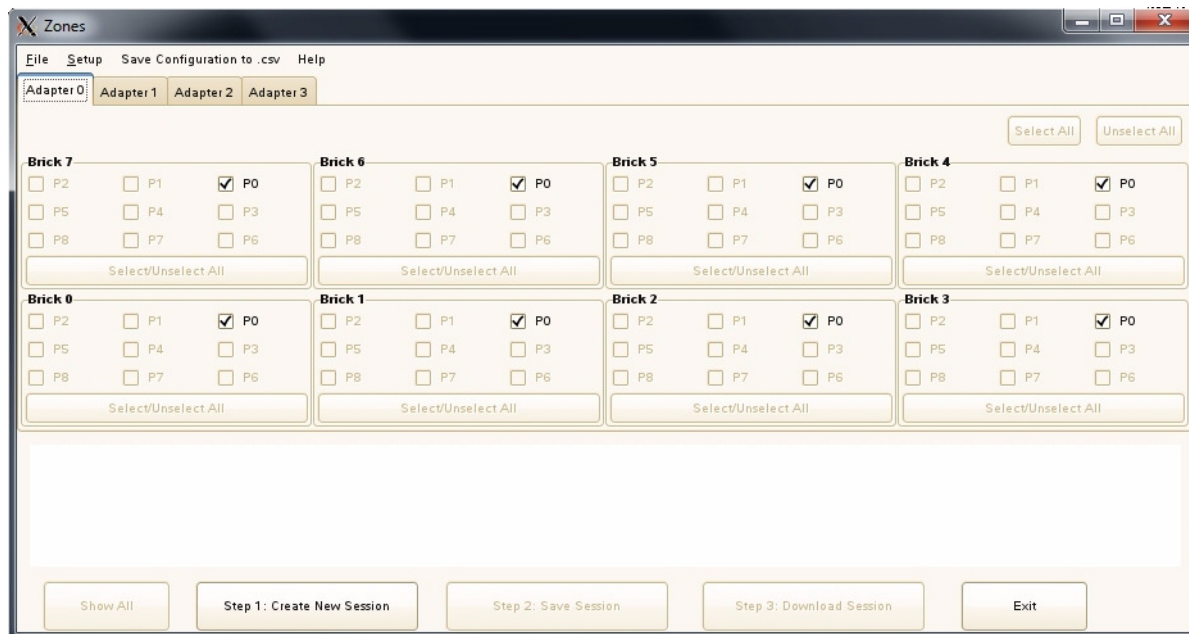


Figure 3-1 Zones 1.4.2 – Current Zoning Configuration

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 news 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 90.)

Opening a Session

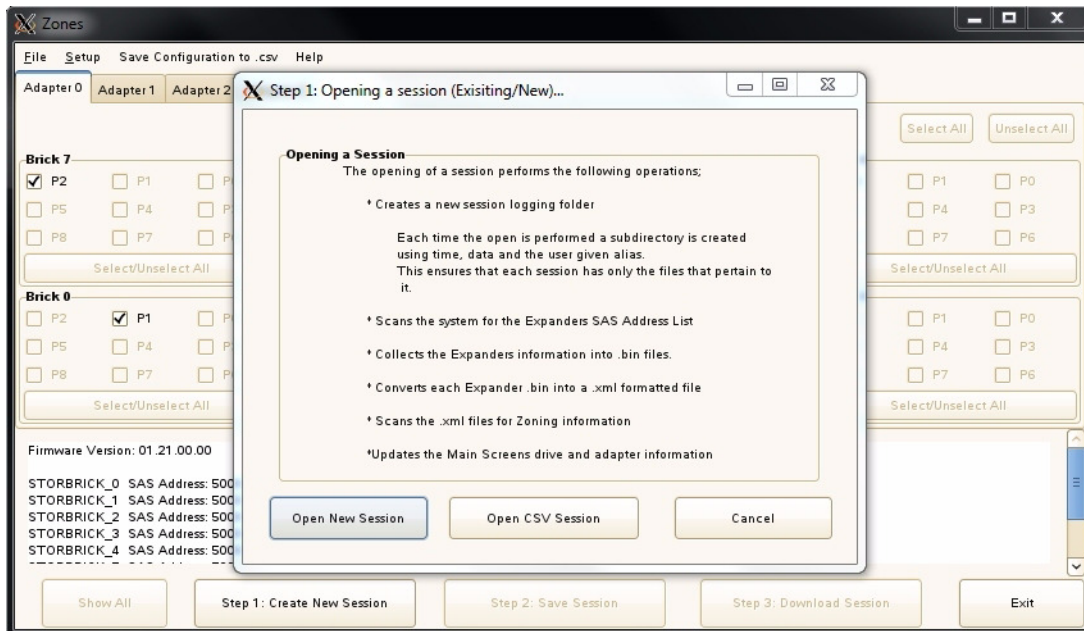


Figure 3-3 Zones 1.4.2 – Step 1: Opening a Session

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**.

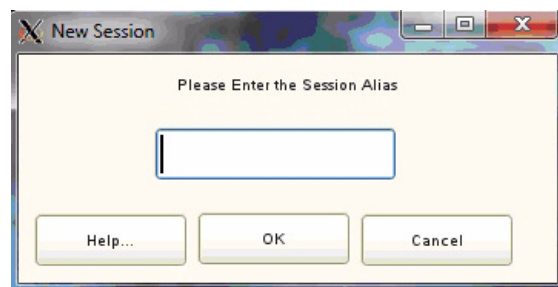


Figure 3-4 Zones 1.4.2 – New Session Alias

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).

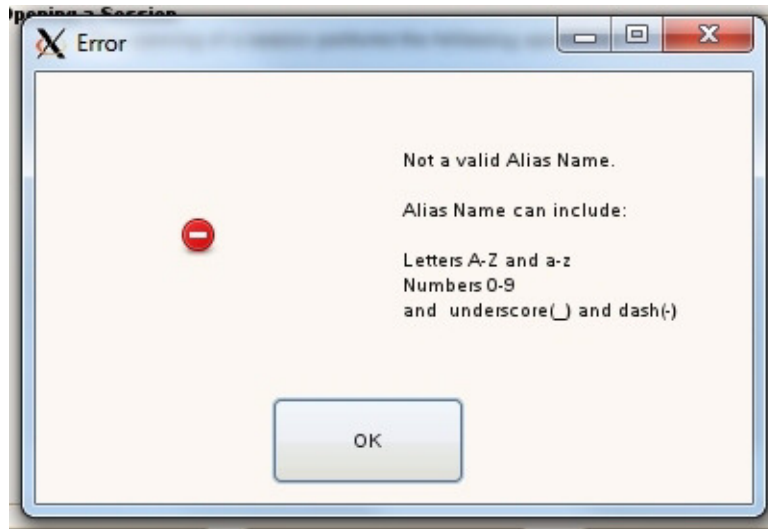


Figure 3-5 Zones 1.4.2 Warning Box – Alias Rules

The Zones Interface will now display the number of adapters present in the system for that StorBrick. In this example (Figure 3-6), all four adapters are present, which can be seen by the four tabs at the top, labeled **Adapter 0**, **Adapter 1**, **Adapter 2**, and **Adapter 3**. This could be the result of having a dual server machine with two adapters per server or a single server with four adapters. 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 the dual server configuration, Zones is unable to determine which adapter belongs to which server. However, this can be used to an advantage. The user can correlate 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 3-7). 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**.

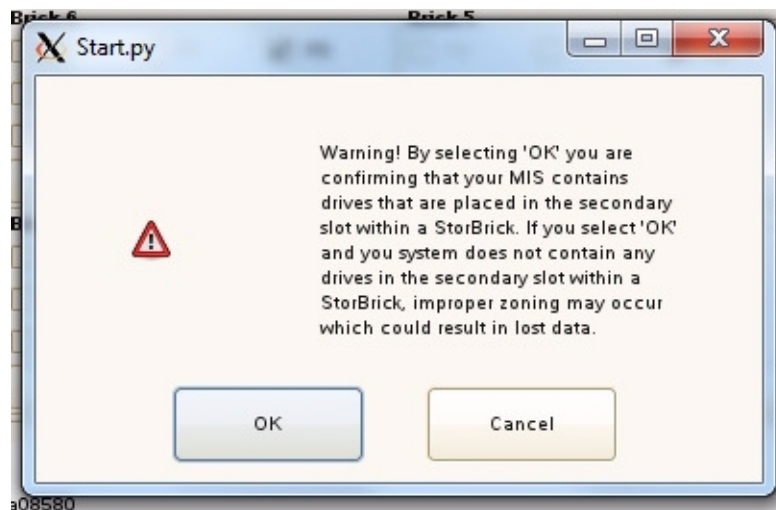


Figure 3-8 Zones 1.4.2 – Warning 144 Drive Configuration

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

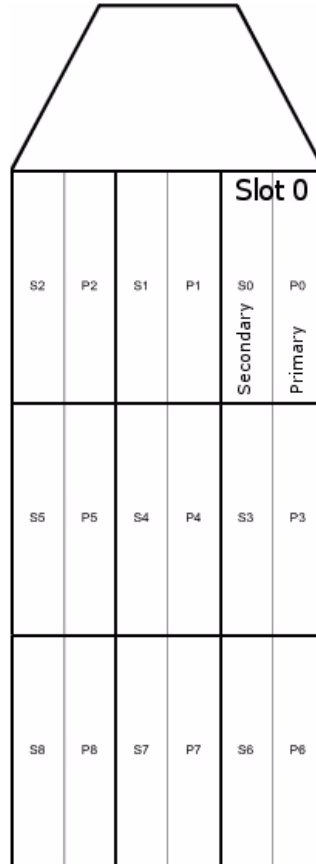


Figure 3-9 StorBrick Diagram for 144-Drive Systems

Saving Zoning Session

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 and forces a reset on each of the selected StorBricks' expander. When complete, the **Brick Selection** window will disappear.

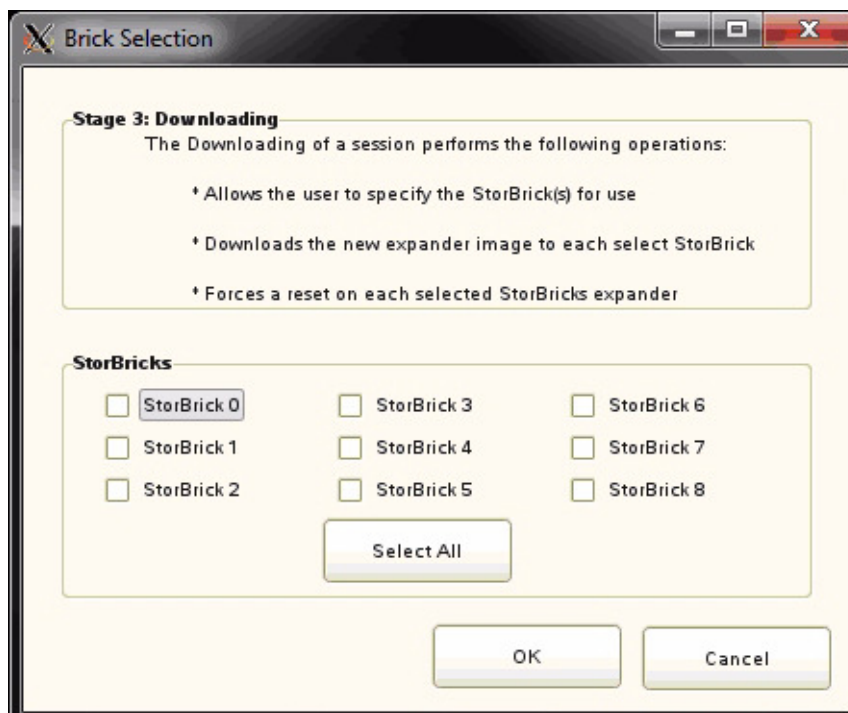


Figure 3-10 Zones 1.4.2 – Brick Selection Window

To complete the zoning process, the machine 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.

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**.

If the JBOD is connected to a different kind of server, it is zoned using the Zones JBOD Zoning Tool, `JGUI.py`. This program can be accessed through Linux by changing to the `/opt/Zones/JBOD` directory and typing `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).

Zones JBOD Interface

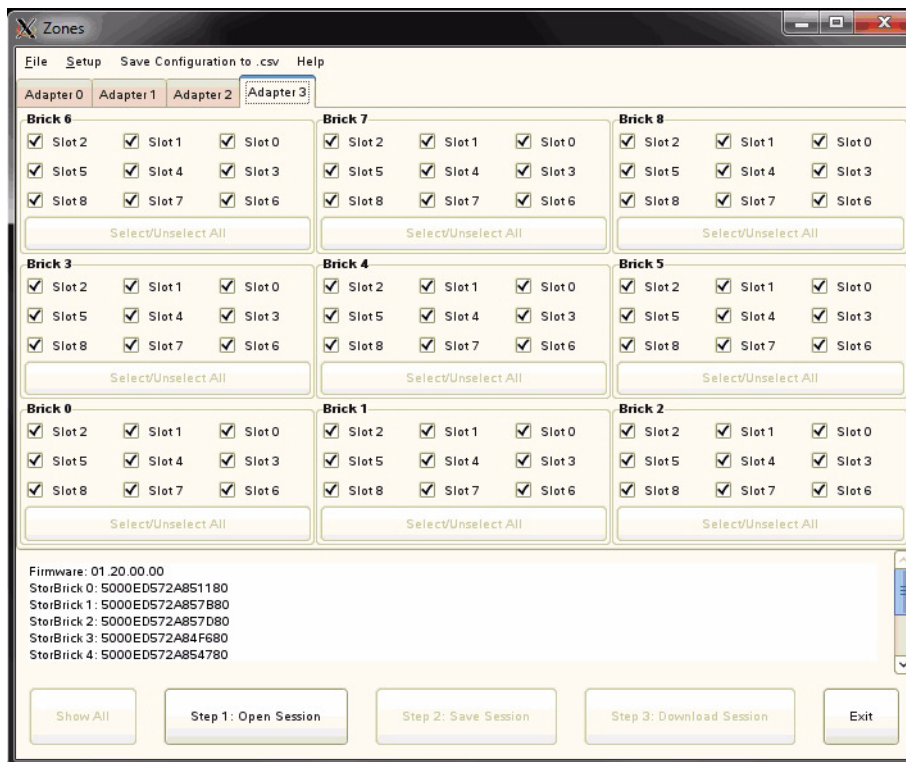


Figure 3-11 Zones 1.4.2 – JBOD Zoning Interface

Opening a Session

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

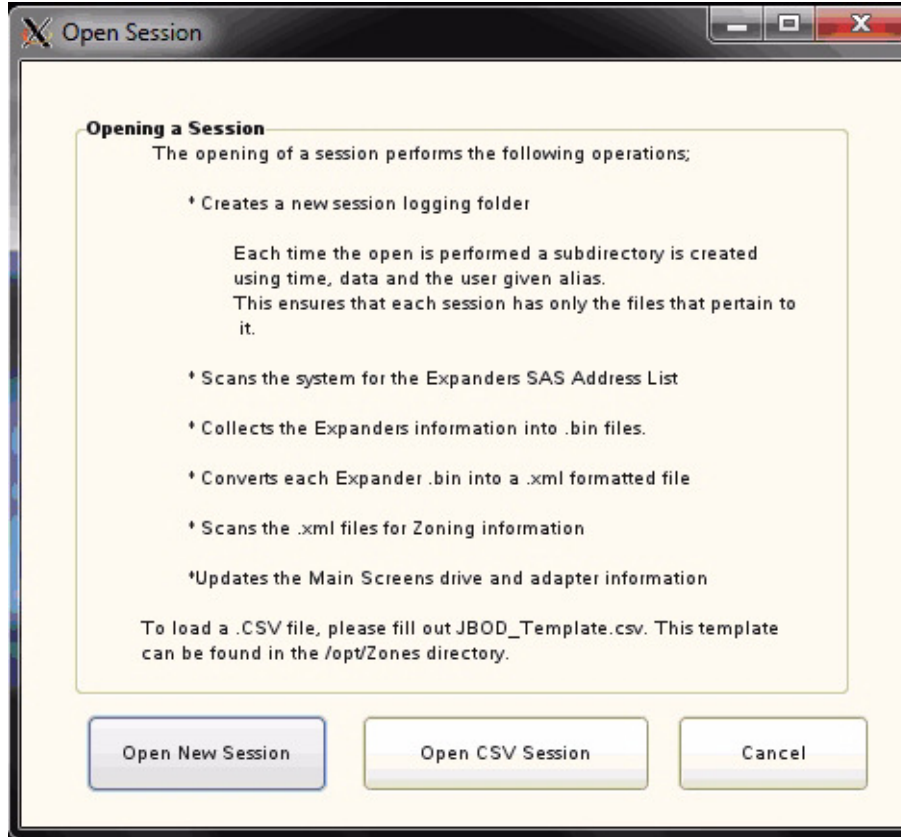


Figure 3-12 Zones 1.4.2 – JBOD Open Session Window

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 90 for more information.) Otherwise, click on **Open a New Session**.

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).

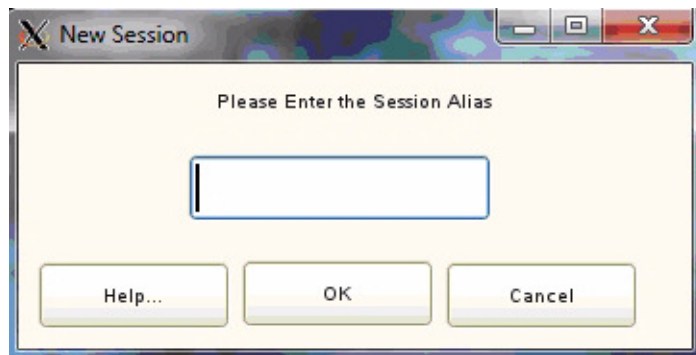


Figure 3-13 Zones 1.4.2 – JBOD New Session Alias

Phy-based zoning allows you to split the drives **between two I/O CARDS** (Figure 1-11 on page 9), 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).

Adapter Tab Assignment

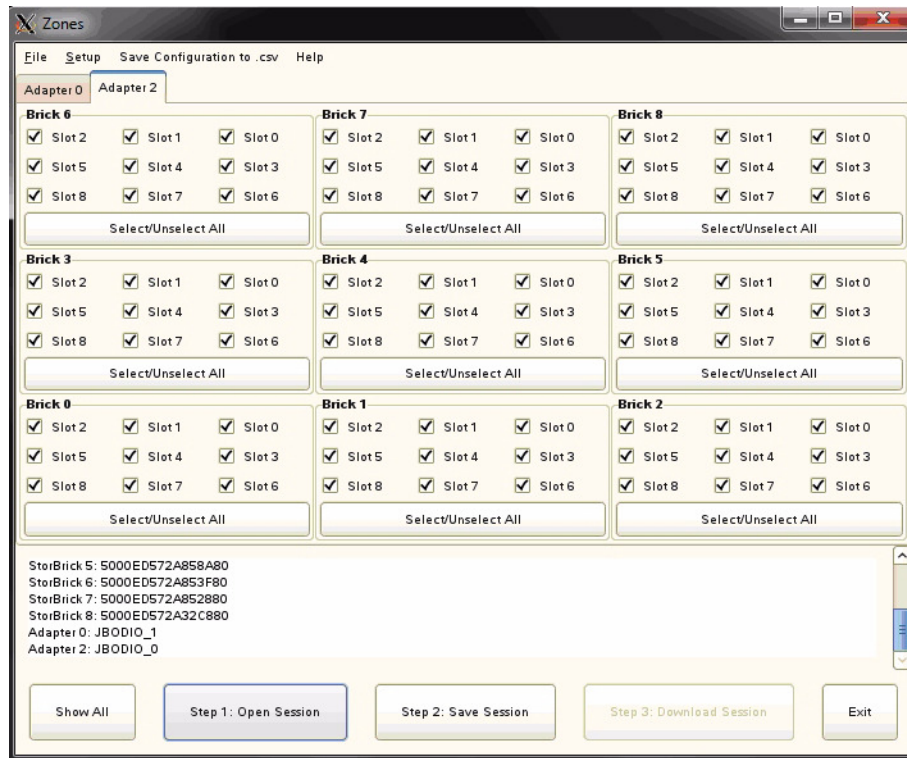


Figure 3-14 Zones 1.4.2 – JBOD Phy-based Zoning Interface

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.

Saving a Session

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 file, and enables the **Step 3: Download Session** button.

Downloading a Session

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 and forces a reset on each of the selected StorBricks' expander. When complete, the **Brick Selection** window will disappear.

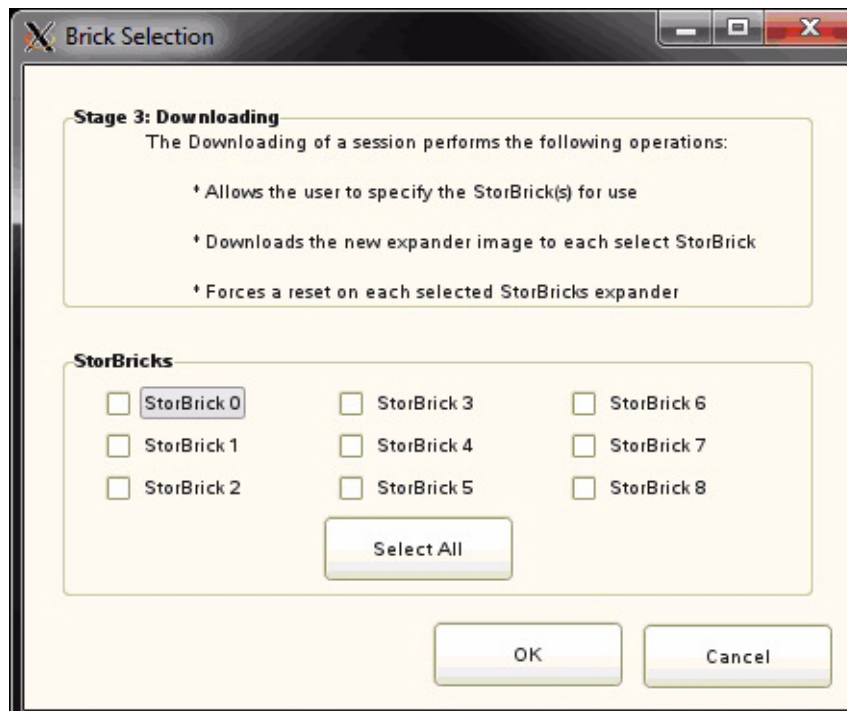


Figure 3-15 Zones 1.4.2 – JBOD Brick Selection Window

Important: 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 JBOD² 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 78).

In the **Open Session** window (Figure 3-3 on page 79) select **Open CSV Session**. This will bring up the file selection window (Figure 3-16).

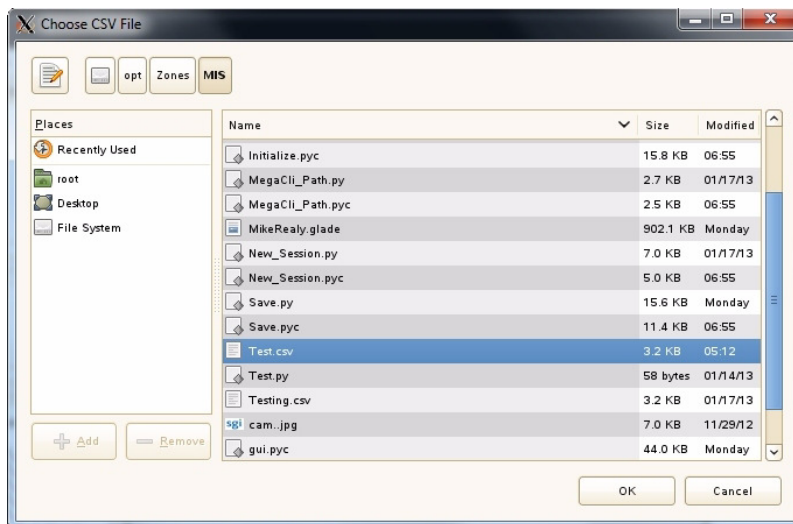


Figure 3-16 Zones 1.4.2 – Choose CSV File Window

² Loading a CSV file always results in Phy-based zoning.

Browse to the location to where the desired CSV file is stored, select it, and click **OK**. A **Warning** box, asking you if you are sure you want to continue (Figure 3-17).




Figure 3-17 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 flash the firmware files and take effect the new settings for each StorBrick.

Creating the Drive Groups in MegaRAID (Linux & Windows)

Power on the machine to complete power-cycle. Open MegaRAID. Right-click on the expander and select **Create a Virtual Drive** (Figure 3-18).




Figure 3-18 MegaRAID – Create a Virtual Drive

A screen will pop up asking you to choose **Simple** or **Advanced** (Figure 3-19). 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.

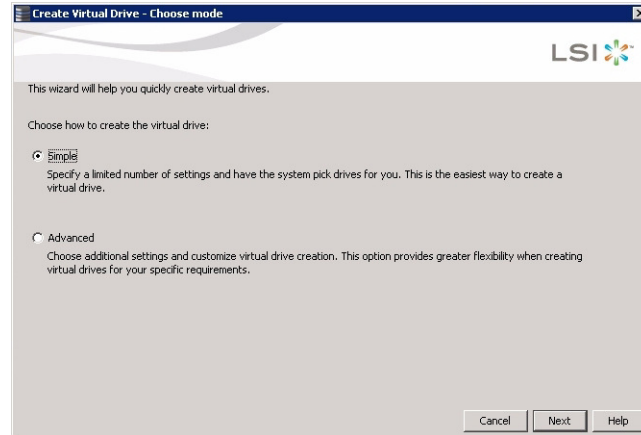


Figure 3-19 MegaRAID – Create Virtual Drive Mode

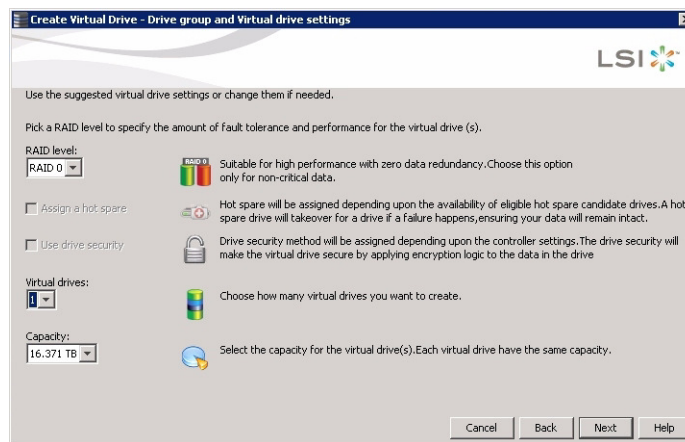


Figure 3-20 Create Virtual Drive – Drive Group Settings

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 3-21).

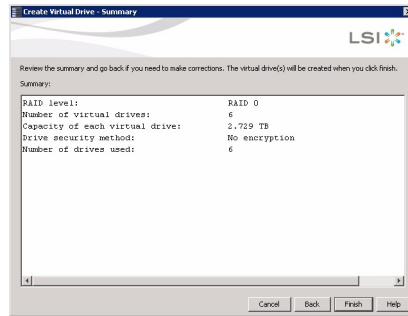


Figure 3-21 Create Virtual Drive – Summary

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

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.

1. Issue YaST2 command (i.e., # `yast2`) to launch the YaST 2 Server Manager GUI (Figure 3-22).




Figure 3-22 Yast2 Server Manager GUI

2. Double-click **Partitioner** to launch.
3. A warning message will appear (Figure 3-23). Click **Yes**.



Figure 3-23 YaST2 – Warning Message

Verify Drives Seen in YaST2

4. Verify that all of your disks have appeared under **Hard Disks** (Figure 3-24).

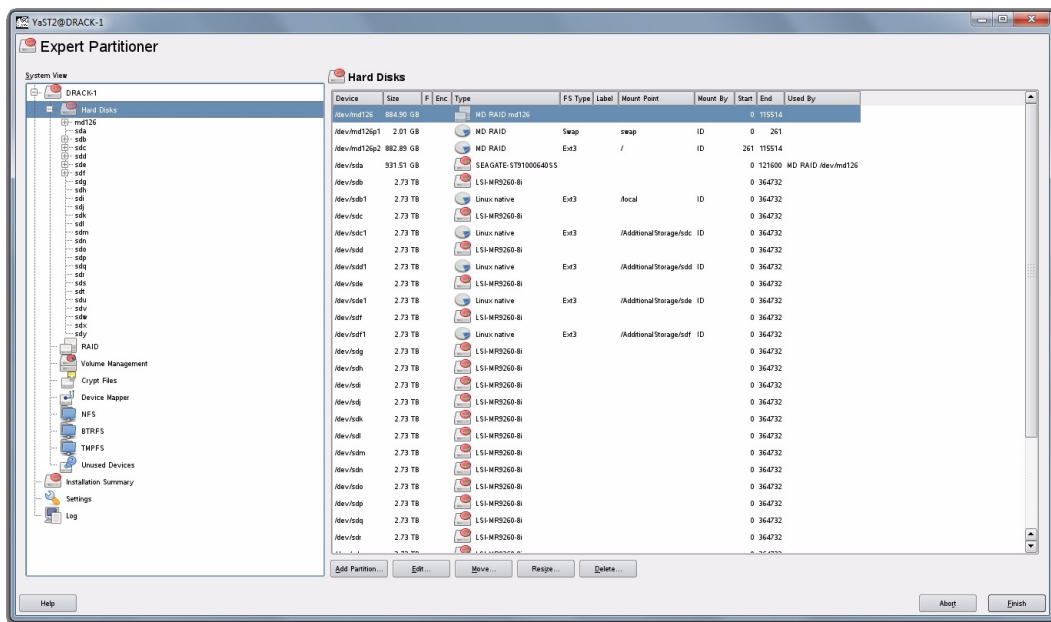


Figure 3-24 YaST2 – Drives Have Appeared

Partitioning the Drives in YaST2

- Under **Hard Disks**, select the disk you would like to partition and click **Add** at the bottom of the screen (Figure 3-25).




Figure 3-25 YaST2 – Add Button

- Select the partition size (Figure 3-26) and click **Next**.




Figure 3-26 YaST2 – Select Partition Size

- Format the partition using `ext3`, mount the disk to your desired folder, and click **Finish** (Figure 3-27).



Figure 3-27 YaST2 – Format & Mount the Drive

- Verify the partition shows up (Figure 3-28) and click **Next**.

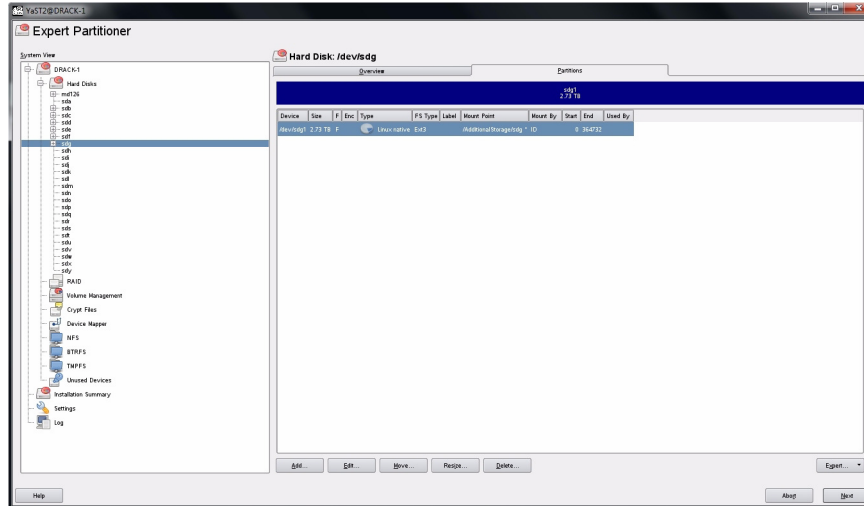


Figure 3-28 YaST2 – Check for Partition

9. Click **Finish** (Figure 3-29).



Figure 3-29 YaST – Click Finish

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




Figure 3-30 YaST2 – Disk Mounting (in process)

Once the disk is mounted (Figure 3-30), the system will return you to the beginning YaST2 GUI.

Windows Server Manager

Verify Drives Seen in Windows Server Manager

In Windows, open Server Manager (Figure 3-31). Verify that the only disks the system sees are system drives, and that they are labeled C : \.

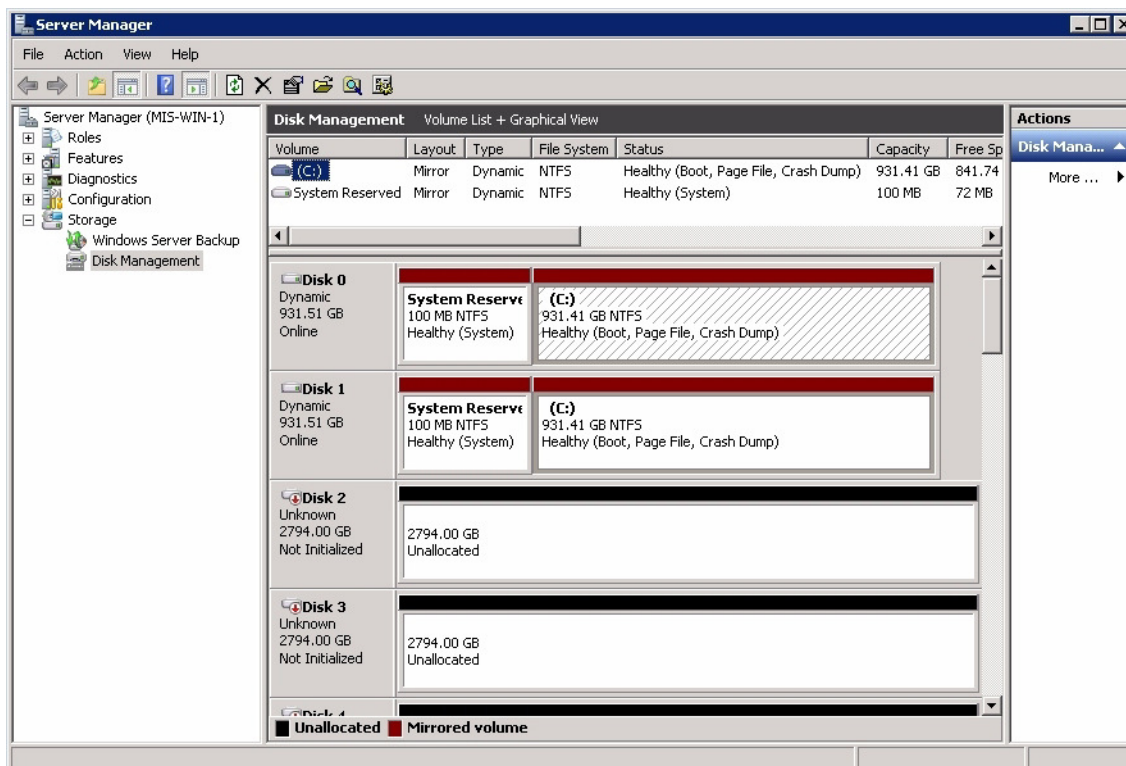


Figure 3-31 Windows Server Manager – Disk Management

Note: Unconfiguring drives removes them from the system.

Formatting the Drives in Windows Server Manager

Drives are formatted using Windows Server Manager (Figure 3-32). 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.

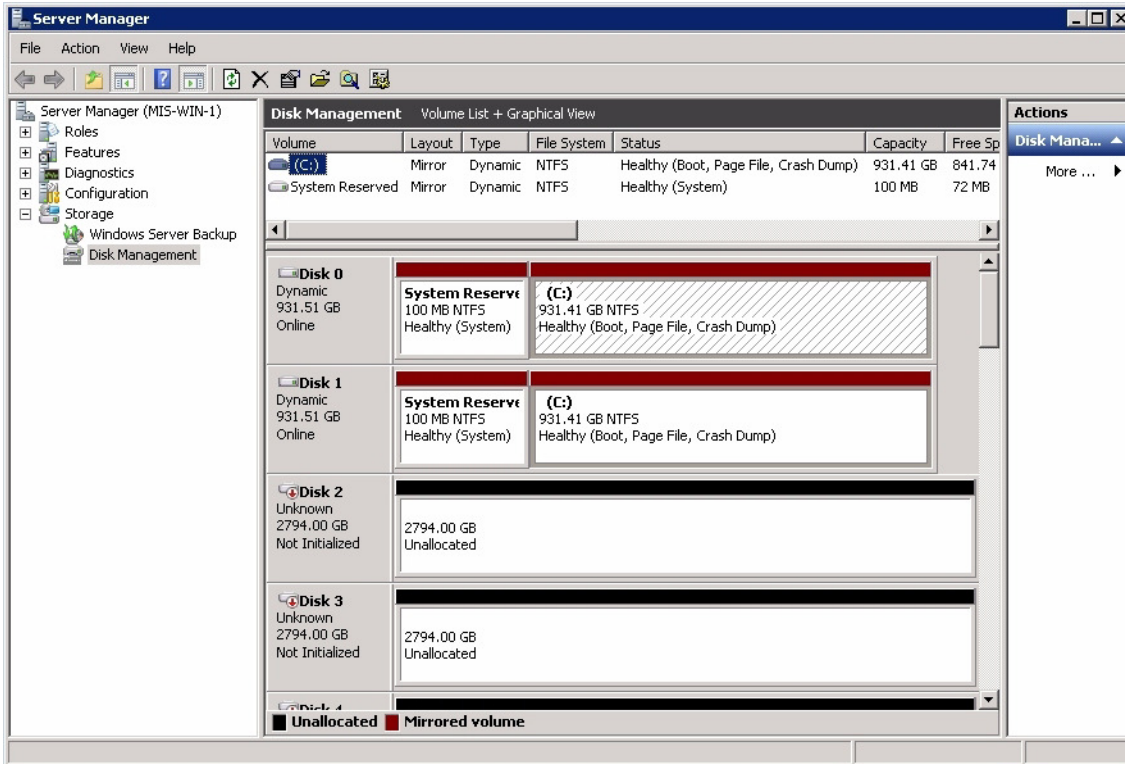


Figure 3-32 Server Manager – Disk Management

Right click in the grey area of the first non-system disk. In the menu that appears choose **Initialize Disk** (Figure 3-33).




Figure 3-33 Server Manager – Initialize Disks

A pop-up window will appear, showing all the uninitialized disks (Figure 3-34).

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



Figure 3-34 Server Manager – Select GPT (GUID Partition Table)

Click **OK**. All the disks should now show as **Online** (Figure 3-35).




Figure 3-35 Server Manager – Disks Initialized and Online

Right click the first non-system disk. Select **New Simple Volume** (Figure 3-36).




Figure 3-36 Server Manager – New Simple Volume

Click **Next** at the **New Simple Volume Wizard** welcome screen (Figure 3-37). Select the size of the volume in MB and click **Next** (Figure 3-38).



Figure 3-37 Server Manager – New Simple Volume Wizard

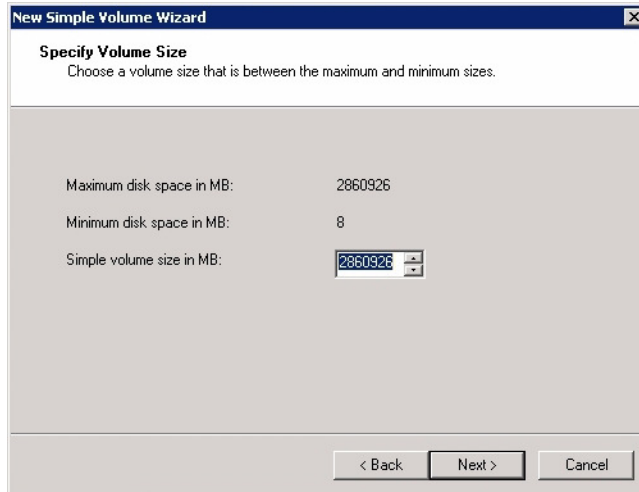


Figure 3-38 New Simple Volume Wizard – Volume Size

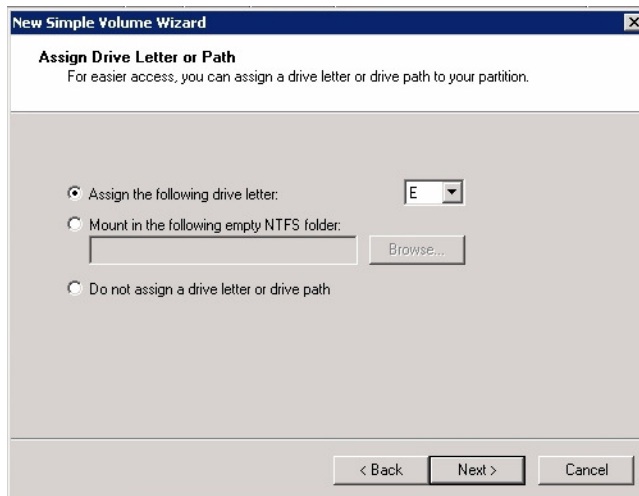


Figure 3-39 New Simple Volume Wizard – Assign Drive Letter or Path

Choose the drive letter to be assigned or click **Next** for the next drive letter available (Figure 3-39).

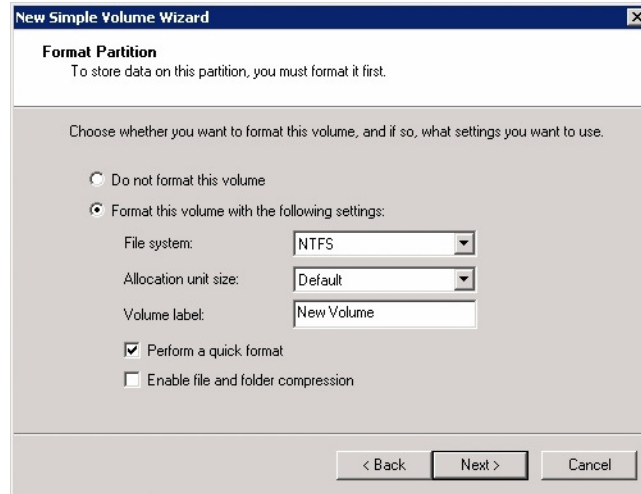


Figure 3-40 New Simple Volume Wizard – Format Partition

Select the format settings to be used, and click **Next** (Figure 3-40).

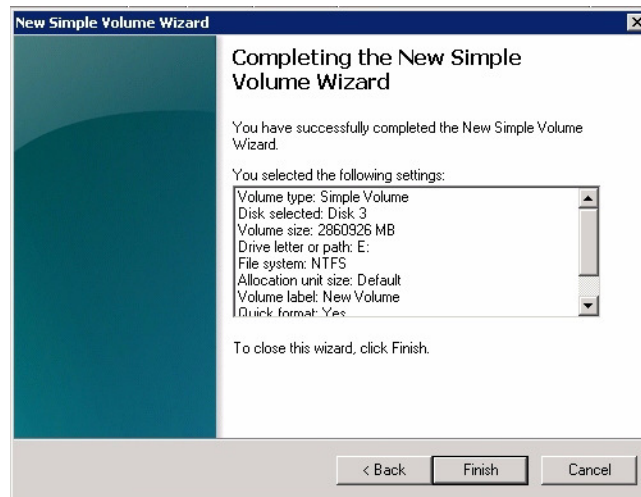


Figure 3-41 New Simple Volume Wizard – Settings Confirmation

Click **Finish** to format the disks (Figure 3-41). New volumes will show in the Disk Management window below the volumes (Figure 3-42).




Figure 3-42 New Simple Volume in Server Manager

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

Before Updating from CLI Zoning Tool 1.0

If you have the first version of the CLI Zoning tool installed on a host machine, remember to save any session files to a new location, and then simply install the new tool over the old one.

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>. 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, “Python for Windows,” prior to running the CLI Zoning Tool.

To install the CLI Zoning Tool software:

1. Go to <http://support.sgi.com>. Log in with your credentials.
2. Download the CLI Zoning Tool for Windows, version 1.4.

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. 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 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
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 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

Drive	Initiator 0	Initiator 1	Initiator 2	Initiator 3
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

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

Table 3-3 Zone Group Implementation (continued)

Zone Group	Description
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 116).

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 3-43). 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 154.) Ensure the MIS system is powered on.



Figure 3-43 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 infomation 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 114).

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 infomation 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; only PHY Based Zoning configurations are available at this time.

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 'pcsv' to modify the 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 'zcsv' to modify the 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.

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

Menu Option	Description
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 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.
6) Change Zoning type (Server <-> JBOD)	This menu option allows you to change the zoning type from SERVER/JBOD to JBOD/SERVER.
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 <code>csv</code> file to update Zoning.
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 3-44 on page 117 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

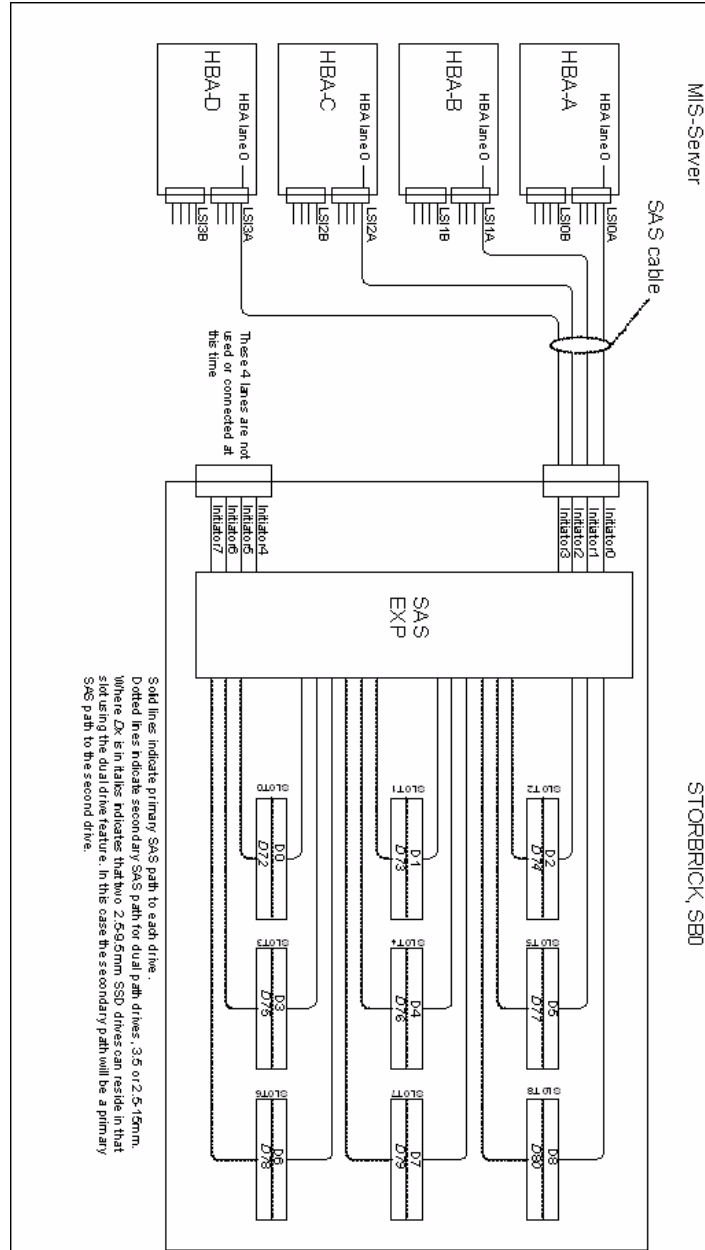


Figure 3-44 Block diagram of MIS Server StorBrick SB0

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	0x0000001e0	0x00000019a	0x00000001e	0x00000000
Phy 07	0x00000007a	0x000000065	0x00000001e	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 0x0000003bc
Phy 05 0x06b084c1 0xec636528 0x025a2c23 0x0000003ab
Phy 06 0x0053e9df 0x00000000 0x004bd5a9 0x000000356
Phy 07 0x016b6a91 0x3237dda3 0x0088112f 0x000000359
Phy 08 0x03cf2b5b 0x01d18ca5 0x76319c04 0x000000001
Phy 09 0x00000000 0x00000000 0x00000000 0x000000000
Phy 10 0x00000000 0x00000000 0x00000000 0x000000000
Phy 11 0x00e45221 0x00717733 0x1afe520c 0x000000001
Phy 12 0x00000000 0x00000000 0x00000000 0x000000000
Phy 13 0x00000000 0x00000000 0x00000000 0x000000000
Phy 14 0x00000000 0x00000000 0x00000000 0x000000000
```

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 CNG	ZONE GRP	ZONE CTRL	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	Time	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	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	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 Windows or Linux 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 3-45).




Figure 3-45 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 3-46).




Figure 3-46 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 RAID5, 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 3-47).




Figure 3-47 RAID 5 or 6 with One Drive per StorBrick

If a larger RAID group is desired, it would have 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 3-48).




Figure 3-48 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 3-49).




Figure 3-49 Three Drive Loss in RAID 6 Require StorBrick Replacement

Spare Drives

The LSI MegaRAID HBA can configure drives to be used as spares. There are general two ways of assigning spares, one 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 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 ServInfo (<http://servinfo.corp/>) or the Technical Publications Library (<http://docs.sgi.com>). Operating System software RAID support: 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
- External Air Ducts

All of the modules are hot-swappable when certain preconditions are met. That is, you can replace the hardware without powering down the platform. Whether or not a module is hot-swappable depends on the configuration of the MIS platform. For example, power supplies require the redundant power option be installed to allow for hot-swappable capability. Boot drives require mirroring enabled to be hot-swappable. Otherwise, all of the following service items are hot items (meaning they can be performed while the system is powered on and in use):

- “Detecting Component Failures” on page 138
- “Sliding the Chassis Forward/Backwards” on page 140
- “Removing the Front or Rear Chassis Cover” on page 140
- “Replacing a Power Supply” on page 141
- “Replacing a Fan Module” on page 143
- “Replacing a Boot Drive” on page 145
- “Checking the System Air Flow” on page 150
- “Installing the External Air Ducts” on page 151

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>). SGI Product Account and log-in credentials required.



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 69 for more information). Capacity drive failures are recognized and reported through MegaRAID Storage Manager (MSM). Boot drive failures are monitored and reported through `mdadm`. Fan and power supply failures are monitored through the fan baseboard and are reported by the BMC, with a few adjustments. See “Drive Monitoring” on page 45, “Power Supply Monitoring” on page 53, and “Fan Monitoring” on page 56 of Chapter 2, “System Interfaces” for more information on how to configure monitoring and alerting of component failures.



Figure 4-1 MIS-S9D Proprietary Network Interface

Selecting Which Side to Access

If you are servicing any of the **fans** or **StorBricks 0-4**, or using the **MIS-S9D Proprietary Network Interface**, access them through the **FRONT** of the rack, and extending the chassis out forward.

If you are servicing the **boot modules**, **StorBricks 5-8**, **server modules**, **JBOD I/O modules**, **internal** or **external cabling**, access them through the **REAR** of the rack, and extending the chassis out backwards.

Table 4-1 Access Points

Target	Access Instructions
Boot Module	REAR access, extend chassis BACKWARDS
Drives in StorBrick 0-4	Extend chassis FORWARD
Drives in StorBricks 5-8	REAR access, extend chassis BACKWARDS
Exterior Air Ducts	FRONT access, extend the chassis BACKWARDS
External Cabling	REAR access, no need to extend the chassis from the rack
Fans	Extend chassis FORWARD
Internal Cabling	REAR access, extend chassis BACKWARDS
JBOD I/O Modules	REAR access, no need to extend the chassis from the rack
MIS-S9D Proprietary Network Interface	FRONT panel, on top of the grille.
Power Supplies	REAR access, no need to extend the chassis from the rack
Server Modules	REAR access, no need to extend the chassis from the rack

Warning: Only extend ONE (1) MIS chassis out of a rack in either direction at any given time (e.g., one chassis extended backwards and one chassis extended forwards – OK; two or more chassis extended out forwards or backwards – BAD). Too much weight in either direction can potentially tip the rack over, causing injury or death.

Sliding the Chassis Forward/Backwards

The railmount of the MIS chassis allows it to be slid forward or backwards 20 inches. 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 railmounts, towards the center of the chassis.
2. Pull the chassis out using the handles. The chassis will latch at the 20" 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

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. This section describes the steps.

1. Extend the chassis from the rack until it latches at the 20" limit.
2. Remove the single security screw from the cover (Figure 4-3).




Figure 4-3 Top Cover Security Screw

3. Push the detent, and slide the cover out and up from the chassis.

Replacing a Power Supply




Figure 4-4 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 4-4).
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 replaced in a system that is already powered on, the power supply should respond to being plugged in and will power on automatically. If the power supply is being replaced in a system that is not powered on, it 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 4-5 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://<DHCP>/service.xml> (for JBODs) from the connected server/laptop. This will bring up the Fan Base Service Page (Figure 4-6).


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		D2	D5	D8	
Power Supply 2 Good						StorBrick 0 Good	D0	D3	D6		
Power Supply 1 o					FAN 0 Good <input type="button" value="On"/> <input type="button" value="Off"/> RPM : 10987 : 10763		D1	D4	D7		
Power Supply 0 o					D2		D5	D8			

Figure 4-6 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 140) and locate the fan module with the illuminated blue LED (Figure 4-5).
- 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 Boot Drive



Drives have two functions in an MIS Server: as capacity drives in the StorBricks, and in the boot drive module which houses the operating system (MIS JBODs are all capacity drives). The following instructions are for replacing a failed boot drive in an MIS Server.



Important: These instructions are for a system with an SGI factory-installed operating system, which features a mirrored boot drive. If the machine has a customer-installed OS, RSTe may not have been enabled in the BIOS, and therefore mirroring may not be enabled. If so, these instructions will not work, and a complete reinstall of the operating system will be necessary (as there is no mirrored data from which to rebuild a lost drive).

1. Turn on the chassis ID LED using the **BMC Web Console** to identify the chassis in the rack (“Virtual Front Panel” on page 73).
2. Once located, extend the chassis out the rear of the rack until it locks at the 20" mark.
3. Remove the chassis cover. (See “Removing the Front or Rear Chassis Cover” on page 140.)
4. Locate the faulted drive with the unilluminated status LED (i.e., the drive *without* a flashing green LED), and remove it from its boot drive bay. (See “Removing a Drive” on page 149.)
5. Replace the faulted drive. (See “Re-installing a Drive” on page 149.)
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. Turn off the chassis ID LED on the front control panel.
9. Check the drive assignment by issuing the `# lsscsi` command. In this example, the new drive appears as `sdd`. This drive assignment is used in the next step to add the drive.

```
# lsscsi
[0:0:11:0]   enclosu SGI  CORP  STORBRICK-7       1500  -
[0:0:15:0]   enclosu SGI  CORP  STORBRICK-0       1500  -
[0:0:19:0]   enclosu SGI  CORP  STORBRICK-2       1500  -
[0:0:23:0]   enclosu SGI  CORP  STORBRICK-3       1500  -
[0:0:27:0]   enclosu SGI  CORP  STORBRICK-6       1500  -
[0:0:31:0]   enclosu SGI  CORP  STORBRICK-5       1500  -
[0:0:35:0]   enclosu SGI  CORP  STORBRICK-4       1500  -
[0:0:39:0]   enclosu SGI  CORP  STORBRICK-1       1500  -
[7:0:0:0]    disk    ATA      INTEL SDSA2CW30  4PC1  /dev/sda
[7:0:1:0]    disk    ATA      TG32C10120GK3001  PRO6  /dev/sdd
```

10. Add the new drive to the container as a spare with the `# mdadm -a` command, and it will replace the missing drive:

```
# mdadm -a /dev/md/md127 /dev/sdd
mdadm: added /dev/sdd
```



Important: If you get the error, "Cannot add disks to a 'member' array, perform this operation on the parent container," make sure you are adding to the container, not the RAID1.

11. Issue the `# cat /proc/mdstat` to monitor the rebuild.

```
# cat /proc/mdstat
Personalities : [raid1] [raid0] [raid10] [raid6] [raid5] [raid4]
md125 : active raid1 sdc[2] sdb[1]
        111357952 blocks super external:/md127/0 [2/1] [U_]
        [>.....] recovery = 3.8% (4257024/111358084)
        finish=9.2min speed=193501K/sec

md126 : active raid1 sdc[2] sdb[1]
        5556224 blocks super external:/md127/1 [2/1] [U_]
        resync=DELAYED

md127 : inactive sdc[1](S) sdb[0](S)
        6306 blocks super external:imsm
```

Note: Finish time is only an approximation since the resync speed will vary according to other I/O demands.

Replacing a Capacity Drive




Figure 4-7 Capacity Drive and Carrier



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

To replace a failed capacity drive:

1. Using MegaRAID Storage Manager, verify the fault (failed unit).
2. Set the drive 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. Turn on the chassis ID LED in the Integrated BMC Web Console to identify the chassis in the rack (“Virtual Front Panel” on page 73).
4. If the drive is in StorBrick 0 through StorBrick 4, extend the chassis out the front of the rack until it locks at the 20" mark. If the drive is a boot drive or in StorBrick 5 through StorBrick 8, extend the chassis out the rear of the rack until it locks at the 20" mark.
5. Remove the chassis cover. (See “Removing the Front or Rear Chassis Cover” on page 140.)

6. Locate the faulted drive with the illuminated yellow LED, and remove it from its StorBrick (or boot drive bay). (See “Removing a Drive” on page 149.)
7. Replace the faulted drive. (See “Re-installing a Drive” on page 149.)
8. Once the drive is replaced, re-install the chassis cover and security screw.
9. Unlock the chassis from the extended position and push it back until it locks into the stowed position.
10. Using the MegaRAID Storage Manager, return the system to a normal state. The new drive will be powered on.
11. If there are no faults at this time, the rebuild or mirroring of the data to the new drive will begin.

Removing a Drive

As shown in Figure 4-7, 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 a 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.

Re-installing a 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.

Important: 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 4-7).
5. Replace the drive carrier into the chassis.
6. Push the drive carrier down to lock it place.
7. Replace the chassis cover and security screw.
8. Unlatch the chassis from the 20" mark and return the chassis back into place in the rack.

Checking the System Air Flow

To check the air flow for an MIS enclosure, perform the following steps:

1. Remove the chassis cover. (See “Removing the Front or Rear Chassis Cover” on page 140)
2. Remove the midspan: unscrew the Phillips screws from either end of the brace, and lift away the brace.
3. Make sure there are no objects wires or foreign objects obstruct air flow through the chassis. Pull all excess cabling out of the airflow path.




Figure 4-8 MIS Chassis Midspan Support Brace

Installing the External 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 4-9 Rack Air Duct Placement

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 154
- “Losing the System’s Setup Configuration” on page 154
- “Safe Power-Off” on page 154
- ““Missing” Firmware Files” on page 155
- “Cannot Receive Email Alerts Using BMC” on page 155
- “Intel BIOS POST error messages and handling description” on page 158
- “Beep Codes” on page 161
- “Not all beep codes signal an error” on page 163
- “Fans don’t spin when power button pressed” on page 163
- “Fans spin when power button is pressed, but no video display is seen” on page 163
- “BIOS or board logo appear, but OS load screen never appears” on page 164
- “OS load screen appears, but nothing further” on page 164
- “OS fully loads, but errors are seen” on page 164
- “Cannot boot after new software or drivers were installed” on page 164
- “Cannot boot after settings were changed in BIOS” on page 164
- “System Gets Stuck in a Reboot Cycle” on page 165
- “Shortening the Cable Management Arm” on page 165

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) Revision 2.3).

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 the speaker to determine if any beep codes exist. Refer to Appendix B, “BIOS Error Codes” for details.

Losing the System’s Setup Configuration

Make sure that you are using a high quality power supply. A poor quality power supply may cause the system to lose the CMOS setup information. If this does not fix the Setup Configuration problem, contact your vendor 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 through the CLI.

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. So 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).

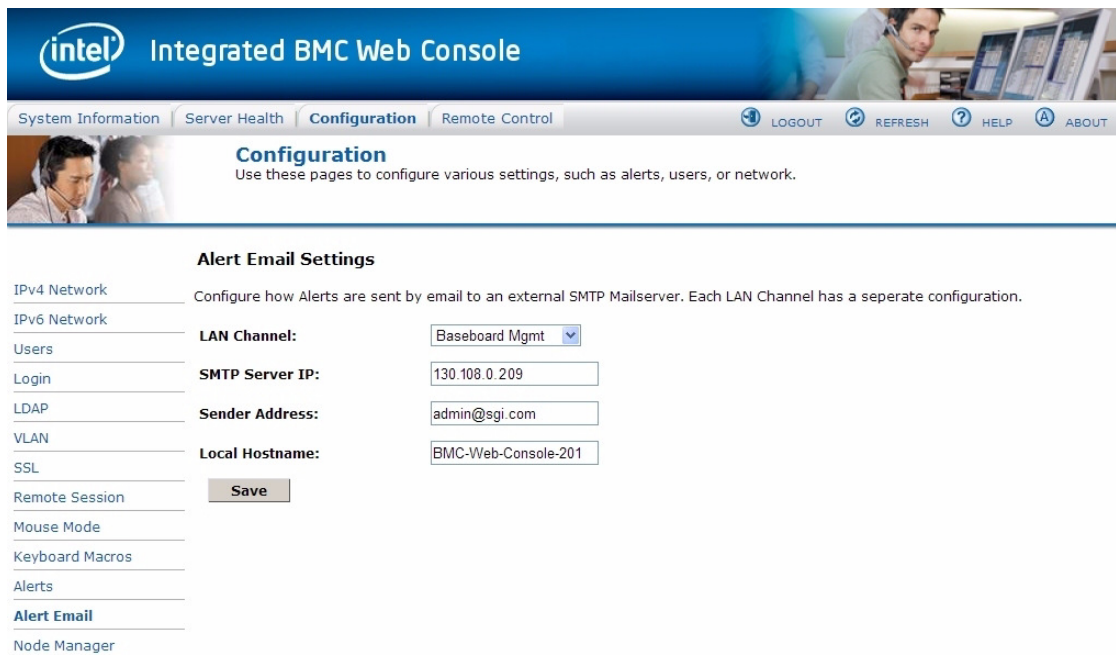


Figure 5-1 BMC Alerts Email Settings Page

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).

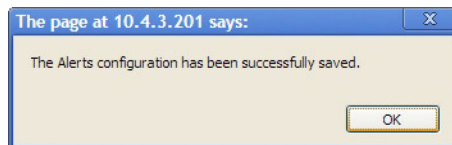


Figure 5-2 BMC Alert Configuration Success Pop-up

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

Alerts

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

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

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@sgl.com

Alert Destination #2:

SNMP Send SNMP Alerts to IP: 0.0.0.0

Email Send Email to:

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 (Figure 5-4).




Figure 5-4 BMC Alert Configuration Success Pop-up

Intel BIOS POST error messages and handling description

Refer to this table if you receive a code during the POST (Power-On Self Test) of booting by 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
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

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

Error Code	Error Message Response	Event Type
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
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

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

Error Code	Error Message Response	Event Type
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
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

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

Error Code	Error Message Response	Event Type
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	WarchDog 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
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 .

Table 5-2 Summary of BMC Beep Codes

Beep Code	Reason for Beep	Associated Sensor
1-5-2-1	No CPUs installed or first CPU socket is empty.	CPU Missing Sensor
1-5-2-4	MSID Mismatch.	MSID Mismatch Sensor
1-5-4-2	Power fault: DC power is unexpectedly lost (power good dropout).	Power unit – power unit failure offset.
1-5-4-4	Power control fault (power good assertion timeout).	Power unit – soft power control failure offset.
1-5-1-2	VR Watchdog Timer sensor assertion	VR Watchdog Timer
1-5-1-4	The system does not power on or unexpectedly powers off and a power supply unit (PSU) is present that is an incompatible model with one or more other PSUs in the system	PSU status

Another source of beep codes may be the LSI MegaRAID card. Table 5-3 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 5-3 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.

Table 5-3 Summary of LSI MegaRAID Card Beep Codes

Beep Code	LSI Firmware State	Cause (Depending on RAID Level)
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 it is yes, there is good power to the modules. Verify all required power cables are correctly plugged into the motherboard. Verify front panel cables are fully seated.

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 141).

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 5-3 on page 162. 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 5-5).




Figure 5-5 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 5-6).



Figure 5-6 Remove Segments from the End of an Empty Arm

Pry the joint apart on one side (Figure 5-7).



Figure 5-7 Pry the Joint Apart

Pull the joint open (Figure 5-7).



Figure 5-8 Pull the Segments from the Arm

The segments are now removed from the end of the arm (Figure 5-8), and the arm can now be attached to the chassis and rack (Figure 5-9).



Figure 5-9 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 5-10).



Figure 5-10 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 5-11).



Figure 5-11 Pry Open the Segment Joint

Do this to the other side, and then pull apart the arm (Figure 5-12).



Figure 5-12 Pull Apart the Cable Management Arm

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



Figure 5-13 Pry Apart the Joint at the Other End

Once complete, separate the segment from the cables by twisting the segment off (Figure 5-14).



Figure 5-14 Removing the Segment from the Cables

Bring the two spliced ends of the cable management arm together, and snap the joint together (Figure 5-15).



Figure 5-15 Join the Spliced Ends of the Cable Management Arm

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

Installing the Cable Management Arm

To install the cable management arm, perform the following steps.

1. First extend the chassis out the rear of the rack (“Sliding the Chassis Forward/Backward” on page 154).
2. Next, lie the cable management arm on a table with the open side up.
3. Lay the cables to be attached to the chassis on top, and starting at one end, press the cables into the arm until they snap inside. Continue to press the cables down the length of the arm until they are all inside.
4. Once complete, slide the arm down the cables until there is enough slack to allow the arm to be bolted into place on the Cable Arm Bracket (Figure 5-5 on page 166).
5. Using the attached bolts, fasten the Cable Management Arm Plate to the Cable Arm Bracket.
6. Then, wrap the arm around itself (as in Figure 5-5).
7. Next, slide the chassis out the front of the rack to be able to access the rack mount EIA side rails.
8. Attach the end arm plate to the EIA rail using the provided hardware.
9. Return the chassis to the stored position in the rack.
10. The cables can now be connected to their corresponding inputs (power, networking).
11. Once cabling is complete, use the provided velcro strips to wrap the cords and hold them tight (Figure 5-16).




Figure 5-16 Velcro Cable Ties

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 mount (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 (50-60KHZ) 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	<ul style="list-style-type: none"> –41° to 95° F (5° to 35° C) –processor cores automatically allowed to run faster than the base operating frequency, if the cores are operating below power, current, and temperature specification limits (< 35°C ambient)
Non-operating temperature range	–40° to 140° F (minus 40° to 60° C)
Operating humidity range	10% to 90% non-condensing
Non-operating Humidity	10% to 95% non-condensing
SGI MIS Server Specifications	
Servers/System	<ul style="list-style-type: none"> –One or two server modules per system –Single- or 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
Max cores	16 per server (32 per system)
Memory	<ul style="list-style-type: none"> –Up to 8 DDR3 DIMMs (4, 8, 16, or 32 GB) for a single-server motherboard configuration –Up to 16 DIMMs for a dual-server motherboard configuration. –Max 128GB per server
Boot drives	<ul style="list-style-type: none"> –Two per server, mirrored using RAID 1 –3.5" or 2.5" (15mm or 9.5mm) –SAS or SATA –Rotational or SSD –Up to 300 GB

Table A-1 Technical Specifications (continued)

Attribute	Specification
Supported operating systems	–RHEL 6.2 –SLES 11 SP1 –VMware ESX 5.0 –Windows 2008 R2 SP1 (WHQL certified, not supported through SGI)
Networking	Up to four user-specified PCIe HBAs, full-height (4.25") and half-depth (3.375"), externally or internally facing
Expansion slots	Single server: 6 x PCIe gen 2 Dual server: 4 x PCIe gen 2 per server (8 total)
RAID controllers	–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 via PCIe cards
Internal storage	–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	–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

