SGI® InfiniteStorage 16000 (IS16000) User's Guide

007-5725-001

Preface

What is in this guide

This user guide gives you step-by-step instructions on how to install and configure the SGI InfiniteStorage 16000 (IS16000) system.

Related Documentation

You should also see the following documents:

- *SGI InfiniteStorage 16000 Quick Start Guide* (007-5689-xxx)
- *SGI InfiniteStorage 16000 and InfiniteStorage 6120 CLUI Command Reference* (007-5726-xxx)

You can find SGI customer documents on the SGI Technical Publications Library (http://docs.sgi.com).

International Standards

The IS16000 complies with the requirements of the following agencies and standards:

- CE
- UL
- CUL
- C-Tick
- FCC

Potential for Radio Frequency Interference

USA Federal Communications Commission (FCC)

This equipment has been tested and found to comply with the limits for a class A digital device, pursuant to Part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference in which case the user will be required to correct the interference at his own expense.

Properly shielded and grounded cables and connectors must be used in order to meet FCC emission limits. The supplier is not responsible for any radio or television interference caused by using other than recommended cables and connectors or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the user’s authority to operate the equipment.

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

European Regulations

This equipment complies with European Regulations EN 55022 Class A: Limits and Methods of Measurement of Radio Disturbance Characteristics of Information Technology Equipments and EN50082-1: Generic Immunity.

Canadian Regulations

ICES-003 Class A Notice - Avis NMB-003, Classe A

This Class A digital apparatus complies with Canadian ICES-003.

Cet appareil numérique de la classe A est conforme à la norme NMB-003 du Canada.

Safe Handling

- Do not lift the IS16000 by the handles at the front and on the power supply modules on the back; they are not designed to support the weight of the enclosure.

Safety

NOTE : If this equipment is used in a manner not specified by the manufacturer, the protection provided by the equipment may be impaired.

Caution ! Safety goggles should be worn when maintaining the equipment.

Warning

The IS16000 *MUST* be grounded before applying power. Unplug the unit if you think that it has become damaged in any way and before you move it.

Caution ! To maintain proper airflow through the system, operate the system with the system top covers closed.

- Plug-in modules are part of the IS16000 and must only be removed when a replacement can be immediately installed. The system must not be run without all modules in place.
- In order to comply with applicable safety, emission, and thermal requirements, the top covers should remain closed while running.
- The IS16000 system must only be operated from a power supply input voltage range of 200 VAC to 240 VAC.
- A faulty power supply or fan module must be replaced with a fully operational module within 24 hours.

WARNING: To minimize the risk of electric shock, disconnect the power from the power supply, either by turning off the switch or by physically removing the power cable, prior to removing the PCM from the enclosure.

- Do not remove a faulty power supply or fan module unless you have a replacement module of the correct type ready for insertion.
- The power connection must always be disconnected prior to removal of the power supply from the IS16000.
- A safe electrical earth connection must be provided to the power cord.

- Provide a suitable power source with electrical overload protection to meet the requirements given in the technical specifications.

Do not remove covers from the power supply module. Danger of electric shock inside. Return the module to your supplier for repair.

Operation of the IS16000 with *ANY* modules missing will disrupt the airflow and the components will not receive sufficient cooling. It is *ESSENTIAL* that all apertures are filled before operating the unit.

Recycling of Waste Electrical and Electronic Equipment (WEEE)

At the end of the product's life, all scrap/ waste electrical and electronic equipment should be recycled in accordance with National regulations applicable to the handling of hazardous/ toxic electrical and electronic waste materials.

NOTE : Observe all applicable safety precautions, such as weight restrictions, handling batteries and lasers etc, detailed in the preceding paragraphs when dismantling and disposing of this equipment.

ESD Precautions

Caution ! It is recommended that you check and fit a suitable anti-static wrist or ankle strap and observe all conventional ESD precautions when handling the IS16000 plug-in modules and components. Avoid contact with backplane components and module connectors.

Data Security

- ***ALL*** the supplied plug-in modules and blanking plates must be in place for the air to flow correctly around the enclosure and also to complete the internal circuitry.
- If the IS16000 is used with modules or blanking plates missing for more than a few minutes, the system can overheat, causing power failure and data loss. Such use may also invalidate the warranty.
- Do not abandon your backup routines. No system is completely foolproof.

Product Support

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

- Refer to http://www.sgi.com/support/
- If you are in North America, contact the Technical Assistance Center at +1 800 800 4SGI or contact your authorized service provider.
- If you are outside North America, contact the SGI subsidiary or authorized distributor in your country.

Reader Comments

If you have comments about the technical accuracy, content, or organization of this document, 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.

Table of Contents

Appendices

CHAPTER 1

Introduction

1.1 Introduction

1.1.1 System Architecture

The IS16000 employs RAID, data integrity and data management software written from the ground up to take advantage of multi-core processors and modern bus architectures. Its highly threaded architecture allows performance to linearly scale with advances in underlying hardware. This same architecture allows the IS16000 to perform in the extreme range of both throughput and IOPS. Designed to house the most scalable unstructured file data, the system supports up to 600 disks of raw storage while enabling a combination of SAS, SATA or SSD disks.

1.1.2 Product Variations

Two couplet configurations (Figure 1) are available with the IS16000 to enable full system performance, capacity, and the highest levels of disk enclosure fault tolerance.

- 5 disk enclosures with up to 300 disk modules at 28U in one rack
- 10 disk enclosures with up to 600 disk modules at 48U in two racks

Figure 1. IS16000 Product Variations

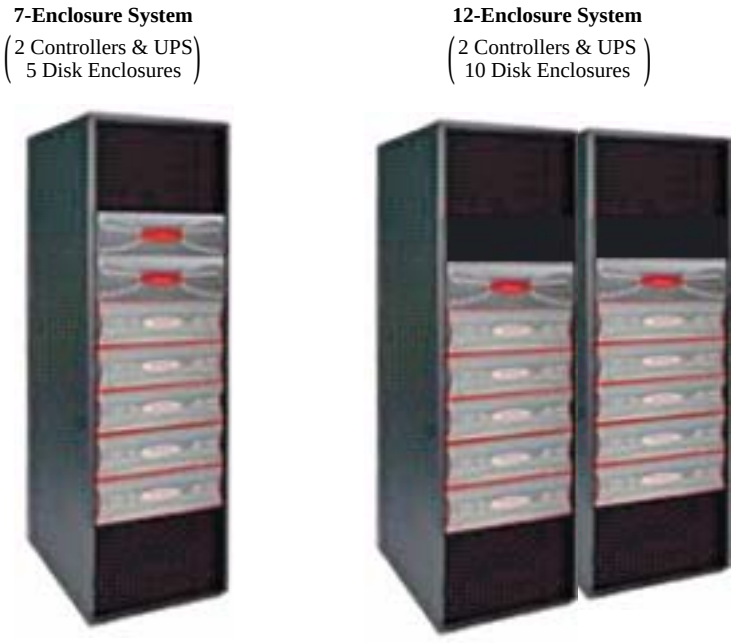

1.1.3 Features of the IS16000

The IS16000 incorporates the following features:

- 8GB/s full-duplex cache link
- Internal SAS Switching (480Gb/s Internal SAS Storage Network)
- Up to 600 SAS, SATA, or SSD disks with full redundant paths
- **InfiniBand™ (IB) or Fibre Channel (FC) Connectivity**
  For IB option, the two IS16000 Controllers provide up to eight (8) QDR InfiniBand host port connections.
  For FC option, the two IS16000 Controllers provide up to sixteen (16) individual 8Gb/s Fibre Channel host port connections, including simultaneous access to the same data through multiple ports. Each FC host port supports point-to-point and switched fabric operation.
- **Active/Active Operation with Cache Coherency and Failover**
  This architecture implements an active/active host presentation model with routing-based data access and full cache coherency.
- **Data Protection**
  The IS16000 RAID stack provides protection against single physical disk failures with RAID 1 or RAID 5 data protection as well as double physical disk failures through the use of high-speed RAID 6 protection.
- **Configurable RAID Group Sizes**
  (5 or 9 disks/RAID5 group, 6 or 10 disks/RAID6 group, 2 disks/RAID1 group)
  This feature allows you to configure the system with the desired RAID and redundancy levels based on the importance of your data. Each RAID group is configured independently and any valid combination for the number of disks in the array is supported.
- **SAS/SATA Storage Pool**
  This feature allows in-box storage pooling, where high performance SAS disks are used for primary data and high capacity SATA disks are used for secondary data or active archiving. Flexible configurations along with intermix of SATA and SAS disks within the same enclosure is supported. Although for performance and reliability reasons, there are best practice guidelines.
- **SATAssure™ Data Protection**
  SATAssure technology (Silent Data Corruption Detection and Avoidance) is designed to improve the reliability of enterprise SATA disks and make sure that data integrity is always mentioned for all I/O operations.
- **Battery Backed Write Back Cache**
  The IS16000 provides a write back cache feature that is used to improve I/O performance. Write back cache data that has not been written to disk is preserved by maintaining power to the cache memory in the event of an AC mains failure long enough to copy the contents of the cache to stable storage. In addition, the IS16000 is designed to tolerate a simultaneous AC mains failure and RAID software failure.
- **Mirrored Write Back Cache**
  Currently, the IS16000 provides the ability to mirror all write back cache data such that the failure of a single controller will not result in data loss. A storage administrator can optionally turn off write back cache mirroring for a RAID set (for higher performance); however data protection is reduced for logical units within that RAID set.
- **Mirrored Transaction Journal**
  RAID write holes are prevented by executing stripe updates as ACID (Atomicity, Consistency,

Isolation, Durability) transactions. If the transactions are interrupted by a power failure, they can be recovered from the transaction journal implemented within the write back cache when power is restored. This journal is mirrored so that when a simultaneous power failure and controller hardware failure occurs, the surviving controller can recover the transactions.

- **Metadata Mirrored n-Ways**
  The IS16000 stores a copy of storage system metadata on 18 physical disks to minimize the likelihood that its metadata is lost or corrupted.

- **Partial Rebuild**
  This feature reduces rebuild times by updating only the data that has changed while the disk was down. The IS16000 tracks the changes made to a RAID set when a member physical disk becomes unavailable. If that member becomes available again within a user-definable timeout, then only the stripes that were modified while the member was missing are rebuilt. This minimizes the mean-time-to-repair for the RAID set and thus limits any performance impact of a disk repair.

- **Hot-Swappable and Redundant Components**
  Adhering to enterprise RAS standard, almost all hardware components (such as physical disks, power supply modules, and fan module) are redundant and hot-swappable.

- **Hot Spares**
  The IS16000 provides pools of spare physical disks that can be automatically used to replace failed physical disks. By replacing a failed RAID set member automatically, the mean-time-to-repair for the RAID set is minimized resulting in improved data reliability.

- **Management Options via RS-232 and Ethernet (SSH)**
  A RS-232 port and Ethernet ports are included to provide local and remote management capabilities.

1.2 The IS16000 System Hardware

This section describes the hardware components of the IS16000.

1.2.1 Controller

The IS16000 Controller is a three-unit (3U), rack-mountable enclosure (Figure 2).

Figure 2. IS16000 Controller Chassis

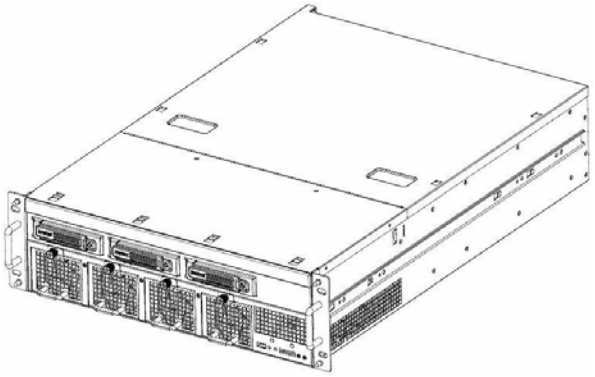

At the front, there are four fan modules, three internal disk modules, control buttons and status LED indicators (Figure 3).

Figure 3. IS16000 Controller Front View without Bezel

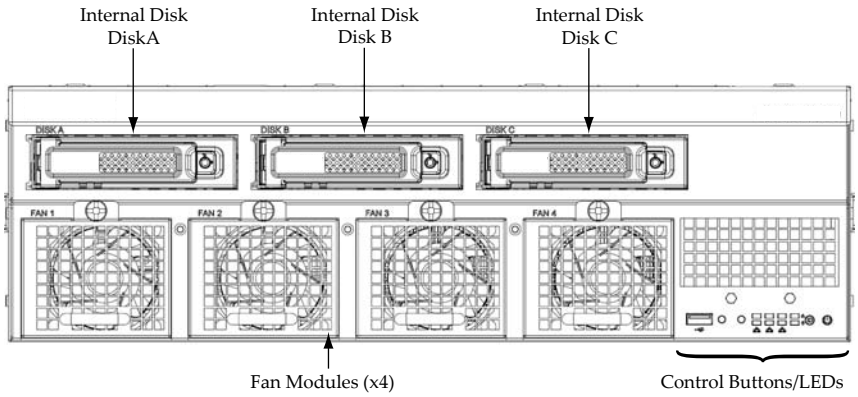

At the back, there are two power supply modules and various I/O connectors (Figure 4).

Figure 4. IS16000 Controller Rear View

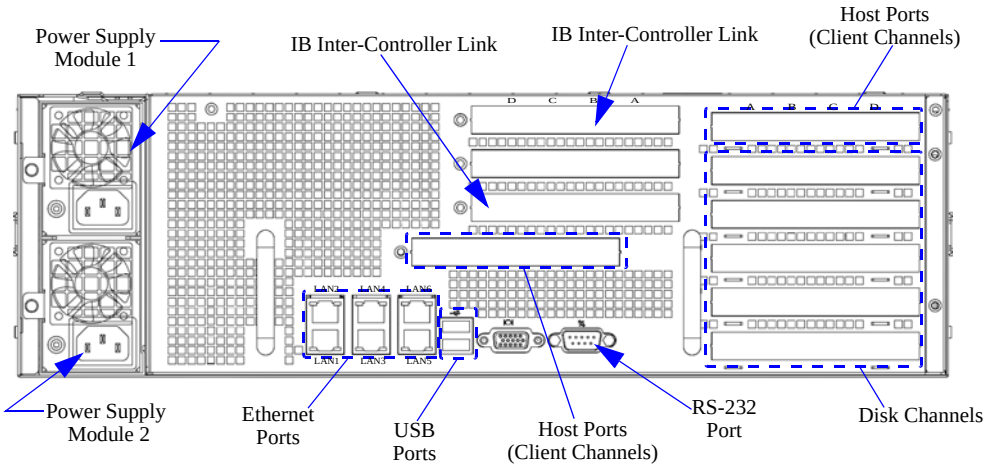

1.2.1.1 Status LED Indicators

Figure 5 below illustrates the positions of all the control buttons and LED indicators on the front panel. The LEDs are defined in Figure 6. The USB port can be used for downloading new firmware and BIOS to the Controller.

Figure 5. Controller Front Panel Control Buttons and Status LED Indicators

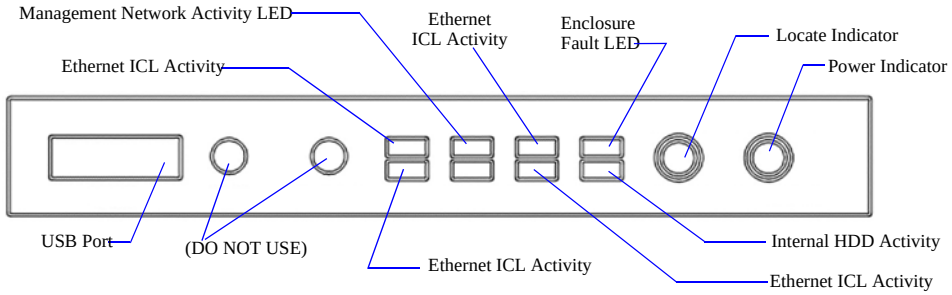

NOTE : Do **NOT** use the Power button under normal operation since doing so may cause data loss. Use the UPS power switch instead if you need to power down the Controller.

Figure 6. Controller Status LED Indicators

Description	Color	Status
Power	Green	ON - power is applied to Controller enclosure OFF - no power is applied to Controller enclosure
Locate	Blue	Flashing at a 2-second interval - receiving “Locate Enclosure” command OFF - NOT receiving “Locate Enclosure” command
Enclosure Fault	Amber	ON - a fan failure, power supply failure, or over temperature condition occurred. A service action is required. OFF - no detectable faults

Figure 6. Controller Status LED Indicators

Description	Color	Status
HDD	Green	Flashes - indicates HDD activity OFF - no HDD activity
Ethernet ICL	Green	Flashes - indicates the corresponding ICL link is up and both Controllers are running OFF - no ICL activity
Network Activity	Green	ON - good connection is established on the corresponding Ethernet port Flashes - indicates activity

1.2.1.2 Fan Module

Each Controller is equipped with four fan modules (Figure 7). These fan modules provide redundant cooling system for the unit. If one module fails, the other three will maintain sufficient cooling for the enclosure. The faulty module will still be providing proper air flow for the enclosure so do not remove it until a new module is available for replacement.

Figure 7. Controller Fan Module

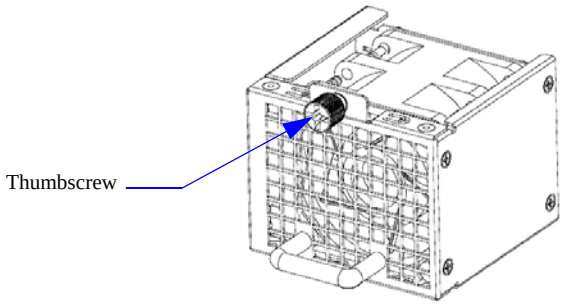

1.2.1.3 Power Supply Module

Each Controller includes two power supply modules (Figure 8). These modules are redundant and hot-swappable. If one module fails, the other module will maintain sufficient power to the enclosure. The faulty module will still be providing proper air flow for the enclosure so do not remove it until a new module is available for replacement.

The LED mounted on the module indicates the status of the power supply. It is green when the module is operating normally and turns off when a fault occurs.

Figure 8. Controller Power Supply Module

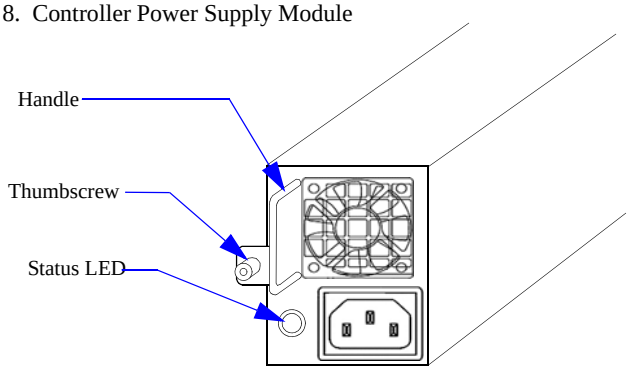

1.2.1.4 Internal Disk Modules

The IS16000 Controller includes three internal hard disks. Disk A and Disk B are mirrored system disks to run the Controller. The information on these 2 disks includes the Linux OS, the Controller software, the event log, and Controller-specific configuration (not storage configuration). Disk C is used for a large trace log[a].

Each disk module comprises a single low profile 1.0-inch high, 3.5-inch form factor hard disk mounted in a carrier. The module handle provides camming of the module into and out of disk bays and positive "spring loading" of the disk/baseplane connector (Figure 9). The handle is released by pressing the handle release button.

Figure 9. Controller Internal Disk Module

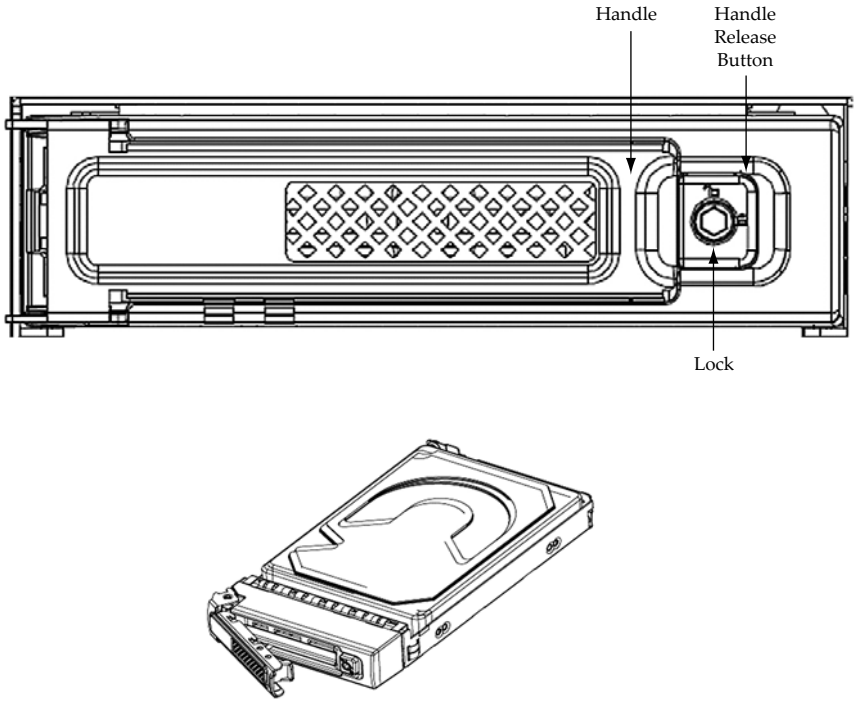

1.2.1.5 RAID Processor

The IS16000 has two parallel, multi-threaded RAID engines that work simultaneously in each Controller for a total of 4 RAID Processors across the redundant Controller pair.

The two Controllers are redundant and hot-swappable which provide the intelligence and active/active data protection features of the IS16000. If a Controller fails, the remaining Controller will assume its functionality and continue to provide data access, at a reduced performance level.

a. The mirrored system disk feature is not yet implemented in V1.3.0.4. Currently, Disk A is the system disk and Disks B and C are not used.

1.2.1.6 I/O Ports

Figure 10 below illustrates the I/O ports on the back of the Controller.

Figure 10. Controller I/O Ports

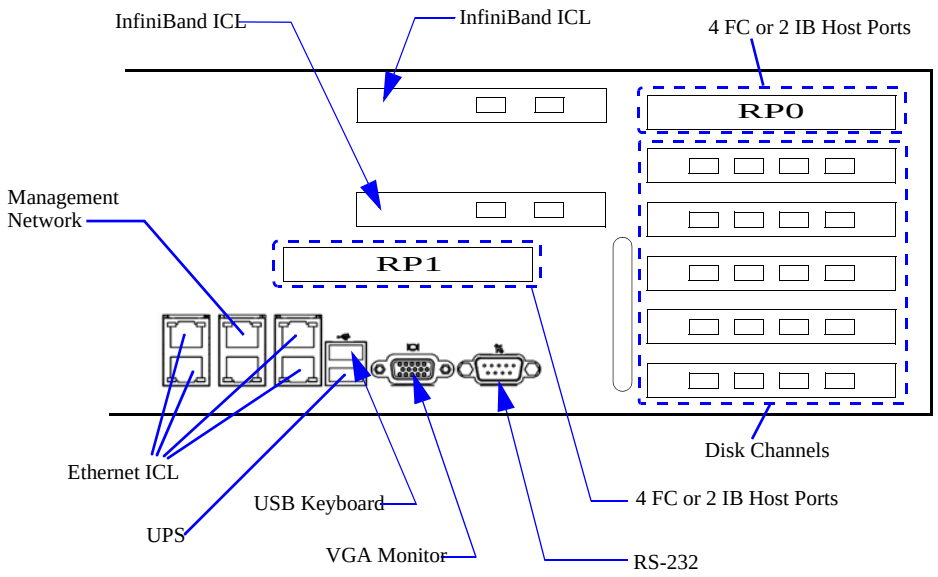

The InfiniBand ICL and Ethernet ICL ports provide connections for Mirrored Write Back Cache and control between the Controllers in a couplet configuration.

The Disk I/O Channels are used for disk enclosure connections.

The RS-232 connector provides local system monitoring and configuration capabilities. The VGA monitor and USB keyboard ports can be used as an alternative to the RS-232 console. However, the RS-232 console is recommended since its output can be logged and its connection can be longer.

The UPS port is used for UPS connection.

The Management Network Ethernet port provides remote monitoring and configuration capabilities.

The RP0 and RP1 host ports (client channels) provide 8 Fibre Channel or 4 InfiniBand host connections on each Controller. The LED color schemes are described in Figure 11 and Figure 12.

Figure 11. FC Host Port LED Color Scheme

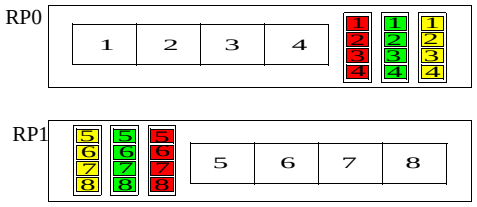

Yellow LED (8 Gbps)	Green LED (4 Gbps)	Red LED (2 Gbps)	Status
Off	Off	Off	Power off
On	On	On	Power on (before firmware init)
Flashing	Flashing	Flashing	Power on (after firmware init)
Yellow, Green, and Red LEDs flashing alternatively			Firmware error
Off	Off	On/Flashing	Online, 2 Gbps link / I/O activity
Off	On/Flashing	Off	Online, 4 Gbps link / I/O activity
On/Flashing	Off	Off	Online, 8 Gbps link / I/O activity
Flashing	Off	Flashing	Locate

Figure 12. IB Host Port LED Color Scheme

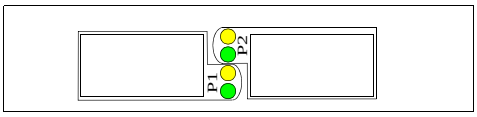

Description	Color	Status
Physical link	Green	ON - good physical link Flashing - indicates a problem with the link
Data activity	Yellow	ON - no data transfer Flashing - indicates data activity

1.2.2 Uninterruptible Power Supplies (UPS)

Each Controller within the IS16000 system is paired with a UPS to provide battery backup power when the AC mains fail. The UPS is a rack-mountable 1U unit.

Figure 13 below illustrates the positions of all the control buttons and LED indicators on the front panel. The DIP switches on the back panel are set at the factory and should only be changed by the SGI field engineer (Figure 14).

Figure 13. UPS Front View

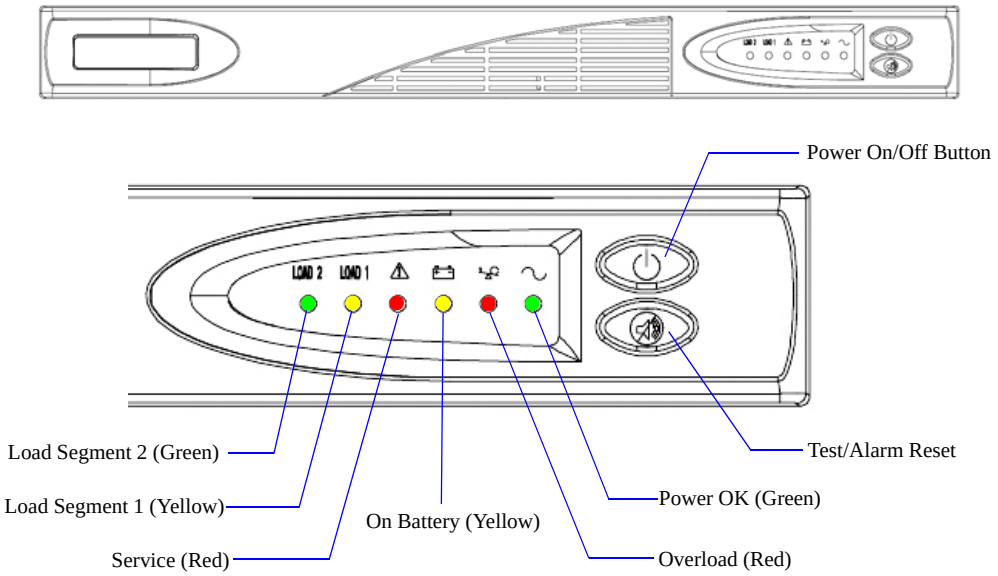

Figure 14. UPS Rear View

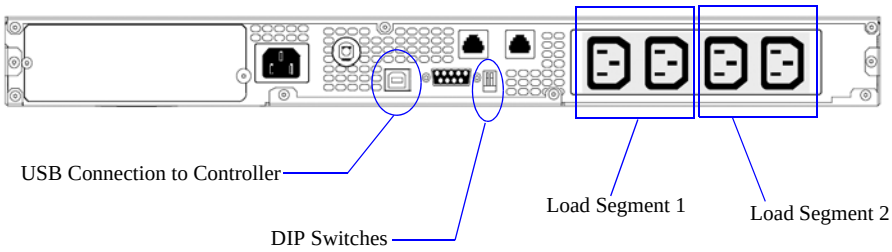

1.2.3 Disk Enclosures

The IS16000 disk enclosure is an ultra dense four-unit (4U) 60-disk enclosure based on 3.5 inch SAS or SATA Hard Disk Drives (HDD)(Figure 15).

The enclosure includes a set of plug-in modules and *(as supplied)* comprises:

- Enclosure Chassis with front panel LED status indicator
- Two (2) Power Cooling (PCM) plug-in modules
- Two Input/Output (I/O) modules
- Up to 60 top loadable hard disk modules in a 5×12 matrix
- Eight (8) SAS Drive Expander Modules (DEM)

Figure 15. Disk Enclosure—Front and Rear Views

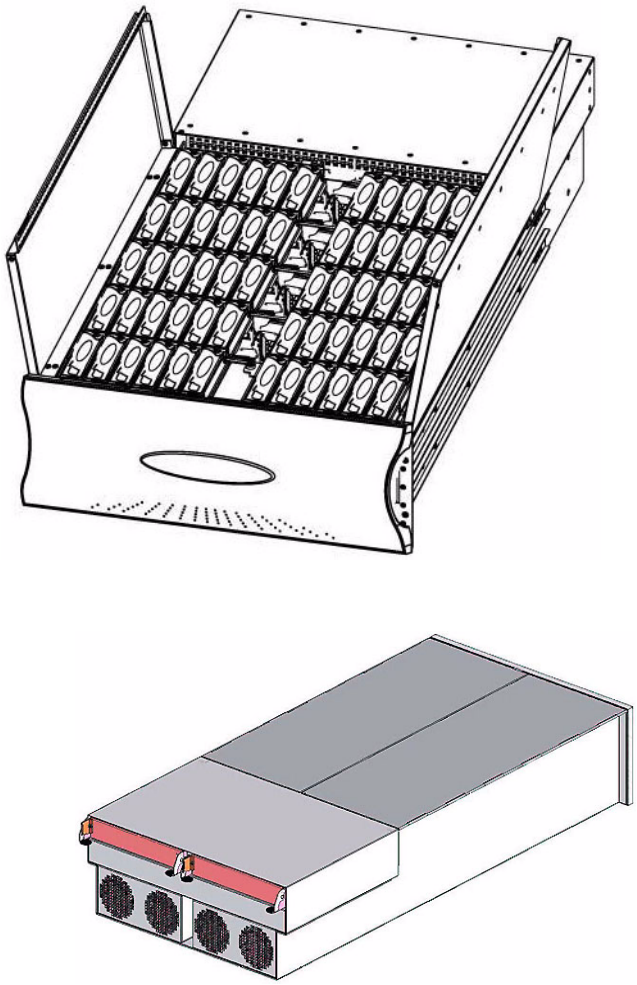

1.2.3.1 Disk Enclosure Chassis

The chassis assembly contains 60 disk bays at the front, each of which accommodates a plug-in disk module capable of holding a 3.5-inch SAS or SATA Hard Disk Drive (HDD) or Solid State Disk (SDD). The 60 disk bays are arranged in five rows of twelve disk modules (5×12) (Figure 16). At the rear, the chassis assembly contains two (2) power cooling modules and two (2) I/O modules (Figure 17).

The chassis is fitted with 19-inch rack mounting features which enables it to be installed into 19-inch wide racks and uses four (4) EIA units of rack space. A mid-plane separates the front and back of the chassis and provides the interconnect system between the power cooling modules, I/O modules, and the baseboard.

Figure 16. Disk Enclosure—Top View

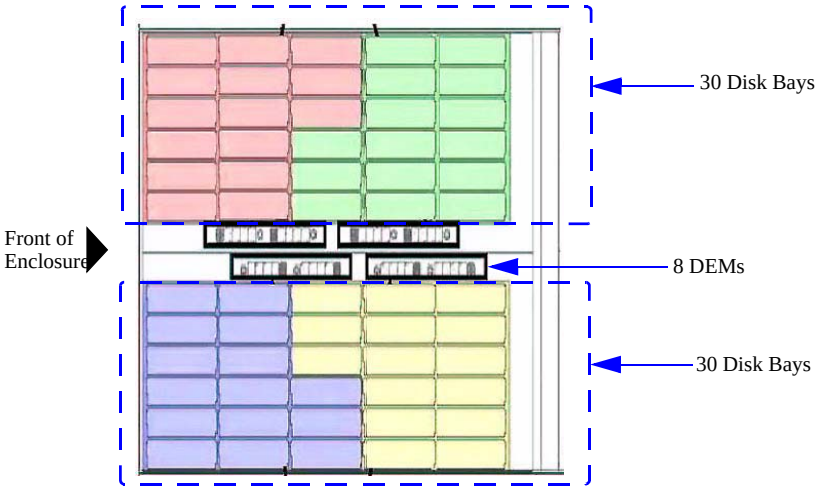

Figure 17. Disk Enclosure—Rear View

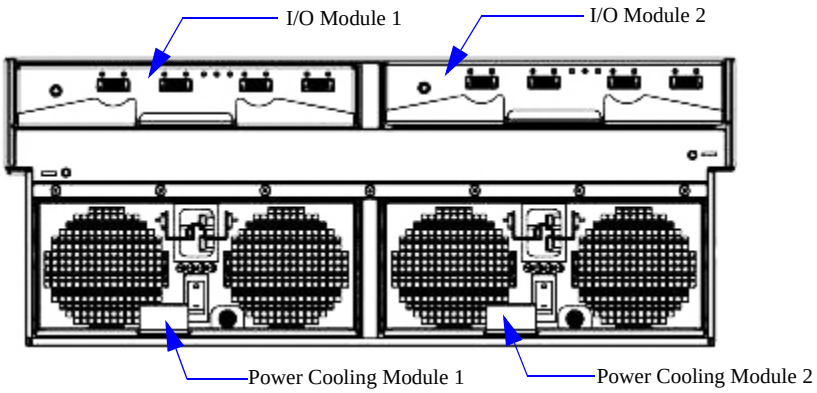

1.2.3.2 Front Panel Indicators

The front panel indicators show the aggregated status of all the modules (Figure 18). The LEDs are defined in Figure 19. The Disk Activity LEDs indicate disk presence and flash during data I/O.

Figure 18. Disk Enclosure—Front Panel LED Indicators

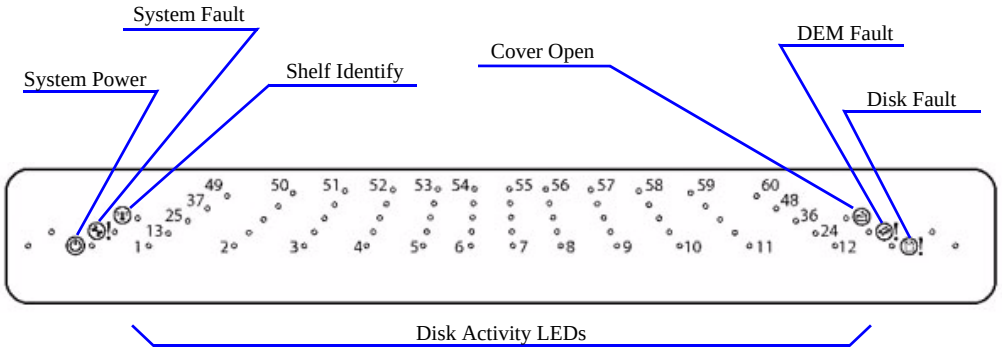

Figure 19. Front Panel LED Description

LED	Definition	Color	Status
	System Power	Green	ON - DC power is present OFF - DC power is not present LED does NOT flash under normal operating conditions
	System Fault	Amber	ON - one or more components within enclosure have failed. A service action is required. Exact failed component has its own amber fault LED lit. OFF - no detectable faults
	Shelf Identify	Blue	Flashing - receiving “Locate Enclosure” command OFF - NOT receiving “Locate Enclosure” command
	Cover Open	Amber	OFF - both cover pieces securely closed and latched in place ON - either of the cover pieces is NOT securely closed and latched in place
	DEM Fault	Amber	OFF - all DEMs operating correctly ON - at least one DEM has failed; service action required
	Disk Fault	Amber	ON - one or more HDDs are failed; SES must determine exact HDD OFF - no detectable disk faults
Individually numbered	Disk Activity	Green	ON - SAS HDD is present Flashing - indicates HDD activity OFF - no HDD activity

1.2.3.3 Disk Modules

The disk module comprises a hard disk mounted in a carrier (Figure 20). Each disk bay can house a single low profile 1.0-inch high, 3.5-inch form factor disk drive in its carrier. A fully loaded enclosure contains 60 disk modules.

The module handle provides the following functions:

- Camming of the module into and out of disk bays
- Positive “spring loading” of the disk/baseplane connector
- Disk Status LED incorporated in handle assembly

NOTE : The disk enclosure design allows for disk bays to be left empty without the need for fitting dummy disk modules.

The disk enclosure supports a SATA interposer card which allows simultaneous access to the SATA HDD from both I/O modules installed in the enclosure. ***Note that a SATA interposer card is required to run SATA HDDs***.

Figure 20. Disk Module

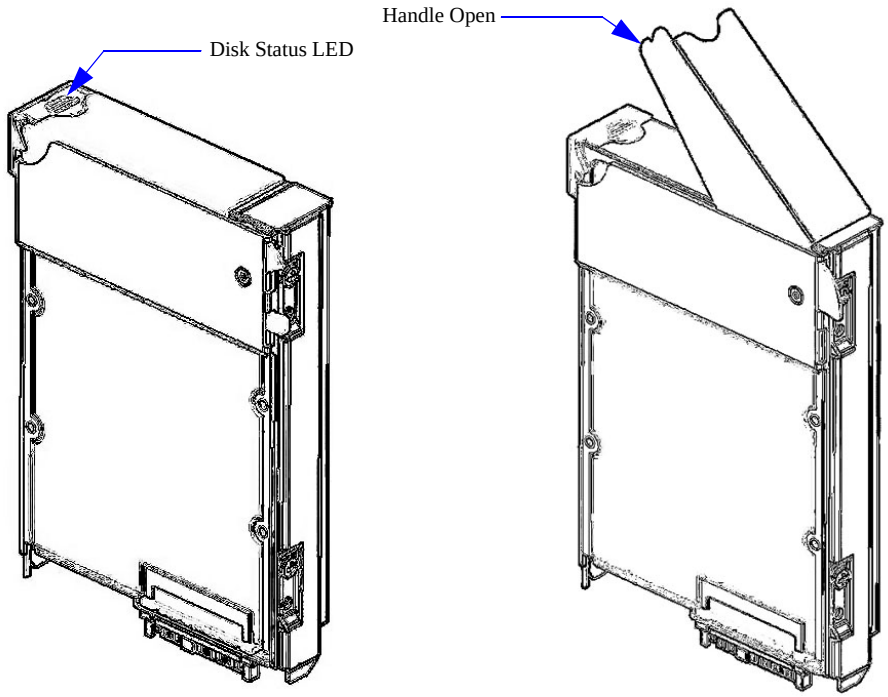

1.2.3.4 Power Cooling Module (PCM)

Two (2) power cooling modules are supplied with the disk enclosure, installed in the rear of the chassis (Figure 21).

Figure 21. Disk Enclosure Power Cooling Module

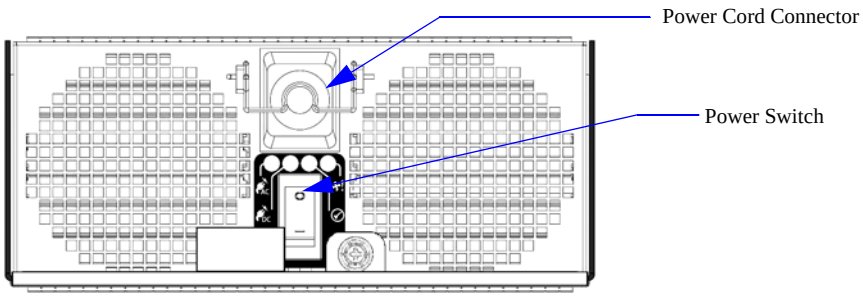

The disk enclosure must always be operated with two PCM installed. Module replacement should only take a few minutes to perform but must be completed within 5 minutes from removal of the failed module. Four (4) LEDs mounted on the PCM indicate the status of the module and the fans (Figure 22). Figure 23 provides a description of the color and status of the LEDs.

Figure 22. Power Cooling Module LEDs

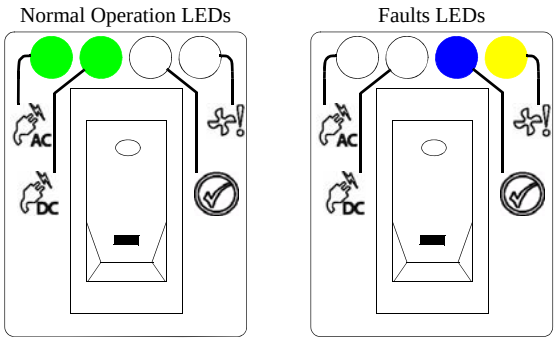

Figure 23. disk enclosure PCM LED Description

LED	Description	Color	Status
	PCM AC	Green	ON - AC input to PCM within tolerances OFF - PCM failed
	PCM DC	Green	ON - DC output of PCM within tolerances OFF - PCM failed
	PCM Fault	Amber	ON - PCM fault detected OFF - no detected PCM faults
	PCM ID	Blue	Flashing - this module is receiving “Locate Power Supply” or “Locate Fan” command OFF - NOT receiving “Locate Power Supply” or “Locate Fan” command

1.2.3.5 I/O Module

The disk enclosure includes two plug-in I/O modules (Figure 24). Processors housed on the I/O modules provide enclosure management and interface to devices on the backplane, PCM, and front display panel in order to monitor internal functions.

Figure 24. A Disk Enclosure I/O Module

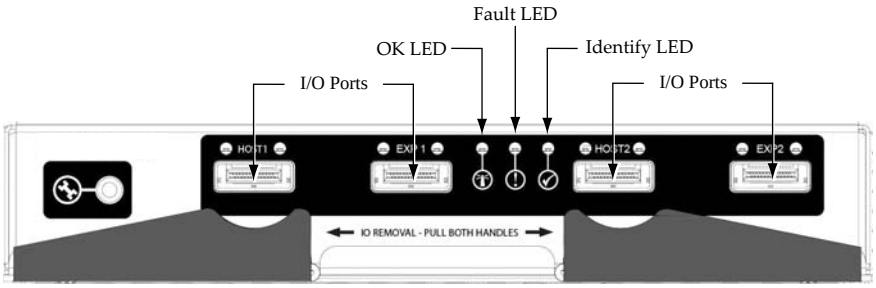

The four I/O ports provide connections to the IS16000 Controllers. The various LED indicators are listed in Figure 25.

Figure 25. I/O Module LED Description

Description	Color	Status
I/O port SAS Link	Green	ON - valid SAS link established OFF - no valid link
I/O port SAS Link Fault	Amber	ON - fault is detected on port OFF - no fault detected
OK	Green	ON - I/O module is functioning properly OFF - internal fault detected
Fault	Amber	ON - a fault is detected on I/O module OFF - no fault detected
Identify	Blue	Flashing - this module is receiving “Locate Expander” command OFF - NOT receiving “Locate Expander” command

1.2.3.6 Internal Indicators

For some components, the failure LED is internal to the enclosure and visible only when the cover is open. The various internal indicators are listed in Figure 26.

Figure 26. Internal LED Description

Description	Location	Color	Status
DEM DC	DEM internal to enclosure	Green	ON - 1.2VDC regulator circuit correctly functioning OFF - faulty 1.2VDC regulator circuit
DEM ID	DEM internal to enclosure	Blue	Flashing - receiving “Locate Expander” command OFF - NOT receiving “Locate Expander” command
HDD ID	HDD internal to enclosure	Blue	Flashing - this slot is receiving “Locate Slot” or “Locate Physical_Disk” command ON - Controller has determined that the HDD in this slot has failed or requires attention OFF - NOT receiving “Locate Slot” or “Locate Physical_Disk” command and HDD is functioning correctly

1.2.3.7 Drive Expander Module (DEM) Card

The disk enclosure contains eight (8) top-loadable, hot-swappable DEMs (Figure 27). The DEMs provide the SAS connectivity between the I/O modules and the HDDs located with the enclosure.

The DEMs are organized into four module pairs, with each module pair controlling a bank of 15 disks. If a DEM card fails, its sister card will provide continuous operation of the unit until a replacement is inserted bringing the unit back to full redundancy.

Figure 27. DEM Card Locations (Viewing from Top)

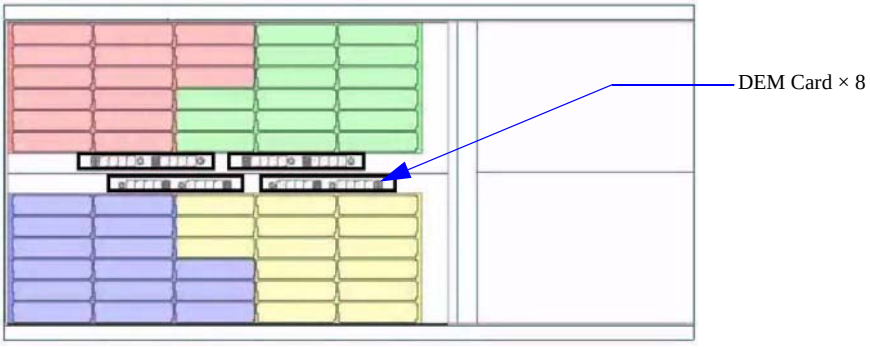

Each DEM (Figure 28) connects to a single port of the HDD based on the location within the enclosure within which it is installed. Each card provides power control signals to each disk slot. It controls HDD identify/service LEDs and monitors the status from partner DEM. Each pair of DEMs supports 15 disk modules.

Figure 28. DEM Card

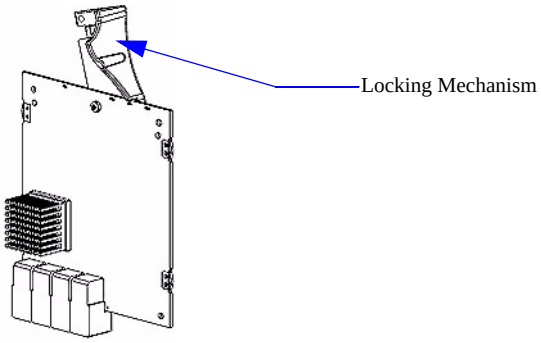

CHAPTER 2

Installation

2.1 Installation Overview

Here is an overview of all the steps needed to be taken to complete an installation and configuration of an IS16000 system.

1. Site preparation.
2. Unpack the IS16000 system.
3. Install the disk modules into the disk enclosures.
4. Verify zoning configuration on disk enclosures.
5. Connect the cables.
6. Power up the system.
7. Validate the hardware.
8. Configure the network interface settings via the RS-232 console.
9. Configure the IS16000 storage (create Storage Pools, Virtual Disks, Spare Pools, and presentations).

2.2 Site Preparation

This section provides information on how to prepare a site prior to installing a IS16000 system. The site preparation steps include:

- Delivery route verification
- Rack location, air flow, and access
- Floor loading
- Cooling supply planning
- AC power supply planning and verification

2.2.1 Delivery Route Verification

Each IS16000 system consists of one to two 19" racks. Each of these racks is shipped in a large crate on a pallet that weighs between 1000 lbs (454.55 kg) and 1500 lbs (681.82 kg) and measures 53" × 40" × 100.5" (1346 mm × 1016 mm × 2553 mm) (Figure 29).

Figure 29. Crate Containing One IS16000 Rack

Each crate includes a ramp that can be used to remove the rack from the crate. A pallet jack must be used to move the crate from the delivery vehicle to a flat location that is at least 5' × 12' (1.52 m × 3.66 m) where the ramp can be installed and the rack can be rolled out of the crate. Due to the weight of the crate, it is essential that either the delivery vehicle is equipped with a lift gate or the destination has a loading dock with a lift.

Once a rack has been removed from the crate, it can be rolled to its destination. Racks are either 42U or 45U. The 42U racks are 81.75" (2077 mm) tall and the 45U racks are 87" (2210 mm) tall. Both type of racks are 28" (712 mm) wide and 45.75" (1162 mm) deep including the front and rear doors.

Note that the physical disk modules are shipped separately from the disk enclosures. The disk modules are packaged in separate boxes that are strapped to pallets. These pallets can be moved to the installation site using a pallet jack.

To Do List:

- Plan the unpacking location where the delivery vehicle will deliver the crate.
- Verify that the route to the installation site does not involve any steps.
- Verify that any elevator that is part of the route has sufficient door height and load capacity.
- Verify that the size of all the doorways along the delivery route is big enough for the pallet jack with the crate to pass through.
- Verify that there is enough space at the unpacking location to set up the ramps and remove the racks from the crates.

2.2.2 Rack Location, Air Flow, and Access

The racks that comprise a single IS16000 storage subsystem must be installed adjacent to one another so that the disk enclosure cables can be properly routed and connected. The airflow through the racks is from front to back. The front of one rack should not be near the back of another rack.

> **NOTE :**Racks with more than 5 disk enclosures have an anti-tip plate installed at the floor level and it is important that it does not cover a cold air grill.

Access to the equipment in the racks is from their fronts and the backs. Front access is for physical disk, disk enclosure, controller chassis, and UPS maintenance. Rear access is for power supply and I/O module maintenance plus cabling.

The recommended service clearances are 48" (1220mm) from the front of a rack and 30" (762mm) between the back of an enclosure and a wall or 30" (762mm) between the back of a IS16000 system and the back of another enclosure that is exhausting hot air (Figure 30). No service clearance is required on the side of the rack.

Figure 30. IS16000 Service Area and Dimensions

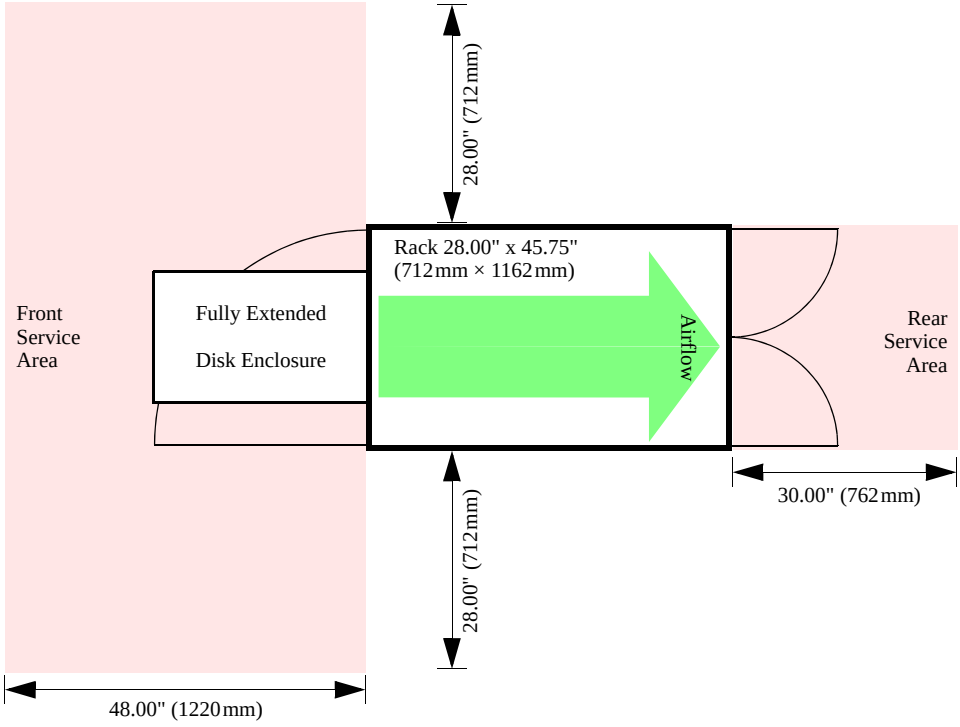

To Do List:

- Plan the location of the racks so that they have sufficient access and proper airflow. A scale drawing is recommended.

The best way to accomplish this is to create rows of racks that are side-by-side where each rack draws cool air from a 60" (1524mm) walking isle and blows hot air into a 48" (1220mm) maintenance isle (Figure 31).

Figure 31. Rows of Racks with Walking Isles and Maintenance Isles

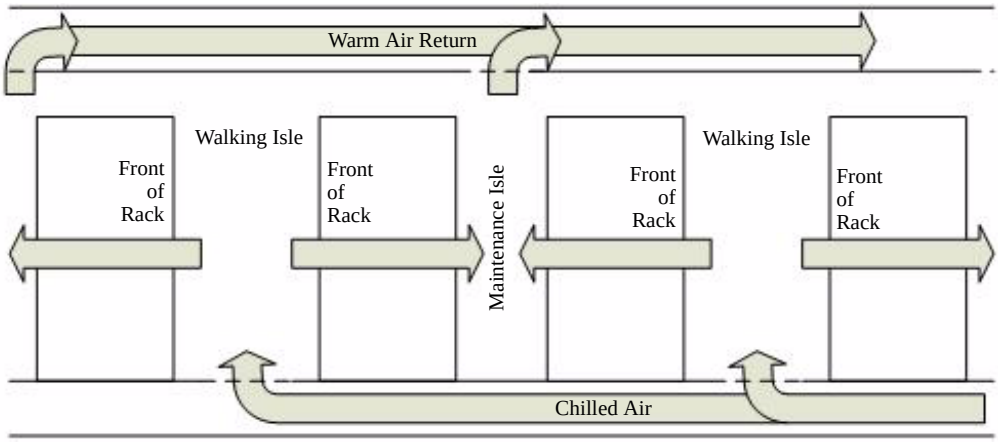

2.2.3 Floor Loading

The IS16000 racks are heavy and it is important to verify that the weight does not exceed the floor specifications. SGI will provide an estimate of the weight of each rack on request.

A full 45U IS16000 rack with ten disk enclosures can weigh up to 2570 lbs (1166 kg). This weight is relatively evenly distributed across the four castors and/or leveling feet, so each will be subject to a load of 645 lbs. While the load from such a rack is 289 lb/ft^2 (2570 lb / 8.9 ft^2) underneath the rack, the average load is less given the space for isles. With the minimum size isles (48" in the front and 30" in the back), the load is 156 lb/ft^2 (289 lb/ft^2 × 45.75" / {45.75" + 48"/2 + 30"/2}).

The more common 42U IS16000 racks for 5 and 10 disk enclosure subsystems weigh approximately 1600 lbs and so the point load is only 400 lbs and the floor loading is 180 lb/ft^2 (1600 lb / 8.9 ft^2). Accounting for the minimum isle space the load is 98 lb/ft^2.

To Do List:

- Verify with the building structural engineer that the floor structure is sufficient to hold the weight of the IS16000 racks configured as planned in Section 2.2.2.

2.2.4 Cooling Supply Planning

Each IS16000 45U rack can draw up to 14.5 KW and generate 50 KBTU/hour of heat.

Each IS16000 42U racks can draw over 8 KW and 25 KBTU/hour of heat.

Sufficient air conditioning must be provided to cool this heat load to a nominal room temperature of 25°C. SGI will provide an estimate of power and heat for each rack on request.

To Do List:

- Verify that the rack locations will have sufficient cooling.

2.2.5 AC Power Supply Planning and Verification

Each IS16000 rack has four to eight power cords that are designed to connect to 208 VAC or 230 VAC. SGI will provide a count of the power cords for each rack on request.

Each AC power cable has a NEMA L6-30P connector (Figure 32). Before the IS16000 racks can be installed, sufficient L6-30R receptacles for these AC power cables must be installed by an electrician. Each L6-30R receptacle should be connected to a 30 Amp circuit breaker.

Figure 32. L6-30R Wiring Diagram

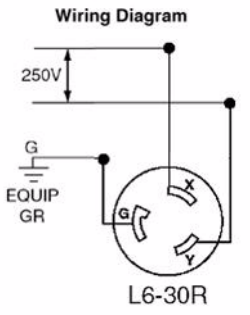

These power cords can either be connected to two separate AC power sources (for example, utility power and UPS power) or to a single AC power source. That is, half the power cords provide sufficient power to operate the system, however, all power cords must be plugged to take advantage of redundant DC power supplies.

Each IS16000 Controller has a UPS that is used to protect its cached data in the event of a power failure. These must be configured with firmware and switch settings depending on the AC voltage and how the site power is grounded. This configuration is best done at the SGI factory and so it is helpful to provide these measurements to SGI as part of the order.

To Do List:

- Each UPS unit has two DIP switches on the back. During installation, verify that the switches are correctly set as described in Section 2.6.7.

2.3 Unpacking the IS16000

The IS16000 system components are already installed in the rack(s).

Warning

Wear an ESD wrist strap or otherwise ground yourself when handling IS16000 modules and components. Electrostatic discharge can damage the circuit boards.

Before you unpack your IS16000, inspect the shipping container(s) for damage. If you detect damage, report it to your carrier immediately. Retain all boxes and packing materials in case you need to store or ship the system in the future.

While removing the components from their boxes/containers, inspect the IS16000 chassis and all components for signs of damage. If you detect any problems, contact SGI Technical Support immediately.

2.3.1 Packing List

The IS16000 ships with the following:

- Installed in the rack(s):
  - Two (2) Controllers
  - Two (2) UPS units
  - Four (4) 1-meter Ethernet ICL cables
  - Two (2) 1-meter InfiniBand ICL cables
  - Two (2) USB cables
  - Twenty (20) SAS cables
  - Five (5) or ten (10) disk enclosures
- Disk modules
- Two (2) serial cables
- Two (2) Ethernet cables

2.4 Installing the Disk Modules

The disk modules are shipped separately from the disk enclosures. To create a more balanced configuration, evenly distribute the disk modules among the disk enclosures. If a mixture of disk technologies, such as SAS and SATA, will be populated into the enclosures, it is best to populate the SAS disks into the front slots and install the SATA disks in the rear slots.

Warning **Wear an ESD wrist strap or otherwise ground yourself when handling the disk modules and components. Electrostatic discharge can damage the circuit boards.**

Follow these steps to install a disk module:

1. Slide the enclosure out from the rack by squeezing the tabs on both rack slides. Keep pulling until the enclosure locks and you hear a clicking sound.

2. Disengage both enclosure cover latches (Figure 33) and open the covers.

Figure 33. Disk Enclosure Cover Latch

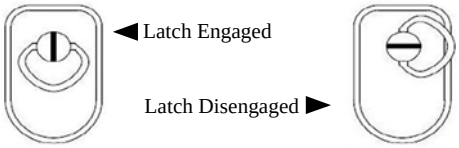

3. On the disk module, slide the latch backward to release the handle (Figure 34).

Figure 34. Release Disk Module Handle

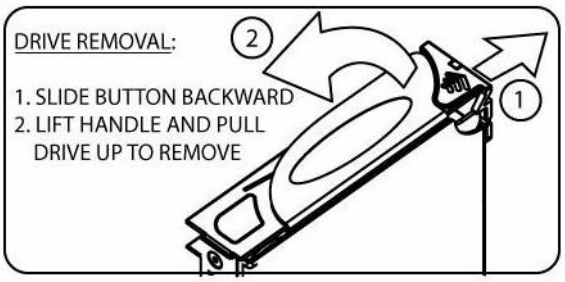

4. Insert the module into a disk bay. Cam the disk module home. The camming foot on the base of the module will engage into the slot in the enclosure.

5. When the module is fully inserted, close the handle. You should hear a click as the latch engages and holds the handle closed.

6. After you have installed all the disk modules in this enclosure, close the enclosure covers and engage both cover latches.

7. Squeeze the tabs on both rack slides and push the enclosure back into the rack.

2.5 Verifying Zoning Configuration on Disk Enclosures

The disk enclosures must be zoned correctly. The zoning configuration is determined by the piano switches located behind the front bezel of the enclosure. Follow these steps to verify the zoning configuration on each disk enclosure:

1. Slide the enclosure out from the rack by squeezing the tabs on both rack slides. Keep pulling until the enclosure locks and you hear a clicking sound.

2. Remove the front bezel by removing the four screws (Figure 35).

Figure 35. Removing Disk Enclosure Front Bezel

3. Once the front bezel is removed, it is possible to see the piano switches through the grey bezel (Figure 36). If the switches are hidden by a Mylar switch cover, remove the cover with a small pair of tweezers. If the switches are still not visible, remove the grey bezel. It is attached to the chassis with 32 small screws (18 on the front and 14 on the bottom).

Figure 36. Zoning Configuration Piano Switches Behind Bezel

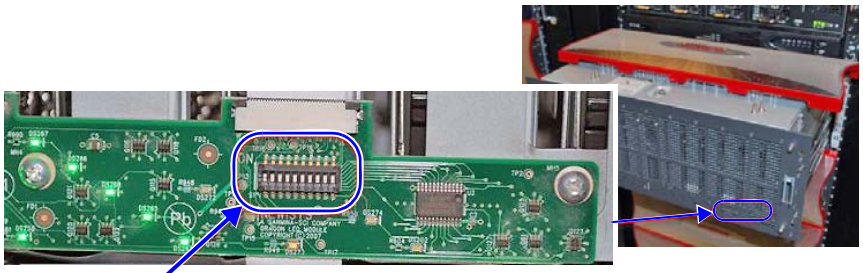

4. Verify that the piano switches are correctly set.

   For system with 5 disk enclosures (zoning configuration 1), the left-most piano switch should be up and the rest down. This divides the disk slots into two logical enclosures.

   For system with 10 disk enclosures (zoning configuration 0), all the piano switches should be down. This connects all disk slots into one logical enclosure.

   If a change is needed, remove the Mylar switch cover with a small pair of tweezers and then use a dental tool to move the switches. The tweezers and dental tool can be inserted through the grey bezel air holes. If changing the switches with a dental tool proves to be too difficult, then the grey bezel can be removed.

5. Replace the front bezel (and the grey bezel, if previously removed).

6. Push the enclosure back into the rack.

7. Repeat the above steps on all disk enclosures.

2.6 Cable Connections

2.6.1 Couplet ICL Cabling

There are two sets of Inter-Controller Link (ICL) connections between the two Controllers.

1. Verify that the two InfiniBand cables are attached to the ICL ports on the two Controllers as shown in Figure 37.

2. Verify that the four short Ethernet cables are attached to the Ethernet ports on the two Controllers as shown in Figure 37.

Figure 37. Inter-Controller Link Connections on Controllers

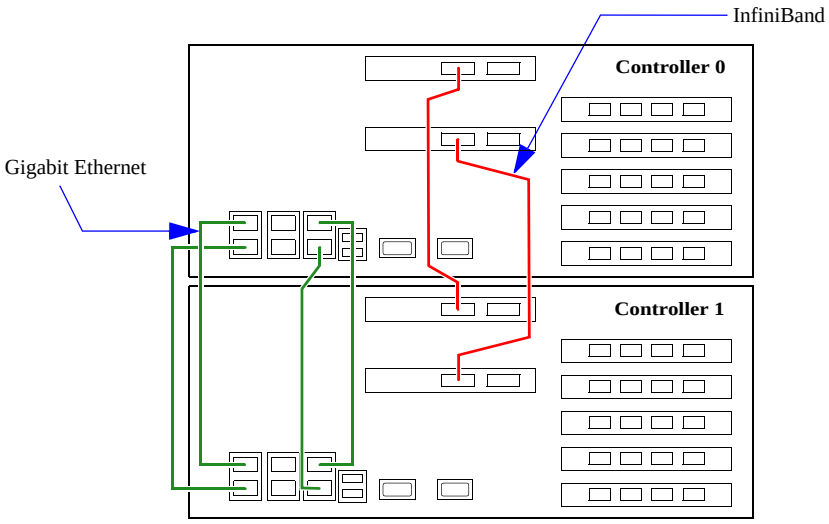

2.6.2 Disk Enclosure Cabling

The IS16000 systems are shipped with the cables attached between the disk enclosures and Controller's I/O channels. Incorrect wiring can prevent the system from operating correctly or from operating at all.

Verify that the disk enclosures are correctly connected to the two Controllers and secured using Figure 38, Figure 39, Figure 40, and Figure 41. The cables are labeled by ports to which they will be connected. In the illustrations, the port numbers are colored to aid in locating matching cables.

Figure 38	Five 60-bay disk enclosures containing *ALL* SAS disks
Figure 39	Five 60-bay disk enclosures containing *ALL* SATA disks or a mixture of SAS and SATA disks
Figure 40	Ten 60-bay disk enclosures containing ALL SAS disks
Figure 41	Ten 60-bay disk enclosures containing *ALL* SATA disks or a mixture of SAS and SATA disks

Figure 38. Connecting 5 Disk Enclosures (with SAS Disks Only) to Controllers

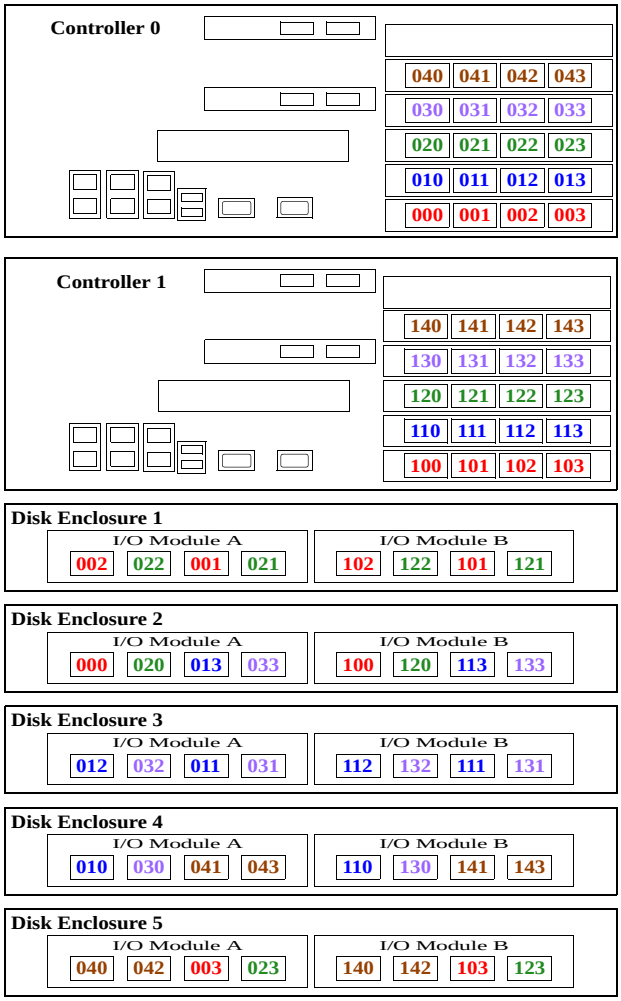

Figure 39. Connecting 5 Disk Enclosures (with SATA and SAS Disks) to Controllers

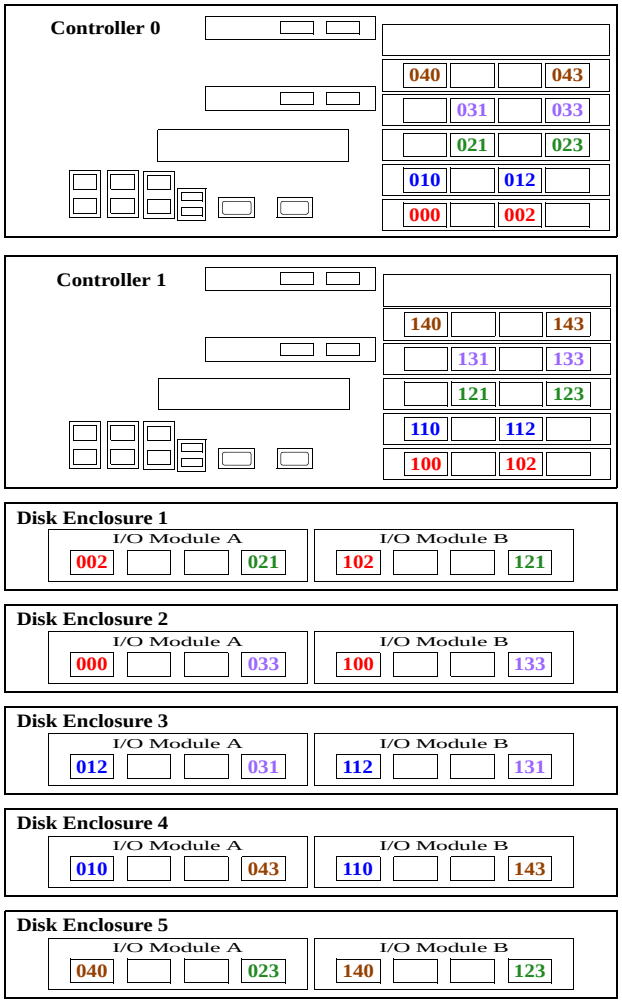

Note: Do NOT connect the cables to the ports that are not labeled.

Figure 40. Connecting 10 Disk Enclosures (with SAS Disks Only) to Controllers

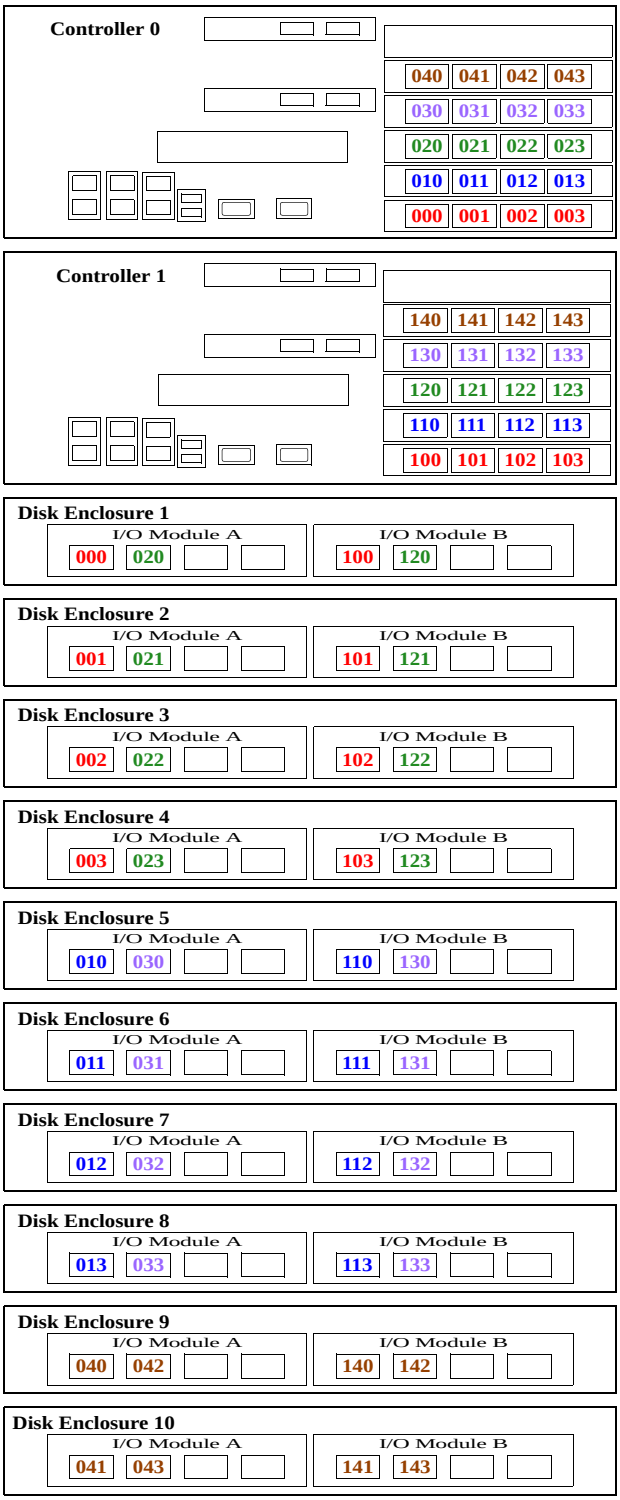

Figure 41. Connecting 10 Disk Enclosures (with SATA and SAS Disks) to Controllers

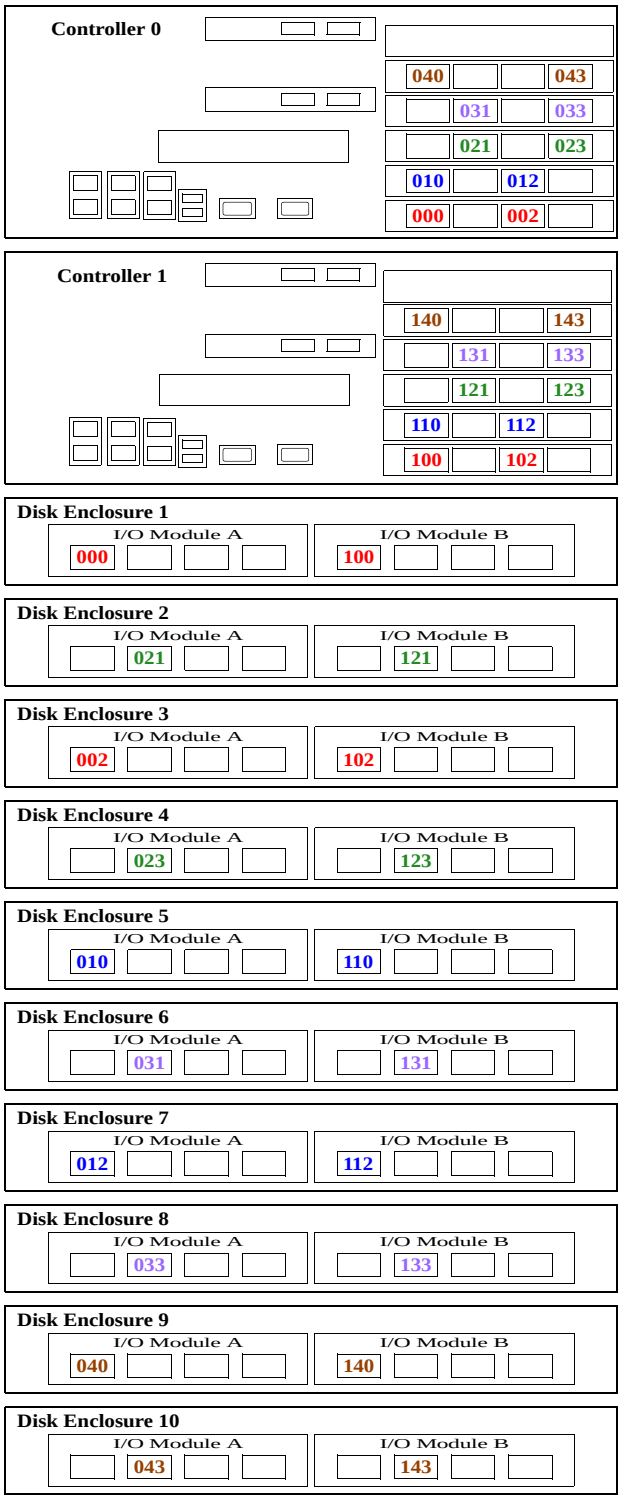

Note: Do NOT connect the cables to the ports that are not labeled.

2.6.3 Management Network Connection

You may remotely monitor the system over your Ethernet network. Connect the Controllers to your network using the Ethernet port as shown in Figure 42. The Ethernet management port supports 10, 100, and 1000BASE-T rates.

Figure 42. Ethernet Connections to Your Network

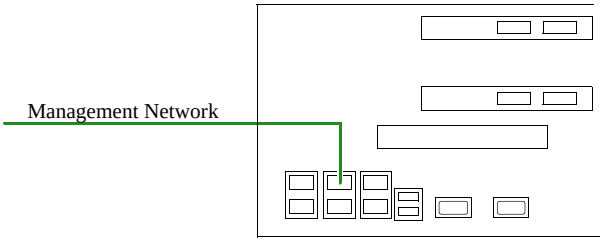

2.6.4 Host Connections

Connect the Controller host ports shown in Figure 43 either directly to your hosts or to a switch that connects your hosts. Depending on your IS16000 model, these connections may be Fibre Channel or InfiniBand.

NOTE : Do NOT use the unused Inter-Controller Link connectors to connect InfiniBand hosts or switches.

Figure 43. Host Ports on Controllers

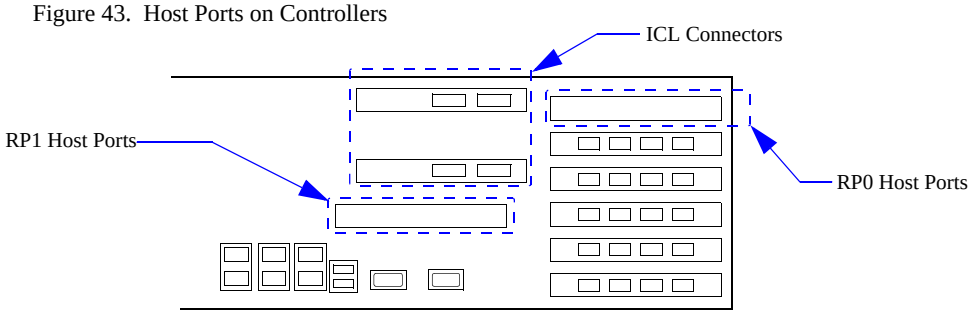

2.6.5 RS-232 Console Connection

Connect a null modem cable between a PC and the RS-232 connector on the back of the Controller (Figure 44). The RS-232 console is used to configure the management network ports during initial configuration. However, it can also be used to log the console output and upgrade the BIOS/BMC/CONFIG firmware.

Figure 44. Connect RS-232 Console to Controller

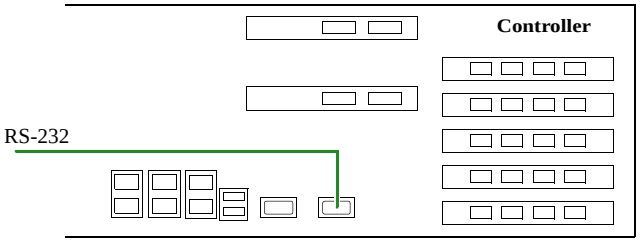

2.6.6 UPS Connection

1. For each Controller, verify that a USB cable is attached to the port on the back of the Controller and its UPS as shown in Figure 45.

2. For each Controller, verify that the two power cables are attached to the Controller's power supplies and its UPS as shown in Figure 45.

Figure 45. Connect UPS to Controller

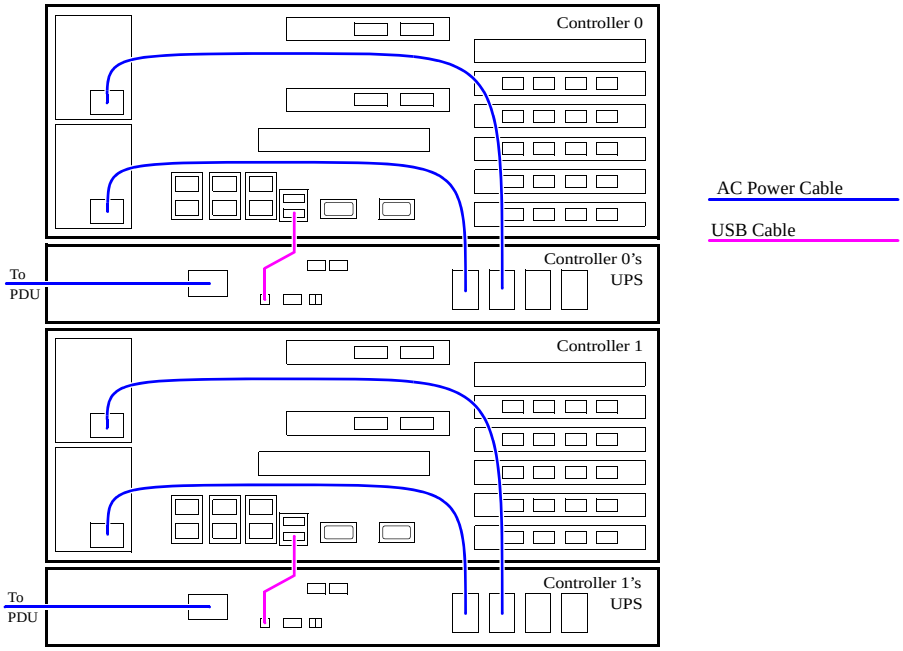

2.6.6.1 UPS Battery

The UPS units are shipped with the batteries disconnected. Follow these steps to connect the batteries:

1. Remove the front bezel of the unit.

2. Fasten the two connectors together (Figure 46). ***Note that it is normal if you see a spark and hear a pop sound as you connect the battery.***

3. Replace the front bezel.

4. Repeat the above steps on the other UPS.

Figure 46. Connect UPS Battery

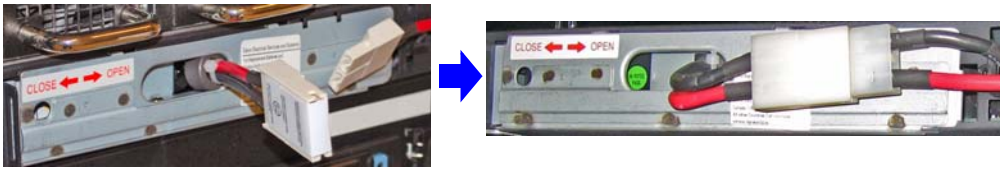

2.6.7 Power Connections

1. Verify your AC power source by measuring the AC voltage.

2. Refer to the illustrations in Figure 47 and verify that the DIP switches on the back of the UPS units are correctly set. If the settings are incorrect, have an SGI field engineer correct them.

NOTE : Changing the voltage of the UPS requires more steps than just changing the DIP switches.

Figure 47. DIP Switch Settings on UPS Rear Panel

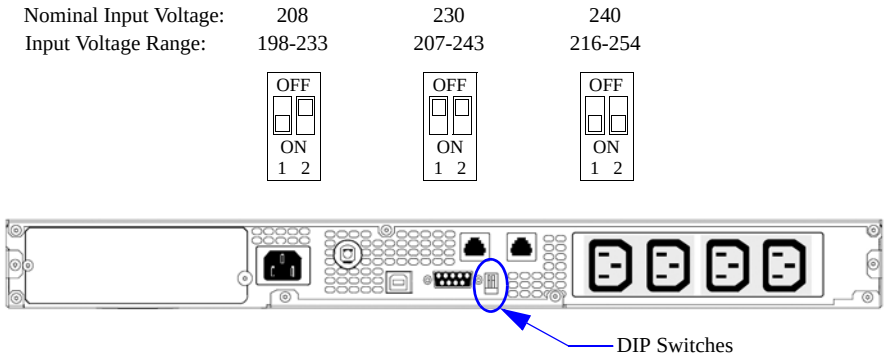

3. Verify that the power switches on all the power distribution units (PDUs) are set to OFF.

4. Connect the PDUs to your AC power source. For maximum redundancy, connect the PDUs to different AC circuits.

2.7 Powering On the System

1. Switch on all the PDUs.
2. Switch on the power supply modules on all the disk enclosures.
3. Press and hold the Power button for 2 seconds on the front panel of both UPS to turn on the units (Figure 48).

Figure 48. UPS Front Panel

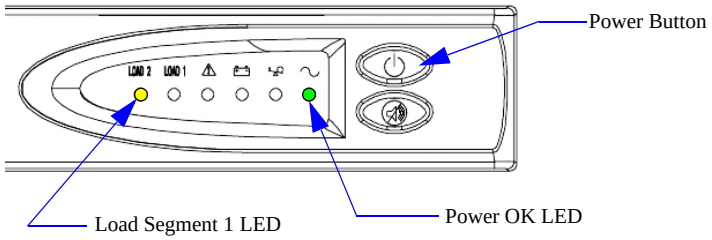

4. Verify that the Power OK LED on both UPS units turns green and the Load Segment 1 LED turns yellow (Figure 48), indicating a successful power application.
5. Verify that the Power indicator on both Controllers turns green (Figure 49). If not, press the Power button once.
6. Verify that the Fault LED on both Controllers is off (Figure 49).

Figure 49. Controller Front Panel LEDs

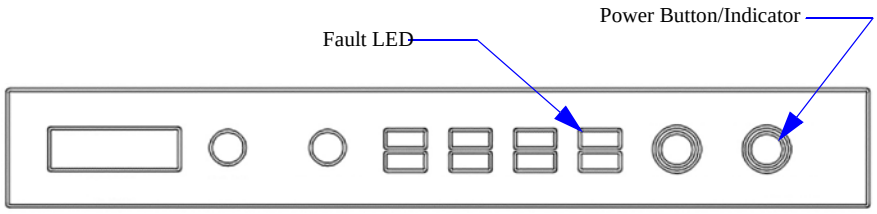

You may now configure the system as described in Section 2.8 "Configuring the IS16000".

2.8 Configuring the IS16000

This section provides information on configuring your IS16000.

NOTE : The configuration examples provided here represent only a general guideline. These examples should not be used directly to configure your particular IS16000.

The CLUI (Command Line User Interface) commands used in these examples are fully documented in Chapter 3—however, exact commands may change depending on your firmware version. To access the most up-to-date commands, use the CLUI's online HELP feature.

The CLUI commands are independent of case. Most of the keywords can be abbreviated and most of the punctuations are optional. For example, "SHOW VD 0 ALL" is adequate for "SHOW VIRTUAL_DISK=0 ALL_ATTRIBUTES"

2.8.1 Planning Your Setup and Configuration

Before proceeding to configure the storage settings for the IS16000, it is necessary to understand the basic organization of the system.

2.8.1.1 Storage Pools and Spare Pools

The disks in the system are categorized into one of the following pools:

- Unassigned Pool - By default, all the disks or replacement disks are initially allocated to this pool. When a disk has failed or been released (from a Spare Pool, for example), it is also moved to this pool.
- Storage Pool (sometimes referred to as a RAID group) - A collection of 2 to 10 disks, ideally of the same capacity and type.
- Spare Pool - This type of pool contains disk(s) that can be used as spare disks in one or more Storage Pools.

2.8.1.2 Storage Pools and Virtual Disks

A Virtual Disk is the storage unit presented to any attached host. Virtual Disks allocate space in 8 GiB increments. For example, 16 GiB of storage space will be allocated when creating a Virtual Disk of 10 GiB.

The IS16000 uses Storage Pools and Virtual Disks to configure disk storage for use by host systems. A Storage Pool (sometimes referred to as a RAID group) is a collection of 2 to 10 physical disks, ideally of the same capacity and type. In RAID5, the capacity of one disk is reserved for parity, allowing data recovery in the event of a single disk loss in the Storage Pool. In RAID6, the capacity of two disks is reserved for parity, allowing data recovery if either one or two disks are down in a Storage Pool. In RAID1, the capacity of one disk is used for data duplication.

A single RAID1 Storage Pool can be configured with 2 disks. A single RAID5 Storage Pool can be configured using 5 or 9 disks. A single RAID6 Storage Pool can be configured using 6 or 10 disks.

NOTE : Although a Storage Pool can be configured with disks of different capacity, in which case the IS16000 will use the lowest capacity for all disks, this is ***NOT*** recommended.

2.8.1.3 WWN Node Names for Controller Ports

If your SAN configuration requires unique world-wide name (WWN) node names on each controller port, set the controller channles (ports) to MAC mode. Use the following command:

APPLICATION SET CHANNEL [0|1|2|3] MODE MAC

See section 3.4.5, "Additional Configuration Considerations for Macintosh Hosts"and other Mac OS notes for details.

2.8.2 Disk Module Layout

Figure 50 shows the layout of disks in a disk enclosure. Disks are always referenced by enclosure ID and slot number as illustrated. The disk layout may also be found inside the covers of the enclosure.

The IS16000 supports any mix of SAS and SATA disks. For most efficient airflow and reduced vibrations, SAS disks should be placed in the front of the enclosure (lower slot numbers) while SATA disks should be placed in the rear of the enclosure (higher slot numbers).

Figure 50. Disk Enclosure Disk Module Layout

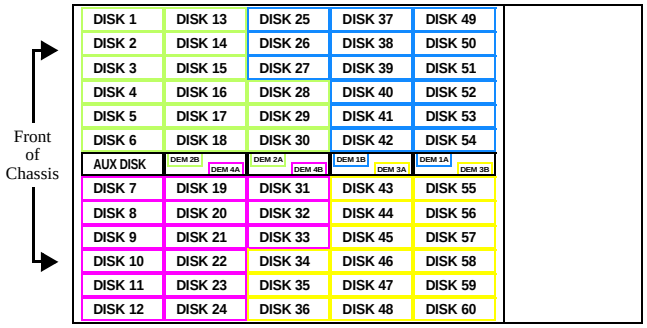

2.8.3 Serial Interface Configuration

1. Connect the supplied null modem cable between a PC and the RS-232 connector on the back of the Controller (Figure 51).

Figure 51. RS-232 Port on Controller

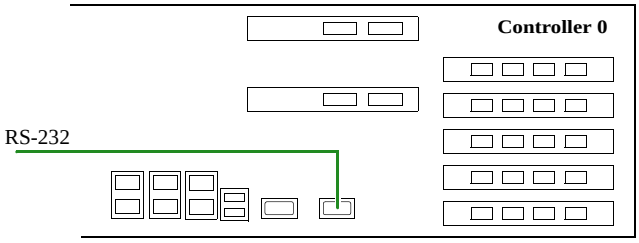

2. Load a serial console program (such as HyperTerminal, minicom, and screen) and use the following settings for the serial connection:

- Bits per second: 115,200
- Data bits: 8
- Parity: None
- Stop bits: 1
- Flow control: None

3. Once connected, press the <Enter> key to bring up the login prompt. Enter the user name **user** and password **user**.

2.8.4 Validate the Hardware

1. Enter: **SHOW CONTROLLER ***
   Verify that both Controller can be seen (Figure 52).

   Figure 52. Show Controller Status

```
RAID[0]$ show controller *

OID: 0x38000000 Index: 0000 Name: A LOCAL PRIMARY
OID: 0x38000001 Index: 0001 Name: B REMOTE SECONDARY

Total Controllers: 2
```

2. Enter: **SHOW ENCLOSURE ***
   Verify that all the disk enclosures can be seen and contain consistent firmware versions.

3. Enter: **SHOW FAN ***
   Verify that no fan faults are found.

4. Enter: **SHOW POWER ***
   Verify that the “AC Fail” and “DC Fail” states are false.

5. Enter: **SHOW TEMPERATURE ***
   Verify that all the temperature sensors do not report over temperature conditions.

6. Enter: **SHOW PHYSICAL_DISK ***
   Verify that all the disks can be seen and are healthy (Figure 53).

   Figure 53. Show Physical Disk Status

```
RAID[0]$ show pd *
                                  |Health|
Encl|Slot|Vendor| ::::::::::::::::State|Idx|State|::::::
-----------------------------------------------------------
   1    1 HITACHI ::::::::::::::::: GOOD    6 NORM  :::::
   1    2 HITACHI ::::::::::::::::: GOOD    1 NORM  :::::
   1    3 HITACHI ::::::::::::::::: GOOD    5 NORM  :::::
:::::::::::::::::::::::::::::::::::::::::::::::::::::::::::

Total Physical Disks:   300
Total Assigned Disks:   0
Total Unassigned Disks: 300
  Total SAS Disks:      300
  Total Member State Normal:  300
```

7. Enter: **SHOW UPS ***
   Verify that the UPS batteries are healthy. Also take note of the charge level as the battery may take several hours to fully charge.

2.8.5 *Clear System Configuration*

Warning

Use of the CLEAR SUBSYSTEM CONFIGURATION command will destroy any existing data.

To ensure that you are starting from a fresh/clean install, enter:

CLEAR SUBSYSTEM CONFIGURATION

Then enter **YES** twice to confirm deletion of current configuration (Figure 54).

Figure 54. Clear Subsystem Configuration

```
RAID[0]$ clear subsystem configuration
Are you sure you want to delete this configuration [No]? yes
Are you sure you want to delete this configuration?
All data will be lost [No]? yes
RAID SUBSYSTEM CONFIGURATION cleared STATUS='Success' (0x0)
```

2.8.6 *Set System Time & Date (NTP)*

NTP (Network Time Protocol) mode is available on the IS16000. You can enter up to four NTP addresses as the time servers.

To enable the NTP mode, enter the command:

SET SUBSYSTEM NTP=[<ip address list, up to 4>|NONE]

Specifying a list of NTP addresses will start the NTP mode on each Controller, using that list of NTP addresses as the time servers. Specifying **NONE** will turn off NTP mode.

To display the current settings, enter the command(Figure 55):

SHOW SUBSYSTEM ALL_ATTRIBUTES

Figure 55. Show NTP Settings

```
RAID[0]$ show subsystem all

******************
*    Subsystem     *
******************

RP Subsystem Name:      co-test-8
UID:                    60001ff0800a30000000000030000000
Subsystem Time:         Mon Sep 27 17:08:59 2010
Locate Dwell Time:      120 seconds
Enabled Licenses:       RAID6 SATASSURE
Fast Timeout:           ON
Pool Verify Priority:   10%
NTP Mode:               ON
                        (10.32.16.24)
                        (10.32.16.25)
                        (10.32.16.26)
                        (10.32.16.27)
Mon Sep 27 17:08:59 2010
```

If you are not using NTP, you can set the time of both Controllers using the **SET SUBSYSTEM DATE_AND_TIME** command. However, once the time is set, the time on the two Controllers is free to drift independent of the other Controller. To change the date and time to March 1, 2010 2:15 pm, for example, type: **SET SUBSYSTEM DATE_AND_TIME=2010:3:1:14:15**

NOTE : If NTP mode is enabled, the **SET SUBSYSTEM DATE_AND_TIME** command will fail with a status of "Setting date/time while in NTP mode".

2.8.7 Set the System Name

You may set the system name using the command (Figure 56):

SET SUBSYSTEM NAME=<name>

To verify the new setting, enter the command:

SHOW SUBSYSTEM ALL_ATTRIBUTES

Figure 56. Set Subsystem Name

```
RAID[0]$ set subsystem name=IS16K
SUBSYSTEM attributes set STATUS='Success' (0x0)

RAID[0]$ show subsystem all
RP Subsystem Name:      IS16K
UID:                    60000000000000000000000030000000
  .
  .
  .
  .
```

2.8.8 Configure Network Interface Settings

The IS16000 can be configured and administered either via serial connection (using the supplied serial cable) or via Ethernet connection using SSH. However, in order to use the SSH connection, it is first necessary to configure the network settings on each Controller. This can only be done using the serial interface as described below.

1. Enter the command:
   UI SET NETWORK_INTERFACE=LOCAL 0 IP_ADDRESS=<ip_address> IP_MASK=<netmask> IP_GATEWAY=<gateway>
   where <ip_address> is an address appropriate to the local network. The netmask and gateway values can also be entered (Figure 57).

   Figure 57. Set Network Interface Example

   ```
   RAID[0]$ ui set network_interface=local 0 ip_address=192.168.0.10
   ip_mask=255.255.255.0 ip_gateway=192.168.0.1
   NETWORK_INTERFACE 0 set with STATUS='Success' (0x0)
   ```

2. To verify the new settings, enter the command (Figure 58):
   UI SHOW NETWORK_INTERFACE=LOCAL 0 *

   Figure 58. Show Network Interface Settings

   ```
   RAID[0]$ ui show network_interface local 0 *
   Network device id 0
       address 192.168.0.10
       netmask 255.255.255.0
       gateway 192.168.0.1
   ```

3. Switch the serial cable to the RS-232 port on second Controller and repeat the above steps to enter a different IP address for second Controller.

At this point, you should be able to use the IS16000 CLUI and GUI via the network, and you will probably find that more convenient for the rest of the steps.

If you want to set up email and SNMP event notification, refer to Section 3.7.5, "Email and SNMP Notification Setup" on page 102 for information.

2.8.9 Create Storage Pools

A Storage Pool on a IS16000 has the following attributes:

- RAID Level (RAID)
  Storage Pools can be configured to use either a RAID1, RAID5 or RAID6 parity scheme. In RAID1, the capacity of one disk is used for data duplication. In RAID5, the capacity of one disk is reserved for parity, allowing data recovery in the event of a single disk loss in the Storage Pool. In RAID6, the capacity of two disks is reserved for parity, allowing data recovery if either one or two disks are down in a Storage Pool. ***For maximum data protection, use RAID6***.

- Chunk Size (CHUNK)
  The chunk size (in KiB blocks) defines the amount of data written to a single disk before proceeding to the next disk in the Storage Pool.

NOTE :RAID1 is a two member RAID set where the data is mirrored on each disk. There is no parity, hence, the chunk size is fixed.

- Disk Count (NUMBER)
  A RAID1 Storage Pool may consists of 2 physical disks. A RAID5 Storage Pool may consist of 5 or 9 physical disks. A RAID6 Storage Pool may consist of 6 or 10 physical disks. For maximum performance, select disks with the same characteristics (such as SAS/SATA, capacity, and RPM).

To create a Storage Pool, use the **`CREATE POOL`** command:

1. At the CLUI prompt, type:
`CREATE POOL RAID_LEVEL=[raid1|raid5|raid6]`
`CHUNK_SIZE=[32|64|128|256] {ASSIGN_POLICY=[SAS|SATA]`
`NUMBER=[2|5|6|9|10] or PHYSICAL_DISK=<list of disks>}`
`INIT_ALLOW=<initialization policy>`
`SATASSURE=[NONE|DATA_INTEGRITY_FIELD|PARITY]`
where <list of disks>=index name of disks; <initialization policy>=Allow_IO, NO_IO, Priority (1-99, where 99 puts all system resources on the initialization).

You may either explicitly select the disks for the pool by using "**`PHYSICAL_DISK=`**" or specify the number of disks in the pool by using "**`NUMBER=`**" in which case the next available disks will be selected (Figure 59). You may disable the SATAssure feature or enable it to use either the data integrity field option or the parity check on read option.

Figure 59. Create Storage Pool Example (1)

```
RAID[0]$ create pool raid_level=raid5 chunk_size=64kb physical_disk=0x6c, 0x6d, 0x6e, 0x6f, 0x70
POOL 0 OID=0x19b60000 create STATUS='Success' (0x0)

RAID[0]$ create pool raid_level=raid6 chunk_size=128kb number=6
POOL 1 OID=0x19b80001 create STATUS='Success' (0x0)
```

If you specify the **`ASSIGN_POLICY`** and **`NUMBER`** parameters together (Figure 60), only the selected type of disks will be used for the Storage Pool and you will not need to enter the individual disk name.

Figure 60. Create Storage Pool Example (2)

```
RAID[0]$ create pool raid_level=raid6 chunk_size=128kb number=6 assign_policy=sas
POOL 2 OID=0x19b80002 create STATUS='Success' (0x0)
```

2. Use the **CREATE POOL** command to add more Storage Pools as needed.
   If you need to delete a Storage Pool, use the **DELETE POOL <pool-id>** command.

NOTE : Storage Pool initialization is a background process and multiple Storage Pools can be initialized simultaneously. However, Virtual Disks on a Storage Pool are NOT accessible until its initialization is complete.

You may assign a name to a Storage Pool using the command:
SET POOL=<pool-id> NAME=<pool name>

To view the list of configured Storage Pools, enter the **SHOW POOL *** command (Figure 61).

Figure 61. Show Pool Information Screen

```
RAID[0]$ show pool *

                                                |Total |Free  |Max   |                    |Disk|  Global  |Spare |
Idx|Name      |State      |Chunk|Raid| Faults |cap GB|cap GB|VD GB | Settings  | Jobs  |T/O |spare pool|Policy|
-----------------------------------------------------------------------------------------------------------------
  0 pool-0      NORMAL       128   5             4104      0      0   W   R F I          10  UNASSIGNED  AUTO
  1 pool-1      NORMAL       128   5             4104      0      0   W   R F I          10  UNASSIGNED  AUTO
  2 pool-2      NORMAL       128   5             4104      0      0   W   R F I          10  UNASSIGNED  AUTO
  3 pool-3      NORMAL       128   5             6192      0      0   W   R F I I        10  UNASSIGNED  AUTO

Total Storage Pools: 6
```

To display the detailed information of a Storage Pool, use the **SHOW POOL <id> ALL_ATTRIBUTES** command (Figure 62).

Figure 62. Show Pool All_Attributes Example Screen

```
RAID$ show pool=0 all
  Index:                0
  OID:                  0x19e50000
  Type:                 STORAGE
  Name:                 pool-0
  Chunk Size:           128KB (0x100 blocks)
  Block Size:           512
  RAID Type:            RAID5
  Free Raid5 Capacity:  0 GB
  Max VD Size:          0 GB
  Total Capacity:       600 GB
  UUID:                 60000000000000000000000180d0000
  Global Spare Pool:    UNASSIGNED
  DiskTimeout(FRT):     10 minutes
  Init Policy:          ALLOW_IO
  Init Priority:        50%
  Full Rebuild Priority:        80%
  Fractional Rebuild Priority:  90%
  Sparing Policy:       AUTOMATIC
  Verify Policy:        DISABLED
  Assign Policy:
    Device Type:        SATA
    Rotation Speed:     NA
    Raw Capacity:       NA
  SATAssure:            None
  Cache Settings:
    ReACT:              FALSE
    IO Routing:         TRUE
    Mirroring:          TRUE
    Read Ahead:         TRUE
    Write Back:         FALSE
  Initializing:         TRUE
  Rebuilding:           FALSE
  Paused:               FALSE
  AutoWriteLock:        FALSE
  Critical:             FALSE
  Forced Write-Thru     FALSE
  Current Home:         0x38000000 0x00000000 (LOCAL)
  Future Home:          0xffffffff  0x00000000
  Preferred Home:       0x38000000 0x00000000 (LOCAL)
  BkgdJob OID:          0x28000003
  BkgdJob Priority:     50%
  Total Phy Disks       5
  State:                NORMAL
  Member Size:          120 GB
    pID     State     UUID
    0x0001  NORM      0x5000cca215c56e47
    0x0002  NORM      0x5000cca215c56456
    0x0003  NORM      0x5000cca215c54c71
    0x0004  NORM      0x5000cca215c5675c
    0x0005  NORM      0x5000cca215c56e55
```

2.8.10 Create Virtual Disks

A Virtual Disk can be created to use all or a part of a Storage Pool.

To create a Virtual Disk, use the **CREATE VIRTUAL_DISK** command:

1. At the CLUI prompt, type:
   CREATE VIRTUAL_DISK POOL=<pool-id> CAPACITY=<capacity>

   where <pool-id> is the Storage Pool to be used by this Virtual Disk; <capacity> is the capacity of the Virtual Disk in GiB or type "max" to use all available capacity (Figure 63).

   A message is displayed to indicate whether the Virtual Disk creation was successful.

   Figure 63. CREATE VIRTUAL DISKS Example Screen

```
RAID[0]$ create virtual_disk pool=0 capacity=16
VIRTUAL_DISK 0 OID=0x89ba000 creation STATUS='Success' (0x0)
RAID[0]$ create virtual_disk capacity=32 pool=0
VIRTUAL_DISK 1 OID=0x89bb001 creation STATUS='Success' (0x0)
```

2. Use the **CREATE VIRTUAL_DISK** command to add more Virtual Disks as needed.

   If you need to delete a Virtual Disk, use the **DELETE VIRTUAL_DISK <virtual disk-id>** command.

You may assign a name to a Virtual Disk :
SET VIRTUAL_DISK=<virtual disk-id> NAME=<virtual disk name>

To view the list of configured Virtual Disks, use the command **SHOW VIRTUAL_DISK *** (Figure 64).

Figure 64. Show Virtual Disk Information Screen

```
RAID[0]$ show virtual_disk *

                                                                    |          Home          |Background|
Idx|Name           | State |Pool|Raid|Cap GB|Settings| Jobs       |Current|Preferred|Future|   Job    |
-------------------------------------------------------------------------------------------------------
  0 vd-0_0          READY     0  5     3632   W   I                0(L) 0   0(L) 0   None   INACTIVE
  1 vd-1_1          READY     1  5     3632   W   I                0(L) 0   1(R) 0  1(R) 0  INACTIVE
  2 vd-2_2          READY     2  5     3632   W   I                0(L) 0   0(L) 0   None   INACTIVE
  3 vd-3_3         NOT RDY    3  5     5480   W   I   I            0(L) 0   1(R) 0  1(R) 0 0x28000003

Total Virtual Disks: 4
```

To display the detailed information of a Virtual Disk, use the command:
SHOW VIRTUAL_DISK=<id> ALL_ATTRIBUTES (Figure 65).

Figure 65. Show Virtual_Disk All_Attributes Example Screen

```
RAID[0]$ show vd=0 all
  Index:              0
  OID:                0x880f0000
  Name:               vd-0_0
  Pool Index:         0
  Pool OID:           0x180d0000
   .
   .
   .
  Current Home:       0x38000000  0x00000000 (LOCAL)
  Future Home:        0xffffffff  0x00000000
  Preferred Home:     0x38000000  0x00000000 (LOCAL)
  BkgdJob OID:        INACTIVE
  UUID:               60000000000000000000000000880f0000
```

2.8.11 Create and Assign Spare Pools

The IS16000 supports the concept of Spare Pool. A Spare Pool contains physical disks that can be used as spare disks in one or more Storage Pools.

In the event of disk failure, a disk that is assigned to a Spare Pool is automatically swapped into the Storage Pool to replace the failed disk. A rebuild then occurs immediately, reducing the possibility of data loss due to additional disk failures.

Spare Pools can be shared by Storage Pools or dedicated to a particular Storage Pool in which case the spare disks will only be used by the designated Storage Pool. After a Spare Pool is created, you will need to select disk(s) from the Unassigned Pool and assign it to the Spare Pool.

Follow these steps to create a Spare Pool:

1. At the CLUI prompt, enter the command: **CREATE SPARE_POOL** (Figure 66)

   A message is displayed to indicate whether the new Spare Pool creation was successful.

   Figure 66. Create Spare Pool Example Screen

   ```
   RAID[0]$ create spare_pool
   SPARE POOL 6 OID=0x19bc0006 create STATUS='Success' (0x0)
   ```

2. Assign physical disk(s) to the Spare Pool. Enter the command:
   ASSIGN PHYSICAL_DISK=<disk-id> TO_POOL=<pool-id>
   where <disk-id> is the index name of the disk to be added and <pool-id> is the index name of the Spare Pool.

   A message is displayed to indicate whether the disk has been added to the Spare Pool (Figure 67).

   Figure 67. Assign Physical Disk to Spare Pool Example Screen

   ```
   RAID[0]$ assign physical_disk=648 to_pool=6
   PHYSICAL_DISK 648 OID=0x22b10288 assigned to POOL 6 OID=0x19bc0006STATUS='Success' (0x0)
   ```

   To view the list of available disks, enter the command:
   SHOW UNASSIGNED_POOL PHYSICAL_DISK

3. Use the **CREATE SPARE_POOL** and **ASSIGN PHYSICAL_DISK** commands to create more Spare Pools as needed.

   If you need to delete a disk from the Spare Pool, enter the command:
   ASSIGN PHYSICAL_DISK=<disk-id> TO_POOL=0x1800ffff

   To delete a Spare Pool, enter the command:
   DELETE SPARE_POOL=<pool-id>

4. You may assign a name to a Spare Pool using the command:
   SET SPARE=<spare pool-id> NAME=<name>

5. To view the list of configured Spare Pool(s), enter the command:
   SHOW SPARE_POOL * (Figure 68).

Figure 68. Show Spare Pool Example Screen (1)

```
RAID$ show spare_pool *
                                |Disk|Total |Total|Storage|
Idx|Name            |Blocks|T/O|Cap GB| PDs |Pool Idx|
-------------------------------------------------------
  6 spare_pool-1       512   10    120    1
Total Spare Pools: 1
```

To display the detailed information of a Spare Pool (Figure 69), use the command:
SHOW SPARE_POOL=<id> ALL_ATTRIBUTES

To display the list of disks in a Spare Pool (Figure 69), use the command:
SHOW SPARE_POOL=<id> PHYSICAL_DISKS

Figure 69. Show Spare Pool Example Screen (2)

```
RAID$ show spare_pool=6 all_attributes
  Index:               6
  OID:                 0x19bc0006
  Type:                GLOBAL SPARE
  Name:                spare_pool-1
  Block Size:          512
  DiskTimeout(FRT):    10 minutes
  Total Capacity:      120 GB
  UUID:                0x0000000000000000
  Total Phy Disks      1
  Storage Pool List:
Total Spare Pools: 1

RAID$ show spare_pool=6 physical_disks
                                                        |Health|
Encl|Slot|Vendor|Product ID| ::::::::::|Pool|State | Idx|State|WWN|
-------------------------------------------------------------------
   2   53 HITACHI : : : : : : : : : :    6 SPARE   648 NORM  :::
```

6. Assign Spare Pools to Storage Pools. Enter the command:
SET POOL=<pool-id> SPARE_POOL=<spare-pool-id>
where <pool-id> and <spare-pool-id> are the index names of the Storage Pool and Spare Pool respectively. Repeat this step for each Storage Pool (Figure 70).

Figure 70. Set Storage Pool to Spare Pool Example Screen

```
RAID[0]$ set pool 0 spare_pool 6
POOL 0 OID=0x19b40000 attributes set with STATUS='Success' (0x0)
```

NOTE : Each Storage Pool should have a Spare Pool assigned to it. If a Storage Pool's attributes show "Global Spare Pool: UNASSIGNED", no Spare Pool will serve this Storage Pool.

Use the **SHOW POOL** command to view the new Storage Pool information (Figure 71).

Figure 71. Show Pool Example Screen

```
RAID$ show pool=1 all
  Index:               0
  OID:                  0x180d0000
  Type:                 STORAGE
  Name:                 pool-0
   :
   :
  Total Capacity:       600 GB
  UUID:                 60000000000000000000180d0000
  Global Spare Pool:    6
  DiskTimeout(FRT):     10 minutes
   :
   :
```

2.8.12 Present Virtual Disk to External Host

Virtual Disks are only presented to the hosts that have been given authorized access. A Presentation on a IS16000 has the following components:

- Discovered Initiators
  A Discovered Initiator is defined as a Fibre Channel or InfiniBand initiator that has logged into the IS16000. This is normally a single port on a Fibre Channel or InfiniBand HBA in an external computer.

- Host
  A Host is an abstract client of a Virtual Disk that has all the access to all data within the Virtual Disk. This most commonly represents an external computer with one or more Fibre Channel or InfiniBand initiators. Host attributes are OS (operating system) of which you can select GENERIC, WINDOWS, LINUX or MAC_OS. The OS attribute dictates the way the controller responds to certain Fibre Channel commands where appropriate. For example, Mac OS X expects a different response from an inquiry of LUN0 (the controller LUN) than a Windows host.

- Channel
  A Channel is one of the Fibre Channel or InfiniBand ports on the Controller. Channel attributes are MODE of which you can select either MAC_OS or STANDARD. When MAC_OS is selected, the port's node name will be set differently in order to be visible to a Macintosh system.

- Stack
  A Stack is the input side of the subsystem. In the case of the IS16000, the stack is defined as the Fibre Channel interface. In future versions of the product, different interface stacks may exist.

- Presentation
  A Presentation is the relationship between a Host and a Virtual Disk. A Presentation implies that the related Host has some sort of access to the Virtual Disk. Attributes of a Presentation are:
  - PORT: from which the host will see the Virtual Disk
  - READ_ONLY: controls read only access
  - PRESENT_HOME_ONLY: presents the specified Virtual Disk from its designated home controller only
  - LUN: user-specified LUN number that the Virtual Disk will show to the host.

Follow these steps to set up a presentation:

1. To display the currently available initiators, enter the command (Figure 72):
`APPLICATION SHOW DISCOVERED_INITIATORS *`

Figure 72. Display Available Initiators

```
RAID$ app show discovered_initiators *

                |          Initiator Identifier          | Controller 0  | Controller 1  |
Index |Type|   ID   |    node           |      port         | RP 0  | RP 1  | RP 0  | RP 1  |
-------------------------------------------------------------------------------------------
00001   FC |0x0000e8 0x2001001b32ae096c 0x2101001b32ae096c          |       |     3 |
00002   FC |0x0000e8 0x2001001b32aeb580 0x2101001b32aeb580          |       |       |0
00003   FC |0x0000e8 0x2001001b328e0280 0x2001001b328e0280          |       |    2  |
00004   FC |0x0000e8 0x2001001b32ae176c 0x2001001b32ae176c          |       |       |  1
00005   FC |0x0000e8 0x2001001b328eb580 0x2001001b328eb580         |0      |       |
00006   FC |0x0000e8 0x2001001b32aeb280 0x2001001b32aeb280       2  |       |       |
00007   FC |0x0000e8 0x2001001b328e176c 0x2001001b328e176c          |  1    |       |
00008   FC |0x0000e8 0x2001001b328e096c 0x2101001b328e096c        3|       |       |

Total FC Initiators: 8
```

2. Create a host. Enter the command (Figure 73):
APPLICATION CREATE HOST NAME=<host name>
OSTYPE=[LINUX|WINDOWS|MAC_OS|DEFAULT|GENERIC]
where <host name> is an assigned host name to help make mapping simpler for the user; <os type> is the mode which can be set to characteristics specific to an Operating System, especially for Mac OSX.

   Use the **APPLICATION CREATE HOST** command to create more hosts as needed.

   If you need to delete a host, use the **APPLICATION DELETE HOST=<host-id>** command.

   To display the current settings, enter the command: **APP SHOW HOST ***

   Figure 73. Create Host Examples

```
RAID$ app create host name=server1-port1 ostype=linux
HOST 0 OID=0x18d0000 creation STATUS='Success' (0x0)

RAID$ app create host name=server1-port2 ostype=linux
HOST 1 OID=0x18e0001 creation STATUS='Success' (0x0)

RAID$ app show host *
                          |Stack| Host Mode Attributes  |
Index|Host Name           |Index|OS Type|Characteristics|
-------------------------------------------------------------
00000 server1-port1        00000 LINUX   0x000000000000001
00001 server1-port2        00001 LINUX   0x000000000000001

Total Hosts: 2
```

3. Map a host to a discovered initiator. Enter the command (Figure 74):
APPLICATION IMPORT DISCOVERED_INITIATOR=<initiator-id>
HOST=<host-id>
where <initiator-id> is the index name of the discovered initiator;
<host-id> is the index name of the host.

   Figure 74. Map a Host to a Discovered Initiator

```
RAID$ app import discovered_initiator=6 host=0

INITIATOR 0 OID=0x280f0000 imported from DISCOVERED_INITIATOR 6 oid=0X30000006
STATUS='Success' (0x0)
```

   Use the **APPLICATION IMPORT** command to map the other hosts to the initiators as needed.

   To display the current mappings, enter the command (Figure 75):
APPLICATION SHOW INITIATOR *

   Figure 75. Show Current Imported Initiators

```
RAID$ app show initiator *
           |Host |        Initiator Identifier         |
Index |Type|Index|      node        |      port        |
-------------------------------------------------------
00000   FC  00000 0x2001001b32aeb280 0x2001001b32aeb280

Total FC Initiators: 1
```

4. Present the Virtual Disks to the hosts. Enter the command (Figure 76):
APPLICATION CREATE PRESENTATION HOST=<host-id>
VIRTUAL_DISK=<vd-id> LUN=<LUN-id>

where <host-id> is the index name of the host; <vd-id> is the Virtual Disk to be presented; <LUN-id> is the LUN that the specified Virtual Disk will be presented as or use the default value if it is not specified.

Figure 76. Create a Presentation

```
RAID$ app create presentation host=0 vd=0 lun=0
PRESENTATION 0 OID=0x0x20110000 creation STATUS='Success' (0x0)
```

To simply present the Virtual Disk to all host ports for both Controllers (promiscuous mode), use this command:
APPLICATION CREATE PRESENTATION VIRTUAL_DISK=<vd-id> HOST=ALL

Use the **APPLICATION CREATE PRESENTATION** command to configure other presentations as needed.

To display the current settings, enter the command (Figure 77):
APPLICATION SHOW PRESENTATION *

Figure 77. Show Current Presentations

```
RAID$ app show presentation *
                                                |      Channel Mask         |
                                                | Controller 0|Controller 1 |
Pres.|Host           Host | VD  |     |Home|Read|  RP 0| RP 1 |  RP 0| RP 1 |
Index|Name           Index|Index| LUN |Only|Only| 0123 | 0123 | 0123 | 0123 |
-------------------------------------------------------------------------------
00000 server1-port1   00000 00000  000  OFF  R/W    f      f      f      f
Total Presentations: 1
```

2.8.12.1 Special Considerations for MAC OS

Apple's Mac OS X Server does not adhere to the FC specification with regards to node naming. Apple expects the node name to vary from port to port while the FC specification calls for it to remain constant. A mode setting is provided to force the node name to be unique on a given client channel (Fibre Channel port). This mode can be set for one or more client channels. While the feature is aimed at Mac SANs, it can also be used for connections to heterogeneous client SANs.

When creating the host, specify the **OSTYPE** to be MAC_OS:
APPLICATION CREATE HOST OSTYPE=MAC_OS

Refer to Section 3.4.5, "Additional Configuration Considerations for Macintosh Hosts" for more information.

2.8.12.2 Selective Presentation

By default, a LUN is accessible via all initiators/ports. You may mask a presentation and choose the specific port on which the initiator may have access to the LUN.

You may mask a presentation when you create it:
APPLICATION SET PRESENTATION HOST=<host-id> VIRTUAL_DISK=<vd-id> LUN=<LUN-id> ENABLE=<mask option>

By selecting the **ENABLE** option, you are choosing the specific port on which the initiator may have access:

Controller 0 RP0: **ENABLE=0**
Controller 0 RP1: **ENABLE=1**
Controller 1 RP0: **ENABLE=2**
Controller 1 RP1: **ENABLE=3**

To mask a LUN to all ports, use **ENABLE=ALL**.

To change a mask, you must first reset it with **ENABLE=NONE**.

To mask a LUN to multiple ports, you must set one port at a time. For example, to mask the same LUN 4 to Controller 0, RP0 and Controller 1, RP1:

```
APP SET PRESENTATION=4 ENABLE=0
APP SET PRESENTATION=4 ENABLE=3
```

2.8.13 *Storage Pool Initialization*

When a Storage Pool is created, initialization begins automatically and will continue in the background until it is completed. While its Virtual Disk(s) can be accessed immediately, the performance of the virtual disks will be degraded while it continues to initialize.

To monitor the progress of a job, use the **SHOW JOB * ALL_ATTRIBUTES** command. It will display the type of job and the percentage of completion (Figure 78).

Figure 78. Show Job Example Screen

```
RAID[0]$ show job * all_attributes
  OID:                0x2e040001 (Index:1)
  Target:             POOL:0x1aa00001 (Index:1)
  Sub-Target:         NA
  Type:               INITIALIZE
  State:              RUNNING
  Completion Status:UNKNOWN
  Priority:           50
  Fraction Complete:71%
  Time:               NA
```

You may also check if a Virtual Disk is ready for access using the **SHOW VIRTUAL_DISK <id> ALL_ATTRIBUTES** command (Figure 79).

Figure 79. Show Virtual Disk Example Screen

```
RAID[0]$ show virtual_disk 0 all_attributes
  OID:                0x89ba0000
  Name:               89ba0000
  POOL OID:           0x19b40003
  Capacity:           16384 MBs
  Offset:             0x0
  State:              READY
  Raidlevel:          RAID5
  IO ROUTING:         TRUE
  WBC:                TRUE
  MWBC:               FALSE
  Initializing:       FALSE
  Paused:             FALSE
  AutoWriteLock:      FALSE
  Data Lost:          FALSE
  Present Home Only:FALSE
   .
   .
   .
```

CHAPTER 3

Administration

3.1 Managing the IS16000

The IS16000 provides a set of tools that enable administrators to centrally manage the network storage and resources that handle business-critical data. These include Configuration Management, Performance Management, and Firmware Update Management.

NOTE : The command descriptions given in this User Guide are specific to Firmware Version 1.3.0.4.xxxx. To access the most up-to-date commands, use the CLUI's Online Help feature.

3.1.1 Management Interface

SAN management information for the IS16000 can be accessed locally through a serial interface, or remotely through SSH.

Locally via Serial Interface

Any RS-232 terminal or terminal emulator (such as Hyperterminal) can be used to configure and monitor the IS16000.

1. Connect the supplied serial cable to the RS-232 port on the Controller (Figure 80). Connect the other end of the cable to a serial port on a standard PC.

Figure 80. RS-232 Port on Controller

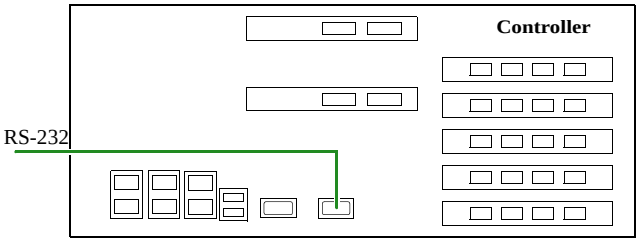

2. Load a serial console program (such as HyperTerminal, minicom, and screen) and use the following settings for the serial connection:
   - Bits per second: 115,200
   - Data bits: 8
   - Parity: None
   - Stop bits: 1
   - Flow control: None
3. Once connected, press the <Enter> key to bring up the login prompt. Enter the user name **user** and password **user**.

Remotely via SSH

To configure and monitor the IS16000 remotely, connect the system to your Ethernet network. Please refer to Section 3.7 for further information on how to set up the Controller's network interface.

3.1.2 User Logins

RS-232 Login

For a terminal session, press the <Enter> key to bring up the login prompt. Enter the user name **user** and password **user**.

Login via SSH

For a SSH session (Figure 81), enter the user name **user** and password **user**. Both are case sensitive.

Figure 81. SSH Login Screen

```
login as: user
user@10.23.23.16's password:

Linux (none) 2.6.25-sgi-016620-3 #2 SMP Mon Aug 16 10:38:29 MST 2010 x
SGI IS
```

Logout

To logout, enter the command **QUIT**. For SSH connection, the current session will be disconnected.

3.1.3 Available Commands

The CLUI commands are independent of case. Most of the keywords can be abbreviated and most of the punctuations are optional.

Entering the "**?**" character at the CLUI prompt will display the available commands within the IS16000 CLUI (Figure 82). To get help information on a command, type the command followed by a question mark.

For example, **ASSIGN ?**<Enter> will display help on "assign" options on the system.

Figure 82. Help Screen

```
RAID[0]$ ?
Select one of the following subjects for additional commands:
   APPLICATION      Selects the application processor as the subject
   RAID             Selects the raid processor as the subject
   UI               Selects the user interface as the subject

Additional commands:
   HELP             Provides information on how to use the user interface help
   QUIT             Exits the Command Line User Interface

Object class options available for default subject RAID:
Available keywords:
   ASSIGN           Assign raid object
   CLEAR            Clear raid object states
   CREATE           Create raid objects
   DELETE           Delete raid objects
   LOCATE           Locate raid objects
   MOVE_HOME        Locate raid objects
   PAUSE            Pause raid objects
   REPLACE          Replace raid objects
   RESUME           Resume raid objects
   SET              Set raid objects
   SHOW             Display raid objects
   SHUTDOWN         Shutdown raid objects
   SYNCHRONIZE      Synchronize raid objects
   UPDATE_FIRMWARE  Update firmware on raid objects
   VERIFY           Start a Verify background job(s)
```

3.1.3.1 Basic Key Operations

The command line editing and history features support **ANSI** and **VT-100** terminal modes. The command history buffer can hold up to 64 commands. The full command line editing and history only work on main CLI and SSH sessions when entering new commands. Basic Key Assignments are listed in Figure 83.

Simple, not full command, line editing only is supported when the:

- CLUI prompts the user for more information.
- alternate CLUI prompt is active. (The alternate CLUI is used on the RS-232 connection during an active SSH session.)

NOTE : Not all SSH programs support all the keys listed in Figure 83.

Figure 83. Basic Key Assignments

Key	ANSI CTRL or Escape Sequence	Description
Backspace	Ctrl-H	deletes preceding character
Delete	Del, or Esc [3~	deletes current character
Up Arrow	Esc [A	retrieves previous command in the history buffer
Down Arrow	Esc [B	retrieves next command in the history buffer
Right Arrow	Esc [C	moves cursor to the right by one character
Left Arrow	Esc [D	moves cursor to the left by one character
Home	Esc [H or Esc [1~	get the oldest command in the history buffer
End	Esc [K or Esc [4~	get the latest command in the history buffer
Insert	Esc [2~	toggles between insert mode and overtype mode
PgUp	Esc [5~	retrieves oldest command in the history buffer
PgDn	Esc [6~	retrieves latest command in the history buffer
Ctrl-U	Ctrl-U	delete to beginning of line
Ctrl-K	Ctrl-K	delete to end of line
Ctrl-A	Ctrl-A	move cursor to beginning of line
Ctrl-E	Ctrl-E	move cursor to end of line

3.2 Configuration Management

The disks in the system are categorized into one of the following pools:

- Unassigned Pool - By default, all the disks or replacement disks are initially allocated to this pool. When a disk has failed or been released (from a Spare Pool, for example), it is also moved to this pool.
- Storage Pool (sometimes referred to as a RAID group) - A collection of 2 to 10 disks, ideally of the same capacity and type.
- Spare Pool - This type of pool contains Physical Disks that can be used as spare disks in one or more Storage Pools.

3.2.1 Disk Information

To display the list of all the disks installed in the system, enter the command (Figure 84):
SHOW PHYSICAL_DISK *

Figure 84. Show Physical Disk List Example Screen

```
RAID[0]$ show pd *
                                                                                                                |Member|
Encl|Slot| Vendor |     Product ID     |Type|Cap GB |RPM|Revision|    Serial Number   |Pool|Status| Idx|State |       WWN       |
-----------------------------------------------------------------------------------------------------------------------------------
   1    1  HITACHI      HDS725050KLA360 SATA     465  7K K2AOAD1A       KRVN65ZAKAZLKH UNAS  GOOD    57  READY  5000cca20def4516
   1    2  HITACHI      HDS725050KLA360 SATA     465  7K K2AOAD1A       KRVN65ZAKNL79H UNAS  GOOD    32  READY  5000cca20df3a263
   1    3  HITACHI      HDS725050KLA360 SATA     465  7K K2AOAD1A       KRVP65ZAKD8Y5F UNAS  GOOD    42  READY  5000cca20defdc87
   1    4  HITACHI      HDS725050KLA360 SATA     465  7K K2AOAD1A       KRVN65ZAK72BLH UNAS  GOOD     1  NORM   5000cca20ded7e05
   .
   .
   .
Total Physical Disks:      120
Total Assigned Disks:      104
Total Unassigned Disks:    16
  Total SAS Disks:         60
  Total SATA Disks:        60
  Total Member State Normal:   104
```

To display the list of unassigned disks, enter the command:
SHOW UNASSIGNED_POOL PHYSICAL_DISK

To retrieve information about a specific disk, enter the command:
SHOW PHYSICAL_DISK=<disk-id> ALL_ATTRIBUTES (Figure 85).

Figure 85. Physical Disk Information Example Screen

```
RAID[0]$ show pd 40 all

  Index:            40
  OID:              0x27000028
  Pool Index:       UNASSIGNED
  Pool OID:         UNASSIGNED
  Capacity:         448 GB
  Raw Capacity:     465 GB(Base 2)/500 GB(Base 10)
  Block Size:       512
  Enabled Disk Ch:  0xffff  0x21d0  0xffff  0x2054
  Disk Slot:        46  (2:46)
  Vendor ID:        HITACHI
  Product ID:       HDS725050KLA360
  Product Revision: K2AOAD1A
  Serial Number:    KRVN65ZBGB8MTF
  Health State:     GOOD
  Rotation Speed:   7200 RPM
  Device Type:      SATA
  Member State:     UNASSIGNED
  Spare:            FALSE
  Failed:           FALSE
  UUID:             0x5000cca20ec52082
```

If there is a failed disk, use the **SHOW UNASSIGNED_POOL FAILED ALL** command to display the failed disk's information (Figure 86).

Figure 86. Show Failed Disk Example Screen

```
RAID[0]$ show unassigned_pool failed all

  OID:              0x20a7003f
  Pool OID:         UNASSIGNED
  Capacity:         704512 MBs (0x56000000 blocks)
  Raw Capacity:     715404 MBs (0x575466f0 blocks)
  Block Size:       512
  Enabled Disk Ch:  0x27  0x22
  Disk Slot:        1:42
  Vendor ID:        Hitachi
  Product ID:       Hitachi HUA721075KLA330
  Product Revision: GK8OAB0A
  Serial Number:          GTF200P8GBVPXF
  Health State:     FAILED
  Rotation Speed:   7200 RPM
  Device Type:      SATA
  Member State:     UNASSIGNED
  State:            READY
  Spare:            FALSE
  Failed:           TRUE
  UUID:             0x5000cca215c564560
```

3.2.1.1 Disk States

Listed below are the possible disk states:

NORM: Disk is in a normal or functional condition.

AMIS: Disk is Already MISsing. The disk in this state must return prior to the other disks. If the specific disk that is in this state never returns to the pool, the pool will remain inoperative. Replacing this disk to do a rebuild will do nothing.

WTRB: Disk is Waiting To ReBuild. A disk went missing then came back and is rebuilding (either partial or full). Then other disks in the Storage Pool went missing causing it to go inoperative, which halts the rebuild. Now waiting for other disks in the system to become active again.

MNRB: Disk is Missing with No ReBuild. This is the same as missing, but when the disk comes back, there is nothing to rebuild. This could happen when multiple disks fail at the same time making the Storage Pool inoperative. Reseating is an option here to see if the disk would come back. This is similar to the AMIS condition, but in this case it does not matter what order the MNRB disks become active.

MISS: Disk is MISSing. This is either not electrically active or seen, or only one Controller sees it in a couplet system. Reseating would be a first recommendation as a partial rebuild would happen if it then appears to both Controllers. Replacing would cause a full rebuild.

FAIL: Disk is FAILed and is unassigned. This means that the disk timeout has either expired or the disk had errors and had been failed. "FAIL" in the Storage Pool command means that it has been failed out of the Storage Pool, not necessarily the state of the disk. This is being renamed to EMPTY.

RBLD: Disk is ReBuiLDing.

PRTRDY: PaRTial ReaDY state is a condition where only one of the two Controllers is able to communicate with a drive.

3.2.1.2 Visual Indication

LOCATE PHYSICAL_DISK=<id> provides a visual indication of the specified disk. The status LED of the disk module will blink.

LOCATE UNASSIGNED_POOL provides a visual indication of the disk modules that are unassigned. The status LED of the disk modules will blink.

LOCATE UNASSIGNED_POOL FAILED provides a visual indication of the disk modules that have failed. The status LED of the disk modules will blink.

3.2.2 Storage Pool Management

The IS16000 creates centrally-managed and vendor-independent Storage Pooling. It enables different types of storage to be aggregated into a single logical storage resource from which virtual volumes (Virtual Disks) can be served up to multi-vendor host computers. The networked Storage Pools will provide the framework to manage the growth in storage demand from web-based applications, database growth, network data-intensive applications, and disaster tolerance capabilities.

3.2.2.1 Display Storage Pool Information

You can add and remove Storage Pools without affecting system operations. Use the **SHOW POOL *** command to display the current list of Storage Pools (Figure 87).

Figure 87. Show Storage Pool List Example Screen

```
RAID[0]$ show pool *
                                          |Total |Free  |Max   |                  |Disk|  Global  |Spare |
Idx|Name           |State     |Chunk|Raid| Faults |cap GB|cap GB|VD GB | Settings  | Jobs  |T/O |spare pool|Policy|
-----------------------------------------------------------------------------------------------------------
  0 pool-0         NORMAL      128   6              720      0      0 DWMRFI               1  UNASSIGNED  AUTO
  1 pool-1         NORMAL      128   6              720      0      0 DWMRFI               1  UNASSIGNED  AUTO
  2 pool-2         NORMAL      128   6              720      0      0 DWMRFI               1  UNASSIGNED  AUTO
  3 pool-3         NORMAL      128   6              720      0      0 DWMRFI               1  UNASSIGNED  AUTO
  4 pool-4         NORMAL      128   6              720      0      0 DWMRFI               1  UNASSIGNED  AUTO
  5 pool-5         NORMAL      128   6              720      0      0 DWMRFI               1  UNASSIGNED  AUTO

 Total Storage Pools: 5
```

NOTE : If you issue the command, **SHOW POOLS *** , and the pools indicate there is a fault, check the details of the pool. The pool will go into WR thru state on initial startup to allow the batteries the opportunity to charge. The WR Thru mode is necessary to allow for data integrity until battery redundancy on each controller is completed. If there is a battery issue, the system will indicate replacement is required and all VDs will remain in WR Thru mode until the issue is resolved.

The **SHOW POOL=<id> ALL_ATTRIBUTES** command displays the detailed information of the specified Storage Pool (Figure 88).

Figure 88. Show Storage Pool Attributes Example Screen

```
RAID[0]$ show pool 0 all

  Index:                0
  OID:                  0x1a520000
  Type:                 STORAGE
  Name:                 pool-0
  Chunk Size:           128KB (0x100 blocks)
  Block Size:           512
  RAID Type:            RAID6
  Free Raid6 Capacity:  0 GB
  Max VD Size:          0 GB
  Total Capacity:       720 GB
  UUID:                 60001ff0800a3000000000001a520000
  Global Spare Pool:    UNASSIGNED
   .
   .
   .
  BkgdJob OID:          INACTIVE
  Total Phy Disks       6
  State:                NORMAL
  Member Size:          120 GB
    pID     State     UUID
    0x01dd  NORM      0x5000cca00d273618
    0x01df  NORM      0x5000cca00d32293c
    0x0263  NORM      0x5000cca00d324438
    0x01e3  NORM      0x5000cca00d2b4c48
    0x01e1  NORM      0x5000cca005039324
    0x01e7  NORM      0x5000cca00d103e94

Total Storage Pools: 1
```

The **SHOW POOL=<id> PHYSICAL_DISK** command displays the list of disks associated with the specified Storage Pool (Figure 89).

Figure 89. Show Storage Pool Physical Disks Example Screen

```
RAID[0]$ RAID[0]$ show pool 0 pd
                                                                                                                                                |Health|
Encl|Slot| Vendor |     Product ID     |Type|Cap GB | RPM|Revision|    Serial Number   |Pool|State | Idx|State |       WWN       |
-------------------------------------------------------------------------------------------------------------------------------------------------------
   1   37 HITACHI  HUS153014VLS300      SAS     136  15K A410     JFVWMZNC5301           0  GOOD   611 NORM    5000cca00d324438
   1   39 HITACHI  HUS153014VLS300      SAS     136  15K A140     J4V1YYAA5301           0  GOOD   481 NORM    5000cca005039324
   1   51 HITACHI  HUS153014VLS300      SAS     136  15K A410     JFVWL5YC5301           0  GOOD   479 NORM    5000cca00d32293c
   2    1 HITACHI  HUS153014VLS300      SAS     136  15K A410     JFVSU5PC5301           0  GOOD   483 NORM    5000cca00d2b4c48
   2    3 HITACHI  HUS153014VLS300      SAS     136  15K A410     JFV8XYGC5301           0  GOOD   487 NORM    5000cca00d103e94
   2   42 HITACHI  HUS153014VLS300      SAS     136  15K A410     JFVPKHUC5301           0  GOOD   477 NORM    5000cca00d273618
Total Physical Disks:      6
Total Assigned Disks:      6
Total Unassigned Disks:    0
  Total SAS Disks:         6
  Total Member State Normal:   6
```

The **SHOW POOL=<id> VIRTUAL_DISK** command displays the list of Virtual Disks associated with the specified Storage Pool (Figure 90).

Figure 90. Show Storage Pool Virtual Disks Example Screen

```
RAID[0]$ show pool 0 vd

                                                                                |          Home   |Background|
Idx|Name            | State |Pool|Raid|Cap GB|Settings| Jobs  |Current|Preferred|    Job    |
-------------------------------------------------------------------------------------------------------------
  0 vd-0_0            READY     0  5     3632   W   I           0(L) 0   0(L) 0   INACTIVE

Total Virtual Disks: 1
```

3.2.2.2 Creating a Storage Pool

To create a Storage Pool, enter the command:
CREATE POOL RAID_LEVEL=[1|5|6] CHUNK_SIZE=[32|64|128|256] {ASSIGN_POLICY=[SAS|SATA] NUMBER=[2|5|6|9|10] or PHYSICAL_DISK=<list of disks>} INIT_ALLOW=<initialization policy> SATASSURE=[NONE|DATA_INTEGRITY_FIELD|PARITY]

where <initialization policy>=Allow_IO, NO_IO, Priority (1-99, where 99 puts all system resources on the initialization).

You may either explicitly select the disks for the pool by using "**PHYSICAL_DISK=**" or specify the number of disks in the pool by using "**NUMBER=**" in which case the next available disks will be selected. If you specify the **ASSIGN_POLICY** and **NUMBER** parameters together, only the selected type of disks will be used for the Storage Pool and you will not need to enter the individual disk name. You may disable the SATAssure feature or enable it to use either the data integrity field option or the parity check on read option.

Examples:

- To use only the specified disks:
  CREATE POOL RAID_LEVEL=RAID5 CHUNK_SIZE=64KB PHYSICAL_DISK=62 63 64 65 66
- To use the next available disks:
  CREATE POOL RAID_LEVEL=RAID6 CHUNK_SIZE=128KB NUMBER=6
- To use only the selected type of disks:
  CREATE POOL RAID_LEVEL=RAID6 CHUNK_SIZE=128KB NUMBER=6 ASSIGN_POLICY=SAS

3.2.2.3 Storage Pool Initialization

When a Storage Pool is created, initialization begins automatically and will continue in the background until it is completed.

You may use the **SHOW POOL** command to check if the Storage Pool is being initialized (see Section 3.2.2.1, "Display Storage Pool Information" for more information).

You may monitor the initialization progress using the **SHOW JOB** command (see Section 3.3.1, "Background Job Priority" for more information).

3.2.2.4 Initialization Job Failure

Rarely, a job may fail to initialize. If so, the event will be reported as:

```
000737 2010-01-25 20:42:09:0684480  G=61 S=2 T=1 RP=0 VP=1
LOG_ST_SET_FAILED  STATE VIRTUAL DISK SET FAILED CONDITION INDEX:00000003 POOL:00000000 INIT    FAILED
```

If an initialization job fails, enter the command **VERIFY POOL <pool-id> FORCE** to restart the initialization.

Examples:

- **SHOW POOL *** command under failed initialization conditions:

```
                                          |Total |Free  |Max   |                  |Disk|  Global  |Spare |
Idx|Name           |State      |Chunk|Raid| Faults |cap GB|cap GB|VD GB | Settings  | Jobs  |T/O |spare pool|Policy|
-------------------------------------------------------------------------------------------------------------
  0 pool-0         FAILED       128   6        I      720      0      0 DWMRFI              1  UNASSIGNED  AUTO
  1 pool-1         NORMAL       128   6               720      0      0 DWMRFI              1  UNASSIGNED  AUTO
```

- **SHOW POOL 0 ALL** command under failed initialization conditions:

```
Index:               0
OID:                 0x1a520000
Type:                STORAGE
Name:                pool-0
Chunk Size:          128KB (0x100 blocks)
Block Size:          512
RAID Type:           RAID6
Free Raid6 Capacity: 0 GB
Max VD Size:         0 GB
Total Capacity:      720 GB
UUID:                60001ff0800a3000000000001a520000
Global Spare Pool:   UNASSIGNED
DiskTimeout(FRT):    1 minute
Init Policy:         ALLOW_IO
 .
 .
 .
Initializing:        FAILED
Rebuilding:          FALSE
Paused:              FALSE
 .
 .
 .
```

- **SHOW VD *** command under failed initialization conditions:

```
                                                                |      Home       |Background|
Idx|Name           | State |Pool|Raid|Cap GB|Settings| Jobs  |Current|Preferred|   Job    |
-------------------------------------------------------------------------------------------
  0 vd-0_0          FAILED    0  5        8   W M I          1(L) 0   1(L) 0   INACTIVE
  1 vd-1_0          FAILED    0  5        8   W M I          1(L) 0   1(L) 0   INACTIVE
  2 vd-2_0          FAILED    0  5        8   W M I          1(L) 0   1(L) 0   INACTIVE
  3 vd-3_1          READY     1  5        8   W M I          1(L) 0   1(L) 0   INACTIVE
```

- **SHOW VD 3 ALL** command under failed initialization conditions:

```
Index:           0
OID:             0x8a680000
Name:            vd-0_0
Pool Index:      0
Pool OID:        0x1a520000
Capacity:        472 GB
Offset:          0x0
State:           FAILED INITIALIZATION
Raidlevel:       RAID6
IO ROUTING:      TRUE
WBC:             TRUE
MIRRORED:        TRUE
Initializing:    FAILED
Paused:          FALSE
AutoWriteLock:   FALSE
```

3.2.2.5 Verifying a Storage Pool

Pool verification is a feature that goes with SATAssure. It allows you to run a background job that finds and fixes parity issues (when using the **FORCE** option). It is also automatically initiated when SATAssure encounters an error in an attempt to fix the errors.

To set this mode, use the command:
SET SUBSYSTEM VERIFY_POLICY=<0..99>
where the number range indicates the internal resources to use for the process. Zero is used to turn the feature OFF, any other number will turn the feature on.

To use the "one-time" verify function, use the command:
VERIFY POOL [<pool-id>|*] FORCE_CONSISTENCY

3.2.2.6 Naming a Storage Pool

The **SET POOL=<pool-id> NAME="<name>"** command lets you specify a name to identify the Storage Pool (Figure 91).

Figure 91. Set Pool Name Example Screen

```
RAID[0]$ set pool 0 name="RAID 5 Set"
POOL 0 OID: 0x19b40000 attributes set with STATUS='Success' (0x0)
```

3.2.2.7 Deleting a Storage Pool

The **DELETE POOL=<id>** command deletes the specified Storage Pool from the system.

The DELETE POOL command erases all the data on the Storage Pool. You cannot delete a Storage Pool if there are Virtual Disks still present on the pool.

3.2.2.8 Locate a Storage Pool

The **LOCATE POOL=<id>** command provides a visual indication of the specified Storage Pool. The status LED of the disk modules in the specified Storage Pool will blink.

3.2.3 Virtual Disk Management

A Virtual Disk is the storage unit presented to any attached host. A Virtual Disk can be created to use all or just a part of the capacity of a single Storage Pool. Virtual disks allocate space in 8 GiB increments. For example, 16 GiB of storage space will be allocated when creating a Virtual Disk of 10 GiB.

3.2.3.1 Display Virtual Disk Information

The **SHOW VIRTUAL_DISK *** command displays the list of configured Virtual Disks in the system (Figure 92).

Figure 92. Show Configured Virtual Disks Example Screen

```
RAID[0]$ show vd *

                                                                                  |          Home    |Background|
Idx|Name            | State |Pool|Raid|Cap GB|Settings| Jobs  |Current|Preferred|    Job     |
-------------------------------------------------------------------------------------------------------
  0 vd-0_0           READY     0  5     3632   W   I           0(L) 0   0(L) 0   INACTIVE
  1 vd-1_1           READY     1  5     3632   W   I           0(L) 0   1(R) 0   INACTIVE
  2 vd-2_2           READY     2  5     3632   W   I           0(L) 0   0(L) 0   INACTIVE
  3 vd-3_3           READY     3  5     5480   W   I           0(L) 0   1(R) 0   INACTIVE
  4 vd-4_4           READY     4  5     5480   W   I           0(L) 0   0(L) 0   INACTIVE
  6 vd-6_5           READY     5  5     5480   W   I           0(L) 0   1(R) 0   INACTIVE

Total Virtual Disks: 6
```

The **SHOW VIRTUAL_DISK=<id> ALL_ATTRIBUTES** command displays the detailed information of the specified Virtual Disk (Figure 93).

Figure 93. Show Virtual Disk Attributes Example Screen

```
RAID[0]$ show vd 0 all

  Index:              0
  OID:                0x8a680000
  Name:               vd-0_0
  Pool Index:         0
  Pool OID:           0x1a520000
  Capacity:           472 GB
  Offset:             0x0
  State:              READY
  Raidlevel:          RAID6
  IO ROUTING:         TRUE
  WBC:                TRUE
  MIRRORED:           TRUE
  Initializing:       FALSE
  Paused:             FALSE
  AutoWriteLock:      FALSE
  Critical:           FALSE
  Forced Write-thru:FALSE
  Current Home:       0x38000000  0x00000000 (LOCAL)
  Future Home:        0xffffffff  0x00000000
  Preferred Home:     0x38000000  0x00000000 (LOCAL)
      Job OID:        INACTIVE
  UUID:               60001ff0800a3000000000008a680000

 Total Virtual Disks: 1
```

3.2.3.2 Creating a Virtual Disk

To create a Virtual Disk, enter the command:
`CREATE VIRTUAL_DISK CAPACITY=<capacity> POOL=<pool-id>`
where <capacity> is the capacity of the Virtual Disk in GiB or type “max” to use all available capacity; <pool-id> is the Storage Pool to be used by this Virtual Disk (Figure 94).

A message is displayed to indicate whether the Virtual Disk creation was successful.

Figure 94. Create Virtual Disks Example Screen

```
RAID[0]$ create virtual_disk capacity=16 pool=0
VIRTUAL_DISK 0 OID=0x89ba000 creation STATUS='Success' (0x0)

RAID[0]$ create virtual_disk capacity=32 pool=0
VIRTUAL_DISK 1 OID=0x89bb001 creation STATUS='Success' (0x0)
```

3.2.3.3 Naming a Virtual Disk

The **`SET VIRTUAL_DISK=<id> NAME=<name>`** command lets you specify a name to identify the Virtual Disk (Figure 95).

Figure 95. Set Virtual Disk Name Example Screen

```
RAID[0]$ set vd 0 name=LUN ”
VIRTUAL_DISK 0 OID=0x89ba000 attributes set with STATUS=’Success’ (0x0)
```

3.2.3.4 Deleting a Virtual Disk

The **`DELETE VIRTUAL_DISK=<id>`** command deletes the specified Virtual Disk from the system. However, it does not delete presentations for that Virtual Disks. If a presentation is configured for a Virtual Disk, deleting the Virtual Disk will return an error. You must first delete all the presentations configured on that Virtual Disk (see Section 3.4.3, "Presentation Commands").

Warning

When you delete a Virtual Disk, you lose access to all data that was stored on that Virtual Disk.

3.2.3.5 Presentation to Hosts

Virtual Disks are only presented to the hosts that have been given authorized access. Refer to Section 2.8.12, "Present Virtual Disk to External Host" for detailed instructions on configuring Virtual Disk presentations to hosts. Refer to Section 3.4, "Presentations" for more information on commands related to presentation.

To display presentations associated with the specified application host, enter the command:
`APPLICATION SHOW HOST=<host-id> PRESENTATIONS`

To display the initiators associated with the specified application host, enter the command:
`APPLICATION SHOW HOST=<host-id> INITIATORS`

3.2.4 Spare Pool Management

The IS16000 supports the concept of Spare Pool which contains Physical Disks that can be used as spare disks. Spare Pools can be shared by Storage Pools or dedicated to a particular Storage Pool in which case the spare disks will only be used by the designated Storage Pool. After a Spare Pool is created, you will need to select disk(s) from the Unassigned Pool and assign it the Spare Pool.

In the event of disk failure, a disk that is assigned to a Spare Pool is automatically swapped into the Storage Pool to replace the failed disk. A rebuild then occurs immediately, reducing the possibility of data loss due to additional disk failures.

NOTE : Each Storage Pool should have a Spare Pool assigned to it. If a Storage Pool's attributes show "Global Spare Pool: UNASSIGNED", no Spare Pool will serve this Storage Pool.

3.2.4.1 Display Spare Pool Information

Use the **SHOW SPARE_POOL *** command to display the list of configured Spare Pool(s) (Figure 96).

Figure 96. Show Configured Spare Pools Example Screen

```
RAID[0]$ show spare_pool *

                                |Disk |Total |Total|Storage |
Idx|Name                  |Blocks| T/O |Cap GB| PDs |Pool Idx|
-------------------------------------------------------------
  6 spare_pool-6             512    10     528    2
  7 spare_pool-7             512    10     528    2

Total Spare Pools: 2
```

To display the detailed information of the Spare Pool, enter the command:
SHOW SPARE_POOL=<id> ALL_ATTRIBUTES command (Figure 97).

Figure 97. Show Spare Pool Attributes Example Screen

```
RAID[0]$ show spare_pool 6 all

  Index:            6
  OID:              0x19bc0006
  Type:             GLOBAL SPARE
  Name:             19bc0006
  Block Size:       0x200
  DiskTimeout(FRT): 10 minutes
  Total Capacity:   1409024 MBs
  UUID:             0x00
  Total Phy Disks   2
```

To display the list of disks in the Spare Pool, enter the command (Figure 98):
SHOW SPARE_POOL=<id> PHYSICAL_DISKS

Figure 98. Show Spare Pool Physical Disks Example Screen

```
RAID[0]$ show spare_pool 6 pd
                                                                                                |Health|
Encl|Slot| Vendor |     Product ID     |Type|Cap GB | RPM|Revision|    Serial Number   |Pool|State | Idx|State |       WWN       |
-------------------------------------------------------------------------------------------------------------------------------------
   2   53 HITACHI  HUS153030VLS300       SAS    279  15K A410     JHVYKUKC5303          14  SPARE  648 NORM    5000cca00d35c6b0
   2   58 HITACHI  HUS153030VLS300       SAS    279  15K A410     JHVWU54C5303          14  SPARE  652 NORM    5000cca00d3291f4

Total Physical Disks:      2
Total Assigned Disks:      2
Total Unassigned Disks:    0
  Total SAS Disks:         2
  Total Member State Normal:   2
```

3.2.4.2 Creating a Spare Pool

Use these commands to create a Spare Pool, add disk to the Spare Pool, and assign the Spare Pool to a Storage Pool:

- **CREATE SPARE_POOL**
- **ASSIGN PHYSICAL_DISK=<disk-id> TO_POOL <spare-pool-id>**
- **SET POOL=<pool-id> SPARE_POOL=<spare-pool-id>**
  where <pool-id> and <spare-pool-id> are the index names of the Storage Pool and Spare Pool respectively.

3.2.4.3 Naming a Spare Pool

The **SET SPARE_POOL=<pool-id> NAME="<name>"** command lets you specify a name to identify the Spare Pool.

3.2.4.4 Deleting a Spare Pool

To delete a disk from the Spare Pool, enter the command:

```
ASSIGN PHYSICAL_DISK=<disk-id> TO_POOL 0x1800ffff
```

To delete a Spare Pool, enter the command:

```
DELETE SPARE_POOL=<pool-id>
```

3.2.5 Disk Rebuild

The IS16000's automatic disk failure recovery procedures ensure that absolute data integrity is maintained while operating in degraded mode.

3.2.5.1 Full and Partial Rebuilds

Both full and partial (fractional) rebuilds are supported in IS16000. Rebuilds are done at the Storage Pool level, not the Virtual Disk level. Partial rebuilds will reduce the time to return the Storage Pool to full redundancy and can be controlled by adjustable time limit for each Storage Pool. Permitted time limits are 0 (in which case partial rebuild is off) to 240 minutes. The timer will begin when a disk in the Storage Pool is declared missing. If the disk reappears prior to the expiration of the timer, a fractional rebuild will be done. Otherwise, the disk will be declared failed, replaced by a spare and a full rebuild will begin to return the Storage Pool to full redundancy. The default partial rebuild timer (Disk Timeout) setting is 10 minutes.

Under heavy write workloads, it is possible that the number of stripes that need to be rebuilt will exceed the system's internal limits prior to the timer expiration. When this happens, a full rebuild will be started automatically instead of waiting for the partial rebuild timeout.

Use the **SHOW POOL=<id> ALL_ATTRIBUTES** command to display the current Disk Timeout setting (Figure 99).

Figure 99. Show Pool All Attributes

```
RAID[0]$ show pool 0 all
  Index:                0
  OID:                  0x1a520000
  Type:                 STORAGE
  Name:                 pool-0
  Chunk Size:           128KB (0x100 blocks)
  Block Size:           512
  RAID Type:            RAID6
  Free Raid6 Capacity:  0 GB
  Max VD Size:          0 GB
  Total Capacity:       720 GB
  UUID:                 60001ff0800a3000000000001a520000
  Global Spare Pool:    0x1aca000e  (Index 14)
  DiskTimeout(FRT):     1 minute
    .
    .
```

To change the disk timeout setting, use this command:

SET POOL=<id> DISK_TIMEOUT=<timeout>

where <timeout> is in the range of <0..240> minutes. The default setting is 10.

3.2.5.2 Sparing Policy

Each Storage Pool has a sparing policy that determines what happens when a physical disk within the pool fails (or becomes inaccessible). In the event of a disk failure, the IS16000 will automatically initiate a disk rebuild if the sparing policy is set to automatic and a Spare Pool has been assigned to the Storage Pool.

Use the **SHOW POOL=<id> ALL_ATTRIBUTES** command to display the current settings (Figure 100).

Figure 100. Show Pool All Attributes

```
RAID[0]$ show pool 0 all

  Index:                  0
  OID:                    0x1a520000
  Type:                   STORAGE
    .
    .
    .
  UUID:                   0x00
  Global Spare Pool:      0x1aca000e (Index 14)
  DiskTimeout(FRT):       10 minutes
  Init Policy:            ALLOW_IO
  Init Priority:          50%
  Full Rebuild Priority:        80%
  Fractional Rebuild Priority:  90%
  Sparing Policy:         AUTOMATIC
    .
    .
    .
```

To change the sparing policy setting, use this command:

SET POOL=<id> SPARING_POLICY=[AUTOMATIC|MANUAL]

The default setting is automatic which is recommended.

A rebuild operation can take up to several hours to complete, depending on the size of the disk and rate of rebuild. Refer to Section 3.3, "Performance Management" on page 76 for information on how to adjust the rate of rebuild.

3.2.5.3 Manual Disk Replace/Rebuild

If a Storage Pool does not have a Spare Pool assigned to it, it becomes necessary to manually add a disk to the Storage Pool to replace a failed disk.

To add a disk to a Storage Pool to replace a failed disk, use this command:

ASSIGN PHYSICAL_DISK=<disk-id> TO_POOL=<pool-id> SET_SPARE

The disk will be put into the Storage Pool as the spare disk and it will be used in the Storage Pool to replace the failed disk.

3.2.5.4 Manual Fail/Rebuild of a Disk

The **SET PHYSICAL_DISK <disk-id> FAILED** command instructs the system to fail the specified disk. When a non-SPARE disk is specified and it is failing, the disk will not cause a multi-channel failure. The disk is marked as failed. An attempt is made to replace it with a spare disk. When a SPARE disk is specified, it is released, but marked as unhealthy and unavailable.

CLEAR PHYSICAL_DISK <disk-id> FAILED changes the specified disk's "Failed" state to "FALSE". After clearing a disk, use the **ASSIGN PHYSICAL_DISK <disk-id> TO_POOL <pool-id> SET_SPARE** command to add the disk back to the Storage Pool it was failed from. This command will also initiate a rebuild if a spare has not already been assigned to the Storage Pool.

3.2.6 SATAssure

You may enable SATAssure for a Storage Pool.

To display the current setting, enter the command (Figure 101):
SHOW POOL=<pool-id> ALL_ATTRIBUTES

Figure 101. Show Storage Pool Details Example Screen

```
RAID[0]$ show pool 0 all
  Index:                0
  OID:                  0x1a520000
  Type:                 STORAGE
  Name:                 pool-0
   .
   .
  Assign Policy:
    Device Type:        SAS
    Rotation Speed:     NA
    Raw Capacity:       136 GB
  SATAssure:            NONE
    .
    .
```

To enable SATAssure, enter the command:
SET POOL=<pool-id> SATASSURE=[PARITY|DATA_INTEGRITY_FIELD]
where **PARITY** is the parity check on read option and **DATA_INTEGRITY_FIELD** is the integrity field option.

To disable SATAssure, enter the command: **SET POOL=<pool-id> SATASSURE=NONE**

3.2.7 Setting the System's Date and Time (NTP Mode)

NTP (Network Time Protocol) mode is available on the IS16000. It provides a means for the Controllers to synchronize their time across a network, usually within a small number of milliseconds over a long period of time. You can enter up to four NTP addresses as the time servers.

To enable the NTP mode, enter the command:
SET SUBSYSTEM NTP=[<ip address list, up to 4>|NONE]

Examples:
SET SUBSYSTEM NTP=1.2.3.4 1.2.3.5 3.4.5.6 6.5.7.8
SET SUBSYSTEM NTP=1.2.3.4
SET SUBSYSTEM NTP=NONE

Specifying a list of NTP addresses will start the NTP mode on each Controller, using that list of NTP addresses as the time servers. Specifying NONE will turn off NTP mode.

To display the current settings, enter the command(Figure 102):
SHOW SUBSYSTEM ALL_ATTRIBUTES

Figure 102. Show NTP Settings

```
RAID[0]$ show subsystem all
********************
*    Subsystem     *
********************

RP Subsystem Name:      co-test-8
UID:                    60001ff0800a30000000000030000000
Subsystem Time:         Mon Sep 27 17:08:59 2010
Locate Dwell Time:      120 seconds
Enabled Licenses:       RAID6 SATASSURE
Fast Timeout:           ON
Pool Verify Priority:   10%
NTP Mode:               ON
                        (10.32.16.24)
                        (10.32.16.25)
```

When in NTP mode:

- The time is set between the two Controllers under the following conditions:
  - when the Controllers boot and discover each other
  - when a **SET SUBSYSTEM** command is issued, whether it is setting the time or not
- Each Controller will attempt to synchronize with the specified NTP servers. If none of the servers are valid, the time on each Controller is free to drift independent of any other time source (and independent of the other Controller).
- The **SET SUBSYSTEM DATA_AND_TIME** command will fail with a status of "Setting date/time while in NTP mode".
- Once the Clock code has finished calibration (100 seconds), NTP will be started. Both Controllers will have the same NTP settings, and so presumably will have synchronized time once NTP sets the time.
- Issuing a new **SET SUBSYSTEM NTP** command with a new set of IP addresses will stop and restart NTP.
- If NTP finds a time difference of more than 128 msec, it will "jump" the time to the correct time. This will result in a discontinuity in the event log, logdisk, syslog, and anywhere that records a timestamp.
- NTP will always set the Controllers to UTC (Coordinated Universal Time). There is no option to set time zones or otherwise change the offset from UTC.

When not in NTP mode:

- The master Controller uses its time to set the time on the other Controller.
- The **SET SUBSYSTEM DATE_AND_TIME** command is used to set the time of both Controllers together. The system records time using the military method, which records hours from 00 to 24, not in a.m. and p.m. increments of 1 to 12. Valid date settings are between years 2000 and 2104. Settings are automatically adjusted for leap years.

  To change the system date and time to March 1, 2010 2:15:32 pm, for example, type:

  ```
  SET SUBSYSTEM DATE_AND_TIME=2010:3:1:14:15:32
  ```

- Once the time is set, the time on the two Controllers is free to drift independent of the other Controller.

3.2.8 Restarting the IS16000

System Restart

The **SHUTDOWN CONTROLLER [LOCAL|REMOTE|0|1] RESTART** command performs a restart on the specified Controller.

The **SHUTDOWN SUBSYSTEM RESTART** command performs a restart on both Controllers.

These commands will prepare the system to be restarted. The system will halt all I/O requests and save the data to the disks before restarting. The restart process may take several minutes to complete.

System Shutdown

The **SHUTDOWN CONTROLLER [LOCAL|REMOTE|0|1]** command shuts down the specified Controller.

The **SHUTDOWN SUBSYSTEM** command shuts down both Controllers.

If you need to power down the IS16000, use **SHUTDOWN** prior to shutting off the power. This will cause the IS16000 to immediately flush its cache, abort all initialization and rebuild operations, and proceed with an orderly shutdown.

All hosts actively using the IS16000 should be safely shutdown and all users logged out before using this command. The IS16000 will halt all I/O requests and save the data to the disks. The unit can be safely turned off after using this command.

Once shut down is complete, all power supplies must be switched off or unplugged. Power must be removed from the system for at least 10 seconds before it will start up again.

NOTE : Use **SHUTDOWN** whenever you need to power down the IS16000 for maintenance. **SHUTDOWN** flushes any data left in the cache and prepares the IS16000 for an orderly shutdown.

3.3 Performance Management

Initialization and rebuild operations are background processes and their rates can be adjusted to minimize their impact on system performance.

3.3.1 Background Job Priority

To monitor all current jobs (Figure 103), enter command:

SHOW JOB * ALL_ATTRIBUTES

The type of job and percentage of completion are displayed.

Figure 103. Show Job Example Screen

```
RAID[0]$ show job * all
                                                    |Fraction|
 Idx|Type           |Target   (Sub)    |State      |Complete|Priority|Status  |Time
-------------------------------------------------------------------------------------
   0 FULL REBUILD    POOL:0   (NA)      RUNNING         0%      80%             NA

Total Background Jobs: 1
```

You may specify the amount of system resources that should be devoted to a background job. The higher its priority value, the faster the background job will run and the more the background job will impact client I/O performance.

To set the job priority, enter the command:

SET JOB=<id> PRIORITY=<priority>

where <priority> is a number between 1 and 99. Note that PRIORITY is not a percentage or a mathematical fraction of the available resources. For example, two background jobs with priority values of 50 will not use 100% of the system resources. The IS16000 may or may not limit the number of background jobs to keep the total of their priorities below 100.

Pause/Resume a Job

You may pause a job at any time using the **PAUSE JOB=<id>** command (Figure 104).

Figure 104. Pause a Job

```
RAID[0]$ pause job 0
JOB 0 OID=0x2b010000 paused with STATUS='Janus Success' (0x0)

RAID[0]$ show job 0 all
  OID:                0x2b010000 (Index:0)
  Target:             POOL:0x1a520000 (Index:0)
  Sub-Target:         NA
  Type:               REBUILD
  Status:             PAUSED
  Completion Status:  UNKNOWN
  Priority:           50
  Fraction Complete:  11%
  Time:               NA
```

To resume the job, enter the command: **RESUME JOB=<id>** (Figure 105)

Figure 105. Resume a Job

```
RAID[0]$ resume job 0
JOB 0 OID=0x2b010000 resumed with STATUS='Janus Success' (0x0)

RAID[0]$ show job 0 all
  OID:                0x2b010000 (Index:0)
  Target:             POOL:0x1a520000 (Index:0)
  Sub-Target:         NA
  Type:               REBUILD
  Status:             RUNNING
  Completion Status:  UNKNOWN
  Priority:           50
  Fraction Complete:  11%
  Time:               NA
```

3.3.2 Right Side I/O

Access to Virtual Disks uses the notion of a preferred home or “homed” path. If a path to the VD utilizes the peer Controller, additional latency is incurred for every I/O as it must be processed first by the “non-homed” Controller and forwarded over the same communication path used by Mirrored Write Back Cache (Figure 106).

Figure 106. Path to VD Utilizes the Peer Controller

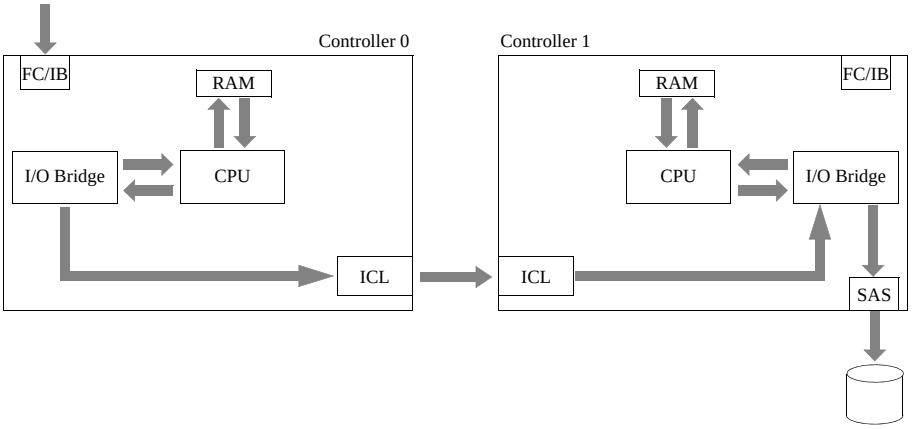

Properly configured and tested host side multi-path drivers will ensure right side I/O will occur. If poor performance is encountered, especially upon initial configuration, then verifying and correcting your primary path to disks is essential.

Verify by reviewing the Virtual Disk counters, enter the command (Figure 107):

```
SHOW VD * COUNTERS ALL
```

NOTE : You need to enter this command three times—first time to initiate the counters, second time to display the results, third time to reset the counters.

Controller 0’s results are displayed on the left and Controller 1’s results are displayed on the right. In Figure 107, if VD 102 is mastered by Controller 0, all of the I/O is passing through Controller 1, thus doing wrong side I/O.

Figure 107. Show Virtual Disk Counters Example Screen

```
RAID[0]$ show vd * counters

Virtual disk Counters: Elapsed time = 12.181 seconds

  Idx  IOs/sec   KiB/sec   KiB/IO Fwd IO/s Fwd KiB/s| IOs/sec   KiB/sec   KiB/IO Fwd IO/s Fwd KiB/s|
-------------------------------------------------------------------------------------------------------
  102        0         0        0         0         0 |    6580    109243       68         0         0 |
  103        0         0        0         0         0 |    7005    116304       68         0         0 |
  104        0         0        0         0         0 |    5158    186392      148         0         0 |
  105        0         0        0         0         0 |    2825    102083      148         0         0 |
  106        0         0        0         0         0 |    3481    210768      248         0         0 |
  107        0         0        0         0         0 |    3641    220500      248         0         0 |
  108        0         0        0         0         0 |       0         0        0         0         0 |
  109        0         0        0         0         0 |       0         0        0         0         0 |
  110        0         0        0         0         0 |       0         0        0         0         0 |
  111        0         0        0         0         0 |       0         0        0         0         0 |
  112        0         0        0         0         0 |       0         0        0         0         0 |
  113        0         0        0         0         0 |       0         0        0         0         0 |
  114        0         0        0         0         0 |       0         0        0         0         0 |
  115        0         0        0         0         0 |       0         0        0         0         0 |
  116        0         0        0         0         0 |       0         0        0         0         0 |
  117        0         0        0         0         0 |       0         0        0         0         0 |
  118        0         0        0         0         0 |       0         0        0         0         0 |
  119        0         0        0         0         0 |       0         0        0         0         0 |
```

3.3.3 Rebuild Policy Priority

Each Storage Pool has its own rebuild policy settings. You may specify the priority values for full rebuild and partial (fractional) rebuild policies. To display the current settings, enter the command **SHOW POOL=<id> ALL_ATTRIBUTES** (Figure 108).

Figure 108. Show Pool Information

```
RAID[0]$ show pool 3 all
  Index:                0
  OID:                  0x1a520000
  Type:                 STORAGE
  Name:                 pool-0
  Chunk Size:           128KB (0x100 blocks)
  Block Size:           512
  RAID Type:            RAID6
  Free Raid6 Capacity:  0 GB
  Total Capacity:       720 GB
  UUID:                 60001ff0800a3000000000001a520000
  Global Spare Pool:    0x1aca000e (Index 14)
  DiskTimeout(FRT):     10 minutes
  Init Policy:          ALLOW_IO
  Init Priority:        50%
  Full Rebuild Priority:        80%
  Fractional Rebuild Priority:  90%
  Sparing Policy:       AUTOMATIC
    .
    .
    .
    .
```

Full Rebuild Priority

To change the full rebuild priority value (Figure 109), enter the command:

SET POOL=<id> REBUILD_FULL_POLICY=<priority>

where <priority> is an integer in the range of 1 to 99. The default value specifies a priority of 80%.

Figure 109. Set Full Rebuild Priority

```
RAID[0]$ set pool 3 rebuild_full_policy=70

POOL 3 OID=0x19b40003 attributes set with STATUS='Success' (0x0)
```

Partial Rebuild Priority

To change the partial rebuild priority value (Figure 110), enter the command:

SET POOL=<id> REBUILD_PARTIAL_POLICY=<priority>

where <priority> is an integer in the range of 1 to 99. The default value specifies a priority of 90%.

Figure 110. Set Partial Rebuild Priority

```
RAID[0]$ set pool 3 rebuild_partial_policy=80

POOL 3 OID=0x19b40003 attributes set with STATUS='Success' (0x0)
```

3.3.4 Cache Coherency

By default, the IS16000 runs in couplet mode, where both Controllers are running simultaneously, communicating through internal Ethernet and InfiniBand connections. This means that the IS16000 is always running with cache coherency enabled.

The cache settings are configurable for each Storage Pool. Available keywords are:

FULL_STRIPE_WRITE_CACHING	Specify full stripe write caching
MIRRORED	Specify cache data mirroring
READ_AHEAD_CACHING	Specify read ahead caching
WRITE_BACK_CACHING	Specify write back caching

To display the current settings, enter the command **SHOW POOL=<id> ALL_ATTRIBUTES**.

To change the settings, use these commands:

```
SET POOL=<id> FULL_STRIPE_WRITE_CACHING=[TRUE|FALSE]
SET POOL=<id> MIRRORED=[TRUE|FALSE]
SET POOL=<id> READ_AHEAD_CACHING=[TRUE|FALSE]
SET POOL=<id> WRITE_BACK_CACHING=[TRUE|FALSE]
```

NOTE : Turning off mirroring may result in data integrity issues.

3.3.4.1 Cache Protection

The IS16000 UPS serves as a battery, which allows the controller to flush the contents of cache to non-volatile storage (internal disk) during a power event. When cache has been completely flushed, the Controller will shut itself down. When power is restored, the flushed cache is replayed back from the non-volatile storage and committed to disk.

Mirrored Write Back Cache (MWBC) provides a mechanism of cache protection by copying, or mirroring, the contents of cache from one singlet to the peer singlet (and vice-versa). In the event of a singlet failure, the copy of cache sitting in the surviving peer can be committed to disk by the surviving peer. Although there is a latency involved in the cache mirroring process, it is recommended as additional protection of data.

3.4 Presentations

A Presentation on a IS16000 has the following components:

- Discovered Initiators
  A Discovered Initiator is defined as a Fibre Channel or InfiniBand initiator that has logged into the IS16000. This is normally a single port on a Fibre Channel or InfiniBand HCA in an external computer.

- Host
  A Host is an abstract client of a Virtual Disk that has all the access to all data within the Virtual Disk. This most commonly represents an external computer with one or more Fibre Channel or InfiniBand initiators. Host attributes are OS (operating system) of which you can select GENERIC, WINDOWS, LINUX or MAC_OS. The OS attribute dictates the way the controller responds to certain Fibre Channel commands where appropriate. For example, Mac OS X expects a different response from an inquiry of LUN0 (the controller LUN) than a Windows host.
  - OSTYPE: “Standard” uses current settings and “Custom” allows for requests for different adjustments from OEMs.

- Channel
  A Channel is one of the Fibre Channel or InfiniBand ports on the Controller. Channel attributes are MODE of which you can select either MAC_OS or STANDARD. When MAC_OS is selected, the port’s node name will be set differently in order to be visible to a Macintosh system.

NOTE : If your SAN configuration requires unique world-wide name (WWN) node names on each controller port, set the controller channles (ports) to MAC mode. Use the command **APPLICATION SET CHANNEL [0|1|2|3] MODE MAC.** See section 3.4.5, "Additional Configuration Considerations for Macintosh Hosts"and other Mac OS notes for details.

- Stack
  A Stack is the input side of the subsystem. In the case of the IS16000, the stack is defined as the Fibre Channel interface. In future versions of the product, different interface stacks may exist.

- Presentation
  A Presentation is the relationship between a Host and a Virtual Disk. A Presentation implies that the related Host has some sort of access to the Virtual Disk. Attributes of a Presentation are:
  - PORT: from which the host will see the Virtual Disk
  - READ_ONLY: controls read only access
  - PRESENT_HOME_ONLY: presents the specified Virtual Disk from its designated home controller only
  - LUN: user-specified LUN number that the Virtual Disk will show to the host.

3.4.1 Discovered Initiator Commands

To map a host to a discovered initiator, use the command:
APPLICATION IMPORT DISCOVERED_INITIATOR=<initiator-id> HOST=<host-id>

To display the currently available initiators, use the command:
APPLICATION SHOW DISCOVERED_INITIATORS *

To display the attributes of a specified initiator, use the command:
APPLICATION SHOW DISCOVERED_INITIATORS=<initiator-id> [ALL ATTRIBUTES]

To create an application initiator for the specified host, use the command:
APPLICATION CREATE INITIATOR HOST=<host-id> WWPN=<integer>

To delete the specified initiator, use the command:
APPLICATION DELETE INITIATOR=<initiator-id>

3.4.2 Host Commands

To create a host, use the command:
APPLICATION CREATE HOST [OSTYPE=GENERIC|LINUX|MAC_OS|WINDOWS]

To delete the specified host, use the command:
APPLICATION DELETE HOST=<host-id>

To display the presentation associated with the specified application host, use the command:
APPLICATION SHOW HOST=<host-id> [PRESENTATIONS]

To display the initiators associated with the specified application host, use the command:
APPLICATION SHOW HOST=<host-id> [INITIATORS]

3.4.3 Presentation Commands

To present a Virtual Disk to the specified host, use the command:
APPLICATION CREATE PRESENTATION HOST=<host-id> VIRTUAL_DISK=<vd-id> LUN=<LUN-id>

To delete a presentation, use the command:
APPLICATION DELETE PRESENTATION HOST=<host-id> VIRTUAL_DISK=<vd-id> [FORCE]
Note the optional parameter of **FORCE** deletes without confirmation. The wildcard ***** deletes all presentations.

To set the specified attribute to the specified value, use the command:
APPLICATION SET PRESENTATION=<object-id> [<attribute-name>=<value>]
Attributes are:

ENABLE=ALL|NONE|<channel-id> where the channel-id is the object -id of an Enabled Client Channel.

LUN=<integer> where the integer is a Logical Unit Number (LUN) that will be used to present the associated Virtual Disk to the associated Host. NOTE: Each LUN integer entered for a presentation is cumulative and does not replace the previous entry.

HOME_ONLY=[TRUE|FALSE] which enables/disables the home_only parameter.

READ_ONLY=[TRUE|FALSE] which enables/disables the read_only parameter.

3.4.4 Configuration of Presentations of Virtual Disks to Hosts

---

NOTE : If you are configuring a storage subsystem with any “MAC OSx presentations”, it is absolutely required that you read and understand Section 3.4.5, "Additional Configuration Considerations for Macintosh Hosts" prior to actually doing your configuration.

There are three objects/relationships that must be properly established in order to create a presentation of a LUN (Virtual Disk) to a host:

- **Establish a host object** for EACH host that wishes access to the IS16000 storage system Virtual Disk.

NOTE : On the storage subsystem you ONLY need to create a single host object for any/all Virtual Disk presentations to that host.

- **Import a discovered initiator** into an established relationship with a host object. An association between a discovered initiator FC Port/Node WWN (World Wide Name) to a host object is established in this operation. This association of host/port WWN to host object is maintained persistently within the subsystems configuration information along with all other information. If the host to controller association is subsequently disrupted and then re-established, the Controller is able to maintain this relationship until such time that the configuration in the storage Controller is cleared or that relationship is deleted.

TIP : Use of a host based HBA utility such as HBAnyware® or SANsurfer will allow you to examine the FC Port/Node WWNs on the host in the easiest manner.

- **Assign a Virtual Disk to a host object.**

NOTE : You may ONLY present a Virtual Disk ONCE to the same host object.

Warning

You can present a Virtual Disk to multiple hosts; however, this is dangerous. If doing so, presenting them as Read-Only to the other hosts would be appropriate.

NOTE : You may present a Virtual Disk to ALL hosts. This may be appropriate in some limited system environments; however, it is best practice to not MIX the presentations, where some VDs are selectively presented while others are presented all. Managing at the host end may become confusing.

To present a Virtual Disk to ALL hosts, use the command:
APPLICATION CREATE PRESENTATION VD=<VD-id> HOST=ALL

3.4.4.1 Host Object Creation Example

The following example demonstrates presenting 6 Virtual Disks to 3 separate hosts (2 Windows hosts and a Linux host) from a single storage subsystem. Although this configuration will have a FC Switch, the steps are identical.

NOTE : In the example below, the CLUI is operating in the RAID subject mode. The ASM (Application Stack Management) commands must be preceded by the subject application.

Examples:

```
RAID[1]$ application create host name=co-ls1 ostype=Linux
HOST 301 OID=0x1d5e012d creation STATUS='Success' (0x0)

RAID[1]$ application create host name=co-test-d10 ostype=windows
HOST 302 OID=0x1d5f012e creation STATUS='Success' (0x0)

RAID[1]$ app create host name=co-test-d08 ostype=windows
HOST 303 OID=0x1d60012f creation STATUS='Success' (0x0)

RAID[1]$ application show host *

                              | Stack |   Host Mode Attributes    |
Index | Host Name             | Index | OS Type | Characteristics |
------------------------------------------------------------------
00304   co-ls1                 00000    LINUX     0x0000000000000001
00305   co-test-d10            00000    WINDOWS   0x0000000000000001
00306   co-test-d08            00000    WINDOWS   0x0000000000000001

Total Hosts: 3
```

Recommendations:

- When creating host objects, always include a descriptive “name” for the host objects that is easily recognizable. This benefits the storage administrator in managing storage connectivity issues.
- The **OS_TYPE** determines the FC flow control mechanism for host I/O.
- The Linux OS type is currently appropriate for all versions/implementations of Linux and UNIX.
- The Windows OS type is currently appropriate for all versions of the Windows operating systems.
- It is imperative the **OS_TYPE** for Apple/MAC hosts be appropriately set to MAC. The setup requirements involving MAC OSX hosts require additional steps and considerations. Please see Section 3.4.5, "Additional Configuration Considerations for Macintosh Hosts".

3.4.4.2 Identifying Host FC Connections via Ports

The host ports of the Controller dynamically acknowledge and log in any FC Host/Port WWN connection that it can sync up with. When you query the storage subsystem for the discovered initiators, the listing you get will be the current set of connections that are logged in. If a connection

is broken and then re-established, the re-discovered Port/Node WWN will be assigned a new host_index number. This is inconsequential after a specific Port/Node WWN is directly associated with a specific host that is done in a subsequent step.

Example of discovered initiators:

```
RAID$ app show discovered_initiators *

                            |          Initiator Identifier            | Controller 0  | Controller 1  |
Index  |Type|    ID         |     node          |       port           | RP 0   | RP 1  | RP 0  | RP 1  |
-------------------------------------------------------------------------------------------------------
00003    FC  |0x010600 0x20000000c9813cc9 0x20000000c9813cc9          |       |    3  |
00005    FC  |0x010800 0x20000000c9813cc8 0x20000000c9813cc8          |       |       |0
00006    FC  |0x010900 0x20000000c9813a47 0x20000000c9813a47          |       |   2   |
00007    FC  |0x010a00 0x20000000c9813a46 0x20000000c9813a46          |       |       |     1
00011    FC  |0x010300 0x2000001b32827e95 0x2000001b32827e95          |0      |       |
00012    FC  |0x010200 0x2000001b32a27e95 0x2000001b32a27e95          |0      |       |

Total FC Initiators: 6
```

In order to provide an association between a host and a Virtual Disk, you must identify the FC connection (its Node/Port WWN name to the host object). You may use the following techniques:

- Utilize the physical sticker tag information on the HBA.
- Connect a single host at a time (and all its FC connections to the subsystem) to the subsystem (direct connect or via a switch); identify their WWNs as a discovered initiator; then a proceed to plug in subsequent hosts, noting their WWNs.
- Utilize a host based FC Adapter utility such as HBAnyware (for Emulex HBAs) (Figure 111) or SANsurfer (for QLogic HBAs) (Figure 112).

Figure 111. HBAnyware Screen

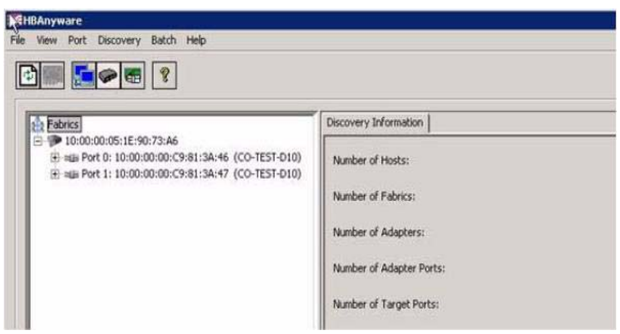

Figure 112. SANsurfer FC HBA Manager Screen

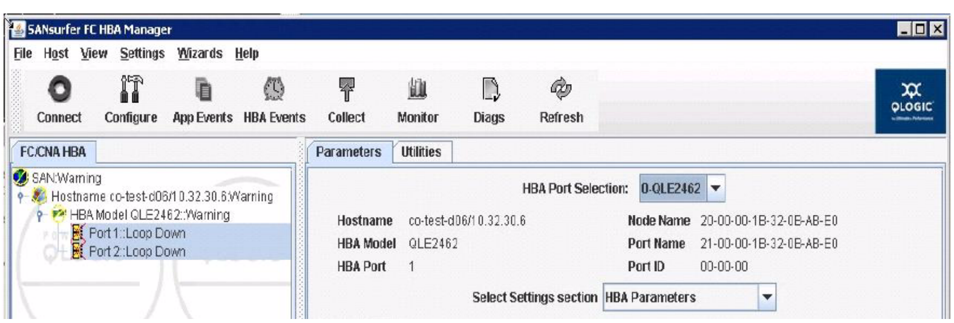

Example for Linux system with driver loaded (example of QLogic):

```
# cd /proc/scsi/qla2xxx
# ls
1 2
# grep adapter-port 1
scsi-qla0-adapter-port=210000e08b9d6149;
bm-dell-09:/proc/scsi/qla2xxx # grep adapter-port 2
scsi-qla1-adapter-port=210100e08bbd6149;
# lsscsi -g | grep -i sgi
[1:0:0:1] disk SGI DD4A-FC008-6620 1.03 /dev/sdc /dev/sg2
[2:0:0:1] disk SGI DD6A-IS16K-10000 1.03 /dev/sdg /dev/sg6
```

Example for Linux system without driver loaded:

```
# cd /sys/class/fc_host
# ls -la
total 0
drwxr-xr-x 4 root root 0 Jul 15 11:47 .
drwxr-xr-x 43 root root 0 Jul 15 11:48 ..
drwxr-xr-x 3 root root 0 Jul 15 11:48 host3
drwxr-xr-x 3 root root 0 Jul 15 11:48 host4
# cd host3
# ls -la
total 0
drwxr-xr-x 3 root root 0 Jul 15 11:48 .
drwxr-xr-x 4 root root 0 Jul 15 11:47 ..
lrwxrwxrwx 1 root root 0 Jul 15 11:47 device ->
../../../devices/pci0000:00/0000:00:04.0/0000:0a:00.0/host3
-r--r--r-- 1 root root 4096 Jul 17 15:31 fabric_name
--w------- 1 root root 4096 Jul 17 15:31 issue_lip
-r--r--r-- 1 root root 4096 Jul 17 15:31 node_name
-r--r--r-- 1 root root 4096 Jul 17 15:31 port_id
-r--r--r-- 1 root root 4096 Jul 17 15:31 port_name
-r--r--r-- 1 root root 4096 Jul 17 15:31 port_state
-r--r--r-- 1 root root 4096 Jul 17 15:31 port_type
-r--r--r-- 1 root root 4096 Jul 17 15:31 speed
drwxr-xr-x 2 root root 0 Jul 15 11:47 statistics
lrwxrwxrwx 1 root root 0 Jul 15 11:48 subsystem -> ../../../class/fc_host
-r--r--r-- 1 root root 4096 Jul 17 15:31 supported_classes
-r--r--r-- 1 root root 4096 Jul 17 15:31 symbolic_name
-rw-r--r-- 1 root root 4096 Jul 17 15:31 system_hostname
-rw-r--r-- 1 root root 4096 Jul 17 15:31 tgtid_bind_type
--w------- 1 root root 4096 Jul 15 11:47 uevent
# cat port_name
0x2100001b3282dc50
```

NOTE : In switch environments you may disable/enable the ports on the switch to identify which physical connection you are dealing with.

Utilizing one or more of the techniques above, you can document an association between the physical host and the discovered initiators on the storage.

Example of showing imported initiators:

```
RAID[1]$ application show initiator *


             | Host    |          Initiator Identifier       |
Index | Type | Index   |   node           |     port         |
-------------------------------------------------------------
00004   FC    0x010600  0x20000000c98107cb 0x10000000c98107cb <-- co-test-d08
00006   FC    0x010900  0x20000000c9813a47 0x10000000c9813a47 <-- co-test-d10
00007   FC    0x010a00  0x20000000c9813a46 0x10000000c9813a46 <-- co-test-d10
00008   FC    0x010300  0x20000000c98107ca 0x10000000c98107ca <-- co-test-d08
00011   FC    0x010200  0x2000001b32827e95 0x2100001b32827e95 <-- co-ls1
00012   FC    0x010800  0x2001001b32a27e95 0x2101001b32a27e95 <-- co-ls1
```

Now identify the WWN and the discovered initiator index number with the host and its index number and import the appropriate discovered initiator to an association with the appropriate host.

Example of importing discovered initiators:

```
RAID[1]$ application import discovered 11 host 304
INITIATOR 12 OID=0x2d66000c imported from DISCOVERED_INITIATOR 11 OID=0x3000000b
STATUS='Success' (0x0)

RAID[1]$ application import discovered 12 host 304
INITIATOR 13 OID=0x2d67000d imported from DISCOVERED_INITIATOR 12 OID=0x3000000c
STATUS='Success' (0x0)

RAID[1]$ application import discovered 4 host 306
INITIATOR 14 OID=0x2d68000e imported from DISCOVERED_INITIATOR 4 OID=0x30000004
STATUS='Success' (0x0)

RAID[1]$ application import discovered 8 host 306
INITIATOR 15 OID=0x2d69000f imported from DISCOVERED_INITIATOR 8 OID=0x30000008
STATUS='Success' (0x0)

RAID[1]$ app import discovered=6 host=305
INITIATOR 16 OID=0x2d6a0010 imported from DISCOVERED_INITIATOR 6 OID=0x30000006
STATUS='Success' (0x0)

RAID[1]$ app imp disc 7 host 305
INITIATOR 17 OID=0x2d6b0011 imported from DISCOVERED_INITIATOR 7 OID=0x30000007
STATUS='Success' (0x0)

RAID[1]$ application show initiator *

             | Host  |            Initiator Identifier       |
Index | Type | Index |      node name     |   port name      |
-------------------------------------------------------------
00012   FC    00304    0x2000001b32827e95  0x2100001b32827e95
00013   FC    00304    0x2001001b32a27e95  0x2101001b32a27e95
00014   FC    00306    0x20000000c98107cb  0x10000000c98107cb
00015   FC    00306    0x20000000c98107ca  0x10000000c98107ca
00016   FC    00305    0x20000000c9813a47  0x10000000c9813a47
00017   FC    00305    0x20000000c9813a46  0x10000000c9813a46

Total FC Initiators: 6
```

3.4.4.3 Presenting a Virtual Disk to a Host

From your list of Virtual Disks that have been created, identify the host that you wish to present this LUN to and create the presentation.

Recommendations:

- Commands entered through the CLI may take a shortened form of the unique word as demonstrated in the above examples
- Spaces may be used in place of "=" sign in command syntax

Example of list of candidate Virtual Disks for presentation:

```
RAID[1]$ show vd *

                                                                |      Home       |Background|
Idx|  Name   | State |Pool|Raid|Cap GB|Settings| Jobs |Current|Preferred| Job |
-------------------------------------------------------------------------------------
 49 vd-49_29   READY   29   6   1400     W I          1(L) 0   1(L) 0   INACTIVE
 50 vd-50_30   READY   30   6   1400     W I          0(R) 0   0(R) 0   INACTIVE
 51 vd-51_31   READY   31   1    344     W I          1(L) 0   1(L) 0   INACTIVE
 52 vd-52_32   READY   32   1    344     W I          0(R) 0   0(R) 0   INACTIVE
 53 vd-53_33   READY   33   5   1400     W I          1(L) 0   1(L) 0   INACTIVE
 54 vd-54_34   READY   34   5   1400     W I          0(R) 0   0(R) 0   INACTIVE
 55 vd-55_35   READY   35   6   3608     W I          1(L) 0   1(L) 0   INACTIVE
 56 vd-56_36   READY   36   6   3608     W I          0(R) 0   0(R) 0   INACTIVE
 57 vd-57_37   READY   37   5   3608     W I          1(L) 0   1(L) 0   INACTIVE
 58 vd-58_38   READY   38   5   3608     W I          0(R) 0   0(R) 0   INACTIVE
 59 vd-59_39   READY   39   1    344     W I          1(L) 0   1(L) 0   INACTIVE
 60 vd-60_41   READY   41   1    896     W I          1(L) 0   1(L) 0   INACTIVE

Total Virtual Disks: 12
```

Example of creating VD presentation to Windows host co-test-d10 (not specifying SCSI LUN_ID):

The next two VDs are presented to a Windows host and the subsystem assigns the SCSI LUN_ID.

```
RAID[1]$ application create presentation vd 51 host 305
PRESENTATION 2482 OID=0x256e09b2 creation STATUS='Success' (0x0)

RAID[1]$ application create presentation vd 52 host 305
PRESENTATION 2483 OID=0x256f09b3 creation STATUS='Success' (0x0)

RAID[1]$ application show presentation *

                                                          |      Channel Mask         |
Pres. | Host           Host |   VD  |     |Home |Read | Controller 0|Controller 1 |
Index | Name          Index | Index | LUN |Only |Only | 0123 | 0123 | 0123 | 0123 |
-------------------------------------------------------------------------------------
02480  co-ls1         00304   00049   060   OFF   R/W   ffff   ffff   ffff   ffff
02481  co-ls1         00304   00050   061   OFF   R/W   ffff   ffff   ffff   ffff
02482  co-test-d10    00305   00051   000   OFF   R/W   ffff   ffff   ffff   ffff
02483  co-test-d10    00305   00052   001   OFF   R/W   ffff   ffff   ffff   ffff

Total Presentations: 4
```

NOTE : For Windows environments, you should let the subsystem specify the LUN_ID for a Virtual Disk.

```
RAID[1]$ application create presentation vd 58 host 306 lun 10
PRESENTATION 2484 OID=0x257009b4 creation STATUS='Success' (0x0)

RAID[1]$ application create presentation vd 59 host 306 lun 12
PRESENTATION 2485 OID=0x257109b5 creation STATUS='Success' (0x0)

RAID[1]$ app show pres *

                                                                |      Channel Mask        |
Pres. | Host           Host |   VD  |     |Home |Read | Controller 0|Controller 1 |
Index | Name          Index | Index | LUN |Only |Only | 0123 | 0123 | 0123 | 0123 |
-------------------------------------------------------------------------------------
02480  co-ls1         00304   00049   060   OFF   R/W   ffff   ffff   ffff   ffff
02481  co-ls1         00304   00050   061   OFF   R/W   ffff   ffff   ffff   ffff
02482  co-test-d10    00305   00051   000   OFF   R/W   ffff   ffff   ffff   ffff
02483  co-test-d10    00305   00052   001   OFF   R/W   ffff   ffff   ffff   ffff
02484  co-test-d08    00306   00058   010   OFF   R/W   ffff   ffff   ffff   ffff
02485  co-test-d08    00306   00059   012   OFF   R/W   ffff   ffff   ffff   ffff

Total Presentations: 6
```

Recommendations:

- At times some hosts may have device discovery issues if there are “gaps” in the SCSI LUN_ID space. Letting the subsystem assign (as it will by default) the lowest SCSI_ID it can from this subsystem, may be considered a best practice.
- Use of a host based HBA utility such as HBAnyware® or SANsurfer greatly enhances the ability to determine perceived connectivity issues between the storage subsystem and the host.
- Your hosts “multipath driver” must be enabled correctly to recognize the SGI Device/Hardware ID.
  - For Windows Server 2008, the Device/Hardware_ID information must be entered into the MPIO stack as an 8-bit / 16-character space-padded field.
  - For Linux, the /etc/multipath.conf file must be edited, and the Device/Hardware_ID is NOT padded with the space character.

    ```
    Device Hardware_ID
    SGI DD6A-IS16K-10000
    ```

3.4.5 Additional Configuration Considerations for Macintosh Hosts

Apple/Mac Operations systems handles Fibre Channels Port/Node WWNs uniquely. Apple hosts expect node names to be unique across all subsystem ports.

All other OS's expect the FC node names to be the same across all subsystem ports.

Because of this uniqueness, VD presentations must be uniquely set up when an Apple/Mac presentation is being made from a IS16000 subsystem.

NOTE : You MUST create the host object with **OSTYPE=MAC** for Mac hosts.

You must identify at least one Controller port (channel) (preferably one Controller port for each IS16000 Controller) on the storage subsystem as operating in the MAC mode:

```
APPLICATION SET CHANNEL [0|1|2|3] MODE MAC
```

NOTE : Setting/changing the mode of a "channel" will require a shutdown/restart of the Controller for the change to actually take place.

If a Mac presentation is being made from the subsystem, NO other VD from this subsystem should be presented to "ALL" hosts !

In a SAN and/or MIXED OS environment of MAC and other operating systems, create the presentations as follows:

```
APPLICATION CREATE PRESENTATION VD=<vd-id> HOST=<host_index> ENABLE=NONE
```

This will disable channels through which to make a presentation. Then you will next enable the specific channels through which you wish to make the presentation:

```
APPLICATION SET PRESENTATION <index> ENABLE {0|1|2|3}
```

To enable two channels, it is necessary to enter two separate commands.

In the current implementation, you can disable all channels, enable all channels, or enable one channel at a time through the CLUI. If you wish to have two channels enabled, use the following commands:

```
APPLICATION SET PRESENTATION <index> ENABLE 0
APPLICATION SET PRESENTATION <index> ENABLE 2
```

Channel 0 is C0RP0. Channel 1 is C0RP1. These are the two RAID processors on Controller 0.

Channel 2 is C1RP0. Channel 3 is C1RP1. These are the two RAID processors on Controller 1.

If you have an existing storage configuration running Windows/Linux presentations and you wish to ADD a VD presentation for MAC:

- You will have to sustain a storage outage now so that you can disable "all" VD presentations thru any channel; and
- Re-enable the appropriate channels one at a time that will be serving VDs to the Windows/Linux hosts.
- You will then set up the appropriate channel(s) that will operate in MAC mode.
- At this point you can restart the storage subsystem, proceed in creating your MAC VD presentations and enable them through their separate channels.

3.5 Firmware Update Management

SGI periodically releases firmware updates to enhance features of the products. Please contact our Technical Support Department to obtain the latest firmware files.

Before performing any updates, please note the following:

- The **UPDATE ENCLOSURE** command starts a background operation so that the CLUI is free to do other operations during the **UPDATE**. The status that comes back from the **UPDATE ENCLOSURE** command only indicates whether the background operation was successfully started. **SHOW ENCLOSURE <id> UPDATE_FIRMWARE_PROGRESS** should be used immediately after the command and then subsequently periodically to check the status of the **UPDATE**.
- When you SCP the Consolidated Download File (.SGI) to the firmware directory, it must be copied to the PRIMARY Controller. Since the PRIMARY may change, it is probably best to copy it to BOTH Controllers. If **SHOW ENCLOSURE <id> UPDATE_FIRMWARE_PROGRESS** reports status JS_ES_DL_UCODE_OPEN_FAILED, the most likely cause is that the file was not copied to the PRIMARY Controller.
- When the IS16000 is busy rebuilding, initializing, or doing host I/O, **UPDATE ENCLOSURE** operations may fail because the enclosure is too busy to service the UPDATE IOs in a timely fashion. Therefore, make sure that there are no rebuilds, initializations, or host I/O operations being done during the **UPDATE ENCLOSURE** operation.
- In order to cause the new enclosure firmware to start being used you must power cycle the enclosure. In the IS16000, it is best to do this by doing a **SHUTDOWN SUBSYSTEM** followed by power-cycling the enclosure. This is also required to recover from an **UPDATE ENCLOSURE** failure. These power cycles are required.
- The **UPDATE PD** command starts a background operation so that the CLUI is free to do other operations during the **UPDATE**. The status that comes back from the **UPDATE PD** command only indicates whether the background operation was successfully started. **SHOW PD * UPDATE_FIRMWARE_PROGRESS** should be used immediately after the command and then subsequently periodically to check the status of the **UPDATE**.
- When you SCP the Consolidated Download File (.SGI) to the firmware directory, it must be copied to the PRIMARY Controller. Since the PRIMARY may change, it is probably best to copy it to BOTH Controllers. If **SHOW PD * UPDATE_FIRMWARE_PROGRESS** reports status JS_ES_PD_DL_UCODE_OPEN_FAILED for a physical disk, the most likely cause is that the file was not copied to the PRIMARY Controller.
- Online physical disk firmware update is not currently supported. Therefore, make sure that there are no rebuilds, initializations, or host IO operations being done during the **UPDATE PD** operation. This is enforced by making sure the user has issued the **SET SUBSYSTEM OFFLINE** command before allowing an **UPDATE PD** command.
- In order to cause the new physical disk firmware to start being used you must power cycle the enclosure in which the physical disk is installed. In the IS16000, it is best to do this by doing a **SHUTDOWN SUBSYSTEM** followed by power-cycling the enclosure.

3.5.1 Displaying Current Firmware Version

The **SHOW CONTROLLER <id> ALL_ATTRIBUTES** command displays version information of the IS16000's hardware and firmware (Figure 113).

Figure 113. Show Controller Information Screen

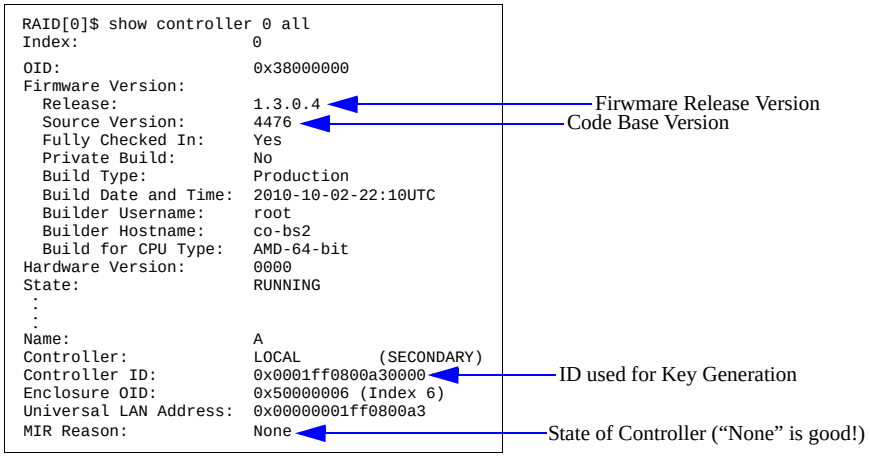

3.5.2 Controller Firmware Update Procedure

The firmware update procedure described here only applies to upgrades from firmware version 1.3.0.4.xxxx or later to a new version. Upgrades from an older version of firmware must be done by a trained technician.

NOTE : The IS16000 now supports an upgrade of the firmware while the system is still online. However, you must correctly follow the upgrade instructions to perform a successful online upgrade.

If you are upgrading from v1.3.0.4.xxxx or later, there are two methods that can be utilized:

- Copy the new firmware image over to the Controller via the network instead of using a USB flash disk for Linux.
- Copy the new firmware image over to the Controller via the network for Windows users.

NOTE : In the examples, the filename of the firmware image is **sgi-flash-2812-opt.tgz** and the IP address of the controller is **10.32.31.240**. Replace these parameters with the filename and IP address appropriate for your installation.

3.5.2.1 Linux Environment-Firmware Update from the Network

Follow these steps to update the firmware:

1. Copy the new firmware to both Controllers using an scp or sftp client. The user name is **`firmware`** and the password is **`Firmware`** (Note that entries are case-sensitive).
   For example:
   `scp sgi-flash-2812-opt.tgz firmware@10.32.31.240:`

2. At the CLUI prompt, enter command:
   `UPDATE_FIRMWARE CONTROLLER=LOCAL FILE="<file-specification>"`

   For example:
   `UPDATE_FIRMWARE CONTROLLER=LOCAL FILE="sgi-flash-2812-opt.tgz"`
   (Note that the file name must be enclosed with double quotation marks.)

3. Upon completion of reboot, login and enter the command:
   `SHOW CONTROLLER LOCAL ALL`
   Verify that the firmware version is correct.

   NOTE : The pools may indicate there is a fault. Issue the command **`SHOW POOLS *`** to check the details of the pool. The pools will go into WR thru state on initial startup to allow the batteries the opportunity to charge. The WR Thru mode is necessary to allow for data integrity until battery redundancy on each Controller is completed. If there is a battery issue, the system will indicate replacement is required and all VDs will remain in WR Thru mode until the issue is resolved.

   For example:
   `CLUI unable to get RAID SUBSYSTEM NAME. STATUS='MIR:Firmware Version Mismatch' (0x30003f8)`

   NOTE : The firmware will be in MIR, firmware version mismatch state for the first Controller that is upgraded. You must ensure that you have a CLUI prompt from the newly upgraded Controller before moving to Step 4.

   ---

4. Upgrade the second Controller by repeating Steps 2-3 above on the other Controller.

3.5.2.2 Windows Environment-Firmware Update from the Network

This method requires a SFTP or SCP client, and involves two steps – copying the firmware image to both Controllers and launching the update.

1. Copy the new firmware to the Controller using an scp or sftp client. You must use an application that supports SFTP or SCP. WinSCP, a free SFTP and SCP client, is used in this example.

2. Launch WinSCP or other SFTP/SCP client. For hostname, use the IP address of your Controller. user name is **`firmware`** and password is **`Firmware`**. Both user name and password are case-sensitive.

3. Select Login .

   At the login screen (Figure 114), use the user name **`admin`** with the password **`password`** to log into the system. User name and passwords are care sensitive.

Figure 114. Login Screen

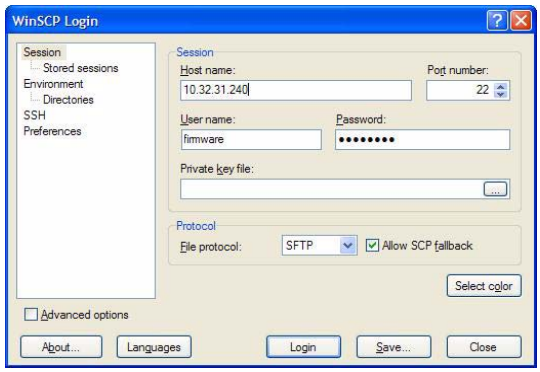

4. Find the image file and drag this to the destination window and select copy.
Copy operation commences.

To launch the update:

1. Login to the Controller via an SSH session with the user name **`user`** and password **`user`**.

2. At the CLI prompt, enter:
`update_firmware controller=<controller-id> FILE="<file-specification>"`

   For example:
`update_firmware controller=LOCAL FILE= "sgi-flash-2812-opt.tgz"`

NOTE : The quotes around the filename of the firmware image are required.

---

3. Upon completion of reboot, login and enter the command:
`SHOW CONTROLLER LOCAL ALL`
Verify that the firmware version is correct.

NOTE : The pools may indicate there is a fault. Issue the command **`SHOW POOLS *`** to check the details of the pool. The pools will go into WR thru state on initial startup to allow the batteries the opportunity to charge. The WR Thru mode is necessary to allow for data integrity until battery redundancy on each Controller is completed. If there is a battery issue, the system will indicate replacement is required and all VDs will remain in WR Thru mode until the issue is resolved.

For example:
`CLUI unable to get RAID SUBSYSTEM NAME. STATUS='MIR:Firmware Version Mismatch' (0x30003f8)`

NOTE : The firmware will be in MIR, firmware version mismatch state for the first controller that is upgraded. You must ensure that you have a CLUI prompt from the newly upgraded controller before moving to Step 4.

---

4. Upgrade the second Controller by repeating Steps 2-3 above on the other Controller.

3.5.3 Disk Enclosure Firmware Update Procedure

Before starting this procedure, verify that a power cycle or any perturbations to the system is NOT required for the next twohours. This procedure requires that all host IO is shutdown prior to its use. Any outstanding formats or rebuilds must not be running. Check the status of any outstanding formats or rebuilds and verify that no disk IO actiivy is in process before updating your firmware.

There are two (2) parts to update the enclosure firmware:

1. Upload the enclosure firmware to the enclosure using either Linux Environment or Windows.
2. Update Disk Shelf using CLUI commands.

3.5.3.1 Linux Environment-Firmware Upload

Obtain the firmware from SGI. Firmware must be running on the IS16000 and must be connected to the enclosure to be upgraded.

1. Copy the new firmware to the Controller using an scp or sftp client. The user name is **`firmware`** and the password is **`Firmware`**. (Note that entries are case-sensitive.)
   Use the secured copy program (scp) to transfer the file to the expander with the command: **`scp <firmware file name> firmware@<ip_address_of_IS16000>:`**

   NOTE : Ensure the colon (:) is at the end of the above command.

2. Enter the password **`Firmware`**. (Note that entries are case-sensitive).

The file will be copied to the Controller.

3.5.3.2 Windows Environment-Firmware Upload

Obtain the firmware from SGI. Firmware must be running on the IS16000 and must be connected to the enclosure to be upgraded.

1. Use the putty secured copy program (pscp) to transfer the file to the expander with the command: **`pscp <firmware file name> firmware@<ip_address_of_IS16000>:`**

   NOTE : Ensure the colon (:) is at the end of the above command.

2. Enter the password **`Firmware`**. (Note that entries are case-sensitive).

The file will be copied to the Controller.

3.5.3.3 Disk Shelf Upgrade

Once the file is copied from either procedure above (Linux or Windows), follow the steps below to complete the upgrade:

1. Log onto the Controller.
   User name is **`user`**. Password is **`user`**.

2. At the CLUI prompt, enter the command:
   UPDATE ENCLOSURE <enclosure num> FILE=<file name uploaded>

   The upgrade will take approximately 2 hours. You can have multiple upgrades going simultaneously so that all disk enclosures can be upgraded at the same time to expedite the upgrade.

3. To monitor the upgrade progress, enter the command:
   SHOW ENCLOSURE 1 UPDATE_FIRMWARE_PROGRESS

   For example:

```
RAID[0]$ show enclosure 1 update_firmware_progress
ES download progress for enclosure 1.
  CDF file name           /tmp/janus_update//IS_D02_011.SGI
  CDF is for Vendor ID    SGI
  CDF is for Product ID   IS OS
  CDF package version     D02.011
  CDF release date        01-SEP-2010
  Download is in progress and 15 percent complete.
  Download consists of 2 images.
  Image 2 is in progress and 13 percent complete.
```

4. When the download is complete, you must power-cycle the enclosure.

5. To verify that the update was successful, enter the command:
   SHOW ENCLOSURE 1 UPDATE_FIRMWARE_PROGRESS

   For example:

```
RAID[0]$ show enclosure 1 update_firmware_progress
ES download progress for enclosure 1.
  Download not in progress -- last download completed successfully.
```

6. Verify that the upgrade was successful by using the command: **SHOW EXPANDER**

NOTE : The expanders are Sub Index 3 through 10. Ensure that the Firmware version and the Init String version are identical.

For example:

```
RAID[0]$ show expander 1 * all
Sub Index:              1
Sub OID:                0x78000001
Enclosure Index:        1
Enclosure OID:          0x50000001
Position:               1
SES Status:             OK
Present:                TRUE
Predicted Failure Ind:  OFF
Locate Indicator:       OFF
Firmware version:       0078

Sub Index:              2
Sub OID:                0x78000002
Enclosure Index:        1
Enclosure OID:          0x50000001
Position:               2
SES Status:             OK
Present:                TRUE
Predicted Failure Ind:  OFF
Locate Indicator:       OFF

Sub Index:              3
Sub OID:                0x78000003
Enclosure Index:        1
Enclosure OID:          0x50000001
Position:               3
SES Status:             OK
Present:                TRUE
Predicted Failure Ind:  OFF
Locate Indicator:       OFF
Part number:            TCA-00300-01-A
Serial number:          MXSCI00089HVD12B
Firmware version:       D02.011
Init string version:    D02.011
FPGA version:           7

Sub Index:              4
Sub OID:                0x78000004
Enclosure Index:        1
Enclosure OID:          0x50000001
```

```
Position:              4
SES Status:            OK
Present:               TRUE
Predicted Failure Ind: OFF
Locate Indicator:      OFF
Part number:           TCA-00300-01-A
Serial number:         MXSCI00089HVD110
Firmware version:      D02.011
Init string version:   D02.011
FPGA version:          7

Sub Index:             5
Sub OID:               0x78000005
Enclosure Index:       1
Enclosure OID:         0x50000001
Position:              5
SES Status:            OK
Present:               TRUE
Predicted Failure Ind: OFF
Locate Indicator:      OFF
Part number:           TCA-00300-01-A
Serial number:         MXSCI00085QVD2CB
Firmware version:      D02.011
Init string version:   D02.011
FPGA version:          7

Sub Index:             6
Sub OID:               0x78000006
Enclosure Index:       1
Enclosure OID:         0x50000001
Position:              6
SES Status:            OK
Present:               TRUE
Predicted Failure Ind: OFF
Locate Indicator:      OFF
Part number:           TCA-00300-01-A
Serial number:         MXSCI00086QVD109
Firmware version:      D02.011
Init string version:   D02.011
FPGA version:          7

Sub Index:             7
Sub OID:               0x78000007
Enclosure Index:       1
Enclosure OID:         0x50000001
Position:              7
SES Status:            OK
Present:               TRUE
Predicted Failure Ind: OFF
Locate Indicator:      OFF
Part number:           TCA-00300-01-A
Serial number:         MXSCI000893VD19D
Firmware version:      D02.011
Init string version:   D02.011
FPGA version:          7

Sub Index:             8
Sub OID:               0x78000008
Enclosure Index:       1
Enclosure OID:         0x50000001
Position:              8
SES Status:            OK
Present:               TRUE
Predicted Failure Ind: OFF
Locate Indicator:      OFF
Part number:           TCA-00300-01-A
Serial number:         MXSCI00085LVD0A5
Firmware version:      D02.011
Init string version:   D02.011
FPGA version:          7

Sub Index:             9
Sub OID:               0x78000009
Enclosure Index:       1
Enclosure OID:         0x50000001
Position:              9
SES Status:            OK
Present:               TRUE
Predicted Failure Ind: OFF
Locate Indicator:      OFF
Part number:           TCA-00300-01-A
Serial number:         MXSCI00089HVD11A
Firmware version:      D02.011
Init string version:   D02.011
FPGA version:          7

Sub Index:             10
Sub OID:               0x7800000a
Enclosure Index:       1
Enclosure OID:         0x50000001
Position:              10
SES Status:            OK
Present:               TRUE
Predicted Failure Ind: OFF
Locate Indicator:      OFF
Part number:           TCA-00300-01-A
Serial number:         MXSCI00089HVD12C
Firmware version:      D02.011
Init string version:   D02.011
FPGA version:          7


Total Expanders: 10
```

3.5.4 Physical Disk Firmware Update Procedure

This procedure requires that all host IO is shutdown prior to its use. Any outstanding formats or rebuilds must not be running. Check the status of any outstanding formats or rebuilds and verify that no disk IO actiivy is in process before updating your firmware.

There are two (2) parts to update the physical disk firmware:

1. Upload the physical disk firmware to the IS16000 using either Linux Environment or Windows.
2. Update the Physical Disk using CLUI commands.

3.5.4.1 Linux Environment-Firmware Upload

Obtain the firmware from SGI.

1. Use the secured copy program (scp) to transfer the file to the IS16000 with the command: **`scp <firmware file name> firmware@<ip_address_of_IS16000>:`**

NOTE : Ensure the colon (:) is at the end of the above command.

2. Enter the password (case-sensitive) **`Firmware`**.

The file is copied to the IS16000.

3.5.4.2 Windows Environment-Firmware Upload

Obtain the firmware from SGI.

1. Use the putty secured copy program (pscp) to transfer the file to the IS16000 with the command: **`pscp <firmware file name> firmware@<ip_address_of_IS16000>:`**

NOTE : Ensure the colon (:) is at the end of the above command.

2. Enter the password (case-sensitive) **`Firmware`**.

The file is copied to the IS16000.

3.5.4.3 Physical Disk Upgrade

Once the file is copied from either procedure above (Linux or Windows), follow the steps below to complete the upgrade:

1. Log onto the Controller.
   User name is **`user`**. Password is **`user`**.
2. At the CLUI prompt, enter command:
   `SET SUBSYSTEM OFFLINE`

NOTE : The subsystem must be offline before proceeding with the remaining steps.

3. At the CLUI prompt, enter command:
UPDATE PD * FILE=<file name uploaded>

The firmware will be updated on any physical disks that match the Vendor Id, Product Id, and Drive Type (SAS/SATA) in the file. The upgrade time varies but typically takes less than a minute.

For example:

```
RAID[0]$ update pd * file="WD2002FYPS_05D07.SGI"
A request has been made to update physical disk firmware. No attempt will be made to preserve
redundancy if the physical disk resides in a pool. If the update to the physical disk
encounters catastrophic failure, then this may result in lost data if the pool loses all
of its redundancy.
It is recommended that a backup is done prior to this operation.
Are you sure you want to update physical disk firmware [Yes]?
PHYSICAL DISK 65535 OID=0xffffffff firmware update started STATUS='An asynchronous command
has been started' (0x3000068)
The controller has initiated the download operation which may take up to several minutes.
Check the progress and status of the download operation with SHOW PD n
UPDATE_FIRMWARE_PROGRESS. Additional information may be found in the event log.
```

4. To monitor the upgrade progress, enter command:
SHOW PD * UPDATE_FIRMWARE_PROGRESS

For example (non-matching drives omitted in output for clarity):

```
RAID[0]$ show pd * update_firmware_progress
Encl|Slot| Vendor  |     Product ID    |Type| Revision| Serial Number  | Pool| Idx|      WWN        |% Complete|   Status
-------------------------------------------------------------------------------------------------------------------------
1      5   HITACHI    HUS153014VLS300    SAS     A410    JFVWMZNC5301      UNAS  271  50014ee2032d2555   56%         In Progress
1      6   HITACHI    HUS153014VLS300    SAS     A410    J4V1YYAA5301      UNAS  272  50014ee2032d196c   56%         In Progress
1     17   HITACHI    HUS153014VLS300    SAS     A410    JFVWL5YC5301      UNAS  274  50014ee2032ca9a9   56%         In Progress
1     18   HITACHI    HUS153014VLS300    SAS     A410    JFVSU5PC5301      UNAS  269  50014ee25882b499   56%         In Progress
```

5. When the upgrade has completed, to verify that the download succeeded, enter command:
SHOW PD * UPDATE_FIRMWARE_PROGRESS

For example (non-matching drives omitted in output for clarity):

```
RAID[0]$ show pd * update_firmware_progress
Encl|Slot| Vendor |     Product ID  |Type| Revision|   Serial Number| Pool|Idx|    WWN          |% Complete|Status
-------------------------------------------------------------------------------------------------------------------------
1     5   HITACHI   HUS153014VLS300  SAS     A410    JFVWMZNC5301     UNAS  271  50014ee2032d2555    Done    JS_GBL_SUCCESS (Success)
1     6   HITACHI   HUS153014VLS300  SAS     A410    J4V1YYAA5301     UNAS  272  50014ee2032d196c    Done    JS_GBL_SUCCESS (Success)
1    17   HITACHI   HUS153014VLS300  SAS     A410    JFVWL5YC5301     UNAS  274  50014ee2032ca9a9    Done    JS_GBL_SUCCESS (Success)
1    18   HITACHI   HUS153014VLS300  SAS     A410    JFVSU5PC5301     UNAS  269  50014ee25882b499    Done    JS_GBL_SUCCESS (Success)
```

6. Shutdown the subsystem, enter command:
SHUTDOWN SUBSYSTEM

7. Power-cycle the enclosures containing the updated physical disks.

8. After powering up the disk enclosures and Controller, verify the upgrade was successful, enter command:
SHOW PD *

The firmware version of the updated physical disks should indicate the new version.

For example (non-matching drives omitted in output for clarity):

```
RAID[0]$ show pd *
                                                                                    |Health|
Encl|Slot| Vendor |    Product ID   |Type| Cap GB | RPM|Revision| Serial Number  |Pool| State|Idx | State |  WWN           |
-----------------------------------------------------------------------------------------------------------------------------
1    5     HITACHI HUS153014VLS300   SAS     1863  5.4K A410      JFVWMZNC5301    UNAS  GOOD  279   READY   50014ee2032d2555
1    6     HITACHI HUS153014VLS300   SAS     1863  5.4K A410      J4V1YYAA5301    UNAS  GOOD  280   READY   50014ee2032d196c
1   17     HITACHI HUS153014VLS300   SAS     1863  5.4K A410      JFVWL5YC5301    UNAS  GOOD  282   READY   50014ee2032ca9a9
1   18     HITACHI HUS153014VLS300   SAS     1863  5.4K A410      JFVSU5PC5301    UNAS  GOOD  277   READY   50014ee25882b499
```

3.6 The System Logs

3.6.1 Displaying Event Logs

To display the event log starting at the start sequence number in ascending sequence number order (Figure 115), enter command:

SHOW CONTROLLER <id> LOG ASCEND

Figure 115. Show Controller Log Example Screen (1)

```
RAID[0]$ show controller 0 log ascend

000001 2010-06-04 01:03:26:6418393 G=0 S=0 T=1 RP=0 VP=1
LOG_ES_UPS_CLIENT_ATTRIBUTE_NAME  ES Controller enclosure 0x1ff0800a30000 UPS attribute 0.......

000002 2010-06-04 01:03:26:6418400 G=0 S=0 T=1 RP=0 VP=1
LOG_ES_UPS_CLIENT_ATTRIBUTE_CHANGED  ES Controller UPS attribute 0 changed to 37......

000003 2010-06-04 01:04:52:8320854 G=3 S=1 T=1 RP=0 VP=1
LOG_LOGDISK_ENABLE_RECEIVED_FROM_STATE LOG RECEIVED FROM STATE
```

To display the event log starting at the start sequence number in descending sequence number order (Figure 116), enter command:

SHOW CONTROLLER <id> LOG DESCEND

Figure 116. Show Controller Log Example Screen (2)

```
RAID[0]$ SHOW CONTROLLER 1 LOG DESCEND

000024 2010-02-11 05:08:48:7027390 G=0 S=0 T=1 RP=0 VP=1 LOG_JOI_TIME_SET JOI TIME
WAS SET BY AN ADMINISTRATOR AT 2010-2-11 5:8:48; NEARBY LOG ENTRIES MAY APPEAR OUT
OF TIME ORDER. OFFSET= 0X1C987C765CD3B2B.

000023 2010-02-06 04:49:20:5069068 G=3 S=1 T=1 RP=0 VP=1 LOG_LOGDISK_ENABLE_RECEIVED
_FROM_STATE LOGRECEIVED FROM STATE

000022 2010-02-06 04:49:20:4952631 G=0 S=0 T=1 RP=0 VP=1 LOG_JOI_TIME_SET JOI TIME
WAS SET BY AN ADMINISTRATOR AT 2010-2-6 4:49:20; NEARBY LOG ENTRIES MAY APPEAR OUT
OF TIME ORDER. OFFSET= 0X1C987C7A302483D.
```

To display the event log starting at the specified start sequence number in ascending sequence number order (Figure 117), enter command:

SHOW CONTROLLER <id> LOG ASCEND START_SEQUENCE <start-sequence-number>

Figure 117. Show Controller Log Example Screen (3)

```
RAID[0]$ SHOW CONTROLLER 1 LOG ASCEND START_SEQUENCE 20

000021 2010-02-06 04:49:20:4731906 G=4 S=2 T=1 RP=0 VP=1 LOG_ST_MIR_STATE STATE
 MIR STATE STATE:000A

000022 2010-02-06 04:49:20:4952631 G=0 S=0 T=1 RP=0 VP=1 LOG_JOI_TIME_SET JOI TIME
WAS SET BY AN ADMINISTRATOR AT 2010-2-6 4:49:20; NEARBY LOG ENTRIES MAY APPEAR
OUT OF TIME ORDER.OFFSET = 0X1C987C7A302483D.

000023 2010-02-06 04:49:20:5069068 G=3 S=1 T=1 RP=0 VP=1 LOG_LOGDISK_ENABLE_
RECEIVED_FROM_STATE LOG RECEIVED FROM STATE

000024 2010-02-11 05:08:48:7027390 G=0 S=0 T=1 RP=0 VP=1 LOG_JOI_TIME_SET JOI TIME
WAS SET BY AN ADMINISTRATOR AT 2010-2-11 5:8:48; NEARBY LOG ENTRIES MAY APPEAR
OUT OF TIME ORDER. OFFSET  = 0X1C987C765CD3B2B.
```

3.6.2 Event Log Structure

Figure 118 below illustrates the structure of an event log.

The severity levels include:

0: Informational
1: Warning
2: Error
3: Fatal

Some additional parameters are:

ES: Enclosure WWN:Slot number (for example, 50001ff101ed0000:60)
IOC: SAS/SATA controller number IOC0 (for example, WWN:5000cca216ed8430)
RC: Reason Code (04 - Device Not Responding)
PHY: Expander PHY device is attached too (for example, PHY:08)
ASCQ: ASC/ASCQ SCSI sense data

Figure 118. Event Log Structure

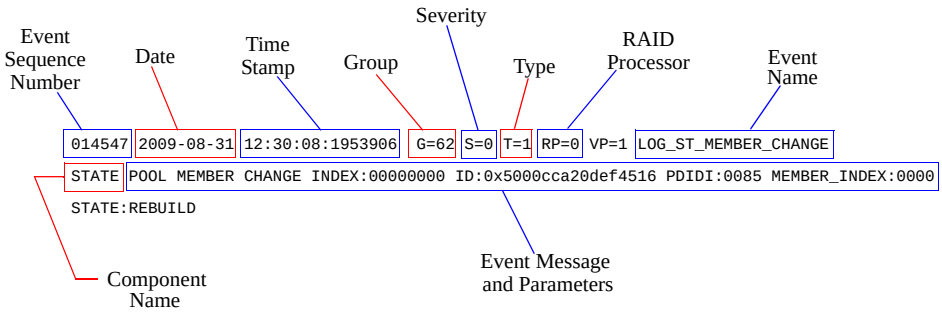

3.6.3 Special Terminology in the Log

- **STATE** – Implements all policy, implements the metadata store, handles all dual-controller issues, and controls all other modules
- **ES** (Enclosure Services) – Monitors all disk enclosures and its local controller enclosure
- **AMPD** and **MPI** – Implements the back-end (SAS/SATA) I/O Controller drivers for access and discovery of Physical Disks and Enclosures
- **CM** (Cache Manager) – Implements read and write-back cache
- **RAID** – Implements RAID-5, RAID-6, RAID-1, and rebuilds
- **MAD** – Data transfer and parity calculations
- **RT** (Routing) – Routes VD request to the right CM
- **IOF** (I/O Forwarding) – Routes PD requests to the right AMPD
- **DUCK** – Dual-controller communication
- **JOI**, **JEX**, **JIPC**, **MIS**, **JTS** – Infrastructure, communication, and testing

3.7 Remote Management of IS16000

The IS16000 can be managed locally through the RS-232 interface, or remotely via SSH. The CLUI is the same regardless of the management interface (RS-232 or SSH).

3.7.1 Network Connection

Connect the Ethernet port on the back of the Controllers to your Ethernet network (Figure 119). Then configure the network interface as described below.

Figure 119. Ethernet Connections to Your Network

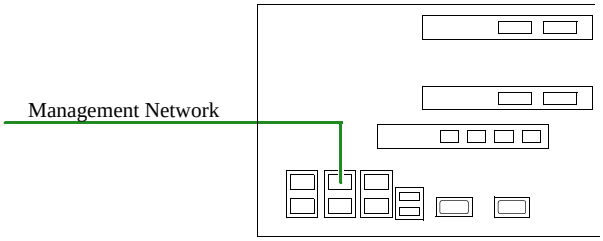

NOTE : Currently, the IS16000 does not support network configuration protocols such as DHCP or BOOTP.

3.7.2 Display Network Interface Attributes

To display the current network interface settings, enter the command (Figure 120):
`UI SHOW NETWORK_INTERFACE [LOCAL|REMOTE|0|1] *`
where **`LOCAL`** gives you information on the Controller that you are currently logged into.

Figure 120. Current Network Interface Settings Screen

```
RAID[0]$ ui show network_interface=local 0 *
Network device id 0
    address 10.32.31.31
    netmask 255.255.240.0
    gateway 10.32.16.2
```

3.7.3 Change Network Interface Settings

NOTE : Initial network interface settings must be configured using the serial interface. Refer to Section 2.8.8, "Configure Network Interface Settings" for more information.

To change the network interface settings on the Controller you are connected to, enter the command (Figure 121):
`UI SET NETWORK_INTERFACE=LOCAL 0 IP_ADDRESS=<ip_address> IP_MASK=<netmask> IP_GATEWAY=<gateway>`

Figure 121. Set Network Interface Example

```
RAID[0]$ ui set network_interface=local 0 ip_address=10.32.31.31 ip_mask=255.255.240.0 ip_gateway=10.32.16.2
NETWORK_INTERFACE 0 set with STATUS='Success' (0x0)
```

3.7.4 Logins

By default, the login name is **user** and its password is **user**. Both are case sensitive.

Only one SSH session is permitted at a time. Once a SSH session is initiated, the RS-232 console switches to a CLI sub-shell. The SSH client should be using port 22 with its local echoing function disabled.

Figure 122. SSH Login Screen

```
login as: user
user@10.23.23.16's password:

Linux (none) 2.6.25-sgi-016620-3 #2 SMP Wed Jan 14 10:38:28 MST 2009 x
SGI InfiniteStorage 16000
```

Logout

To logout, enter command **QUIT**. For SSH connection, the current session will be disconnected.

3.7.5 Email and SNMP Notification Setup

3.7.5.1 Email Setup

Automatic emails will be sent as notification of a selected group of warning and error events that have occurred on the IS16000.

The recipient address is user-configurable and only one address is permitted per system.

To set the email address, enter the command (Figure 123):

UI SET EMAIL IP_ADDRESS=<ip_address> IP_PORT=<port>
FROM= "<email address>" TO="<email address>" SUBJECT="text"

To show all the email address attributes, use the **UI SHOW EMAIL ALL** command (Figure 123).

Figure 123. Email Setup Example Screen

```
RAID[0]$ ui set email ip_address=192.168.0.10
SNMP Agent set with STATUS='Success' (0x0)

RAID[0]$ ui set email ip_port=25
SNMP Agent set with STATUS='Success' (0x0)

RAID[0]$ ui set email to="jdoe@sgi.com"
SNMP Agent set with STATUS='Success' (0x0)

RAID[0]$ ui set email subject="SGI IS16000 Event Notification"
SNMP Agent set with STATUS='Success' (0x0)


RAID[0]$ ui show email

  IP_ADDRESS=192.168.0.10
  IP_PORT=25
  FROM=jdoe@sgi.com
  TO=jdoe@sgi.com
  SUBJECT=SGI IS16000 Event Notification
```

3.7.5.2 SNMP Setup

The Simple Network Management Protocol (SNMP) monitors network attached devices for conditions that warrant administrative attention. The SNMP traps have been implemented to monitor critical and warning events. A management information base (MIB) has also been created to be used to provide inquiry objects and events to the user's monitoring application. The IS16000 SNMP traps expose management data on the managed system in the areas of temperature senor, fans, power supplies, pools, and physical disks as well as a variety of real-time critical and error events.

Each Controller has an SNMP_AGENT. When changing the SNMP settings, you must set the changes on each Controller separately.

To set the IP address, enter the command (Figure 124):

`UI SET SNMP IP_ADDRESS=<ip_address> COMMUNITY="<name>"`

To show all the SNMP trap agent attributes, use **`UI SHOW SNMP`** command (Figure 124).

Figure 124. SNMP Configuration Example Screen

```
RAID[0]$ ui set snmp ip_address=192.168.0.10
SNMP Agent set with STATUS='Success' (0x0)

RAID[0]$ ui set snmp community=public
SNMP Agent set with STATUS='Success' (0x0)

RAID[0]$ ui show snmp

SNMP Trap Agent Attributes

  IP_ADDRESS=192.168.0.10
  COMMUNITY=public
```

3.7.5.3 Inquiry Items and Events

Figure 125, Figure 126, and Figure 127 below detail the additional inquiry items and events that will be trapped in the SNMP MIB and in the Email Agent.

Figure 125. SNMP Inquiry Objects

SNMP Inquiry Objects	
Item	**Returned Values**
Temperature Sensor	Number of temperature sensors, list of temperature sensors. For each sensor : ID, Enclosure ID, Enclosure position, Status (normal, warning, critical).
Fans	Number of fans in the system, list of fans. For each fan: ID, Enclosure ID, Enclosure position, Status (healthy, failure).
Power Supplies	Number of power supplies in the system, list of power supplies. For each power supply: ID, Enclosure ID, Enclosure position, Status (healthy, failure).
Pools	Number of pools, list of pools. For each pool: OID, type (storage, spare, unassigned), access, number of members.
Physical Disks	List of disks, for each disk: WWN, enclosure number, slot, status (normal, failed, predicted failure, unknown)

Figure 126. SNMP Only Events

SNMP ONLY Events	
Error	**Description**
LOG_AMPD_DEVICE_INIT_FAILED	A device failed to initialize.

Figure 126. SNMP Only Events

SNMP ONLY Events	
Error	**Description**
LOG_AMPD_DSK_DEVICE_INIT_FAILED	DISK device initialization failure.
LOG_AMPD_ICL_DEVICE_INIT_FAILED	ICL device initialization failure.
LOG_AMPD_MPI_DSK_SCSI_SENSE_DATA	Disk data received.
LOG_AMPD_MPI_ICL_SCSI_SENSE_DATA	ICL SCSI sense data was received.
LOG_AMPD_MPI_SCSI_SENSE_DATA_EVT	SCSI sense data was received.
LOG_AMPD_MPI_SES_SCSI_SENSE_DATA	SES SCSI sense data was received.
LOG_AMPD_SES_DEVICE_INIT_FAILED	SES device initialization failure.
LOG_ES_BATTERY_FAILURE_PREDICTED	UPS for the Controller is predicted to fail and should be replaced.
LOG_ES_COOLING_ELEMENT_INSERTED	A fan (located in the power supply) has been inserted.
LOG_ES_COOLING_ELEMENT_NORMAL	A fan (located in the power supply) has returned to normal status.
LOG_ES_COOLING_ELEMENT_REMOVED	A fan (located in the power supply) has been removed.
LOG_ES_COOLING_ELEMENT_WARNING	Enclosure cooling element has reported a warning condition.
LOG_ES_DISK_SLOT_ELEMENT_INSERTED	A device has been inserted into a disk slot.
LOG_ES_DISK_SLOT_ELEMENT_REMOVED	A device has been removed from a disk slot.
LOG_ES_ENCL_UPS_WARN_AC	UPS for the Controller has started with no AC power.
LOG_ES_EXPANDER_ELEMENT_INSERTED	A DEM or an IO module has been inserted.
LOG_ES_EXPANDER_ELEMENT_NORMAL	The enclosure expander element has reported a warning condition.
LOG_ES_EXPANDER_ELEMENT_REMOVED	A DEM or an I/O module has been removed.
LOG_ES_EXPANDER_ELEMENT_WARNING	The DEM reports normal.
LOG_ES_POWER_SUPPLY_INSERTED	A power supply has been inserted.
LOG_ES_POWER_SUPPLY_NORMAL	The power supply has returned to a normal status.
LOG_ES_POWER_SUPPLY_REMOVED	A power supply has been removed.
LOG_ES_POWER_SUPPLY_WARNING	Power supply for the Controller issued warning.
LOG_ES_TEMPERATURE_SENSOR_INSERTED	A temperature sensor (located in the power supply) has been inserted.
LOG_ES_TEMPERATURE_SENSOR_NORMAL	A temperature sensor (located in the power supply) has returned to normal status.
LOG_ES_TEMPERATURE_SENSOR_REMOVED	A temperature sensor (located in the power supply) has been removed.
LOG_ES_TEMPERATURE_SENSOR_WARNING	A temperature sensor in the enclosure has reported a warning condition.
LOG_ES_UPS_CLIENT_IS_RESPONSIVE	Controller enclosure UPS is responsive.

Figure 126. SNMP Only Events

SNMP ONLY Events	
Error	**Description**
LOG_RAID_UNCORRECTED_MEDIUM_ERR	Aphysical device medium error could not be corrected due to lack of pool redundancy.
LOG_RAID_UNCORRECTED_SILENT_ERR	A SATAssure error (silent error) could not be corrected due to lack of pool redundancy.
LOG_RT_SNMP_TRAP_EVENT	Event received.
LOG_ST_ MEMBER_CHANGE	A member of a storage pool has changed its status.

Figure 127. SNMP and Email Notification Events

SNMP and Email Notification Events	
Error	**Description**
LOG_AMPD_MPI_IOC_INIT_FAIL	SAS/SATA channel has failed to initialize.
LOG_ES_COOLING_ELEMENT_ERROR	Enclosure cooling element has reported an error condition.
LOG_ES_CTLR_PWR_SRC_CHANGED	Controller power source has changed.
LOG_ES_ENCL_UPS_WARN_BATT	UPS for the Controller indicates that its battery should be replaced.
LOG_ES_ENCL_UPS_WARN_INTF	UPS for the Controller indicates that its interface has failed. Check the cable.
LOG_ES_ENCL_UPS_WARN_UPS	UPS for the Controller indicates that it has failed.
LOG_ES_EXPANDER_ELEMENT_ERROR	A DEM or an I/O module has reported an error condition.
LOG_ES_POWER_SUPPLY_ERROR	The power supply for one of the enclosure within the subsystem reported error condition.
LOG_ES_SET_BATTERY_CHARGER_FAILED	Attempt to set Controller battery charger current and voltage failed.
LOG_ES_TEMPERATURE_SENSOR_ERROR	The enclosure temperature sensor has reported an error condition.
LOG_ES_UPS_CLIENT_IS_NOT_RESPONSIVE	Controller enclosure UPS is not responsive. Check the USB cable between the UPS and Controller.
LOG_ST_HEARTBEAT_MISMATCH	Heartbeat mismatch reboot.
LOG_ST_MIR_STATE	Manual Intervention Required state.
LOG_ST_MIRROR_JOIN_FAIL	Other Controller unable to flush mirror data status.
LOG_ST_NO_CONFIG_READ	Unable to read the configuration from backend drives status.
LOG_ST_NO_CONFIG_WRITE	Unable to write the configuration to backend drives status.
LOG_ST_POOL_CHANGE	Pool state changed.
LOG_ST_SET_AWL	Auto Write Lock condition.
LOG_ST_SET_CRITICAL	Critical condition state.
LOG_ST_SET_FAILED	Failed condition.

Figure 127. SNMP and Email Notification Events

SNMP and Email Notification Events	
Error	**Description**
LOG_ST_SPLIT_BRAIN	The Controllers have lost communications with each other and are operating in split-brain mode.
LOG_SYS_STARTUP	The system has been restarted.

CHAPTER 4

Troubleshooting

4.1 Troubleshooting

This section describes common problems, with possible solutions, which can occur with the IS16000 system.

4.1.1 *"RAID[0]$" Prompt is not shown correctly*

The default CLUI prompt is "RAID[0]$". If "CLUI$" is shown instead of "RAID[0]$", one of the following problems may have occurred:

- The firmware on this Controller is in a "MIR" state and the MIR state needs to be resolved.
- The CLUI has lost communication with the firmware or the firmware is not running. A restart of the system is needed.
- There are multiple CLUI connections to the firmware already and there are not enough resources to allow another connection. Only one CLUI connection is currently supported.

4.1.2 *Event Log*

Some types of issues can be seen from the event log:

- Background jobs affecting performance

```
014547 2009-08-31 12:30:08:1953906  G=62 S=0 T=1 RP=0 VP=1 LOG_ST_MEMBER_CHANGE
STATE POOL MEMBER CHANGE INDEX:00000000 ID:0x5000cca20def4516 PDIDI:0085 MEMBER_INDEX:0000
STATE:REBUILD
```

- Backend issues affecting performance

```
000425 2009-04-28 14:39:12:6963749  G=10 S=0 T=1 RP=0 VP=5
LOG_AMPD_MPI_SAS_DEVICE_DISAPPEARED   AMPD IOC0 WWN:5000cca216ed8430
ES:50001ff101ed0000:60. RC:04 Hndl:0065 PrntId:50001ff101ed017f PHY:08
BT:00:14 DI:00000081 DSK:0000

000430 2009-04-28 14:39:13:8133330  G=10 S=0 T=1 RP=0 VP=5
LOG_AMPD_MPI_SAS_DEVICE_APPEARED   AMPD IOC0 WWN:5000cca216ed8430
ES:50001ff101ed0000:60. RC:03 Hndl:0065 PrntId:50001ff101ed017f PHY:08
BT:00:14 DI:00000081 DSK:0000
```

- ICL problem - link bounce

```
000468 2009-08-25 22:04:20:7729086  G=60 S=0 T=1 RP=0 VP=5 LOG_AMPD_ICL_LINK_STATUS_CHANGE
AMPD Ctx:2f0006c24528 Drv:2f0006c2ed80 Old:DWN New:UP
000469 2009-08-25 22:04:20:7729497  G=60 S=0 T=1 RP=0 VP=1 LOG_DUCK_CONTROLLER_CONNECTED
DUCK Controller Connected.
```

- ICL problem - Controller failure

```
008418 2009-07-29 08:16:50:9043234 G=60 S=0 T=1 RP=0 VP=1 LOG_DUCK_CONTROLLER_DISCONNECTED
DUCK Controller Disconnected.
008419 2009-07-29 08:16:50:9043363  G=60 S=0 T=1 RP=0 VP=1 LOG_ST_OTHER_DIED
STATE OTHER CONTROLLER DIED

RAID[0]$ show controller *
OID: 0x38000000 Index: 0000 Name: A LOCAL PRIMARY
Total Controllers: 1
```

4.2 Manual Intervention Required (MIR) States

Manual Intervention Required (MIR) represents a condition with the Controller that requires the user to provide a solution before proceeding with normal Controller operations. This is to guard against the Controller firmware from executing operations that may not necessarily be the desired operation of the user. These conditions will most likely be seen in a new system installation environment. For example, when a system is booted and the backend physical disks have never been installed behind the Controllers, the Controller firmware has never had a chance to write out configuration metadata. The Controller recognizes that there is no valid metadata and requires the user to acknowledge proceeding or not.

Use the **SHOW CONTROLLER LOCAL ALL_ATTRIBUTES** command to display the MIR condition of the Controller you are logged into (Figure 128).

Figure 128. Show Controller Information Screen

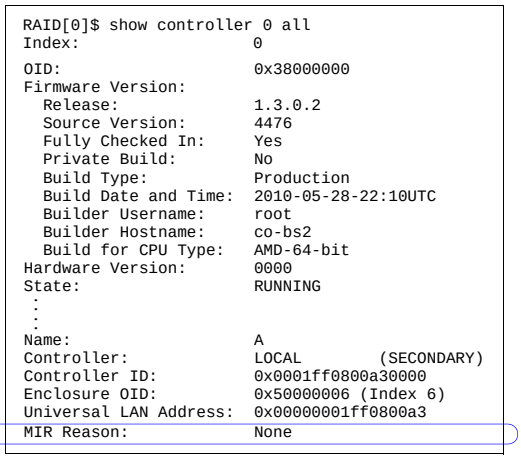

```
RAID[0]$ show controller 0 all
Index:                  0
OID:                    0x38000000
Firmware Version:
  Release:              1.3.0.2
  Source Version:       4476
  Fully Checked In:     Yes
  Private Build:        No
  Build Type:           Production
  Build Date and Time:  2010-05-28-22:10UTC
  Builder Username:     root
  Builder Hostname:     co-bs2
  Build for CPU Type:   AMD-64-bit
Hardware Version:       0000
State:                  RUNNING
 .
 .
 .
Name:                   A
Controller:             LOCAL        (SECONDARY)
Controller ID:          0x0001ff0800a30000
Enclosure OID:          0x50000006 (Index 6)
Universal LAN Address:  0x00000001ff0800a3
MIR Reason:             None
```

Listed below are the possible MIR conditions and their required actions.

MIR_JIS_DISCOVERY_IN_PROG

Initialized Storage discovery is in progress. Please allow time for configuration discovery to complete. If this condition persists, please reboot the system. If this condition continues to persist after the reboot, please contact customer support.

MIR_OTHER_JIS_DISCOVERY_IN_PROG

Initialized Storage discovery is in progress on other controller. Please allow time for configuration discovery to complete. If this condition persists, please reboot the system. If this condition continues to persist after the reboot, please contact customer support.

MIR_NO_BACKEND_DRIVES

This controller cannot find any disk modules on the backend. Install disk modules on the backend or fix the condition that prevents this controller from finding backend disks. Please refer to the installation guidelines for proper setup.

MIR_NO_CONFIG

A configuration could not be created on the backend disks. Fix the condition that prevents this controller from creating a configuration on the backend disks. Please refer to the installation guidelines for proper setup.

MIR_NO_QUORUM

No quorum disks could be found within the disk modules on the backend. Fix the condition that prevents the quorum disks from being seen. This may be caused by the disks never being in the storage array. This will normally happen on new system installations. Use the "**CLEAR SUBSYSTEM MIR_STATE**" command to create an empty configuration. ***NOTE**: This creates a new configuration and the old configuration will be deleted if existed.

MIR_NOT_LAST_CONTROLLER

This controller found a valid configuration, but was not present when another controller owned the configuration. Another controller may have cached data for this configuration. This may happen when controllers are swapped out or if the controllers went down and restarted individually. To use the found configuration on this controller, use the "**CLEAR SUBSYSTEM MIR_STATE**" command.

MIR_MULTIPLE_JIS

Multiple configurations were found on the quorum disks. This may happen if the disks from one system were installed in another system when the systems were powered down and then rebooted. If the disks were installed in the system while running, then this should not be an issue. A list of found configurations will be listed. Use the "**CLEAR SUBSYSTEM MIR_STATE ID=<id>**" command to use the specified ID's configuration.

MIR_DUAL_NO_AGREE

The two connected controllers do not agree on the ID of the configuration. This may occur if one controller saw a subset of the disks and the other controller saw a different subset of disks. Please refer to the installation guidelines for proper setup.

MIR_CONFIG_MISMATCH

The configuration version of this firmware does not match the configuration version of that on media. To proceed, either reload the previous version of firmware and do a backup then upgrade, or delete your configuration to continue. Use "**CLEAR SUBSYSTEM CONFIGURATION**" to create an empty configuration. ***NOTE**: This creates a new configuration and the old configuration will be deleted if existed.

MIR_NO_LOAD_CONFIG

A configuration could not be loaded from the backend disks. Fix the condition that prevents this controller from loading a configuration from the backend disks, or use "**CLEAR SUBSYSTEM MIR_STATE**" to create an empty configuration. Please refer to the installation guidelines for proper setup. ***NOTE**: This creates a new configuration and the old configuration will be deleted if existed.

4.3 Recovery from Disk Failure

When a disk failure occurs on the IS16000, the Storage Pool containing that disk will begin operating in degraded mode. This means that the Storage Pool will continue to handle I/O commands from the host, but there will be no redundancy to protect against additional disk failures on the same Storage Pool. If another disk module fails on the same storage pool before the data on the first disk is rebuilt (to a replacement disk or hot spare), the Storage Pool will go offline.

If one disk in a Storage Pool fails, the data or parity information on the failed disk will be reconstructed from the parity disk and data disks of that Storage Pool.

4.3.1 Automatic Rebuild

A single disk failure in any Storage Pool does not result in data loss. The Virtual Disk(s) on that Storage Pool will continue to operate in degraded mode. If a spare disk is available and automatic rebuild is enabled, the IS16000 will automatically rebuild the data to a spare disk.

NOTE : System performance will be impacted while recovery is taking place.

When a disk failure occurs, the failure is written to the event log. You may monitor the rebuild progress or adjust the rebuild rate to match the user load requirements.

To display the event log (Figure 129), use the **SHOW CONTROLLER LOCAL LOG** command.

Figure 129. Display Event Log

```
RAID[0]$ show controller local log

000041 2009-02-04 15:03:17:7295774  G=4 S=0 T=1 RP=0 VP=1
LOG_ST_POOL_CHANGE  STATE POOL    CHANGE   POOL:0003 STATE:0005

000042 2009-02-04 15:03:17:7295780  G=4 S=0 T=1 RP=0 VP=1
LOG_ST_MEMBER_CHANGE  STATE MEMBER  CHANGE   ID:5000cca215c56e02 PDIDI:0145
POOL:0003 INDEX:0000 STATE:0003

000043 2009-02-04 15:03:17:7295810  G=4 S=0 T=1 RP=0 VP=1
LOG_ST_REBUILD_START  STATE REBUILD START    ID:5000cca215c56e02 PDIDI:0145
POOL:0003 INDEX:0000 FENCE:0000000000000000
```

State Change on Pool:0003

Disk Replacement

Rebuild Initiated on Spare Disk

To look at the failed disk (Figure 130), enter command **SHOW UNASSIGNED_POOL FAILED ALL**

Figure 130. Display Failed Disk

```
RAID[0]$ show unassigned_pool failed all
  OID:               0x20a7003f
  Pool OID:          UNASSIGNED
  Capacity:          704512 MBs (0x56000000 blocks)
  Raw Capacity:      715404 MBs (0x575466f0 blocks)
  Block Size:        512
  Enabled Disk Ch:   0x27  0x22
  Disk Slot:         1:42
  Vendor ID:         Hitachi
  Product ID:        Hitachi HUA721075KLA330
  Product Revision:  GK8OAB0A
  Serial Number:           GTF200P8GBVPXF
  Health State:      FAILED
  Rotation Speed:    7200 RPM
  Device Type:       SATA
  Member State:      UNASSIGNED
  Spare:             FALSE
  Failed:            TRUE
  UUID:              0x5000cca215c564560
```

Physical Location
Enclosure 1 Disk Slot 42

To monitor the rebuild progress (Figure 131), enter command **SHOW JOB * ALL_ATTRIBUTES**

Figure 131. Display Rebuild Progress

```
RAID[0]$ show job * all_attributes

  OID:              0x2b050003
  Target:           0x19b40003
  Type:             REBUILD
  Status:           RUNNING
  Priority:         80
  Fraction Complete:12%
```

To show the Storage Pool information (Figure 132), enter command **SHOW POOL <id> ALL_ATTRIBUTES**. Once the rebuild is complete, the status of Storage Pool will return to "NORMAL".

Figure 132. Display Storage Pool Information

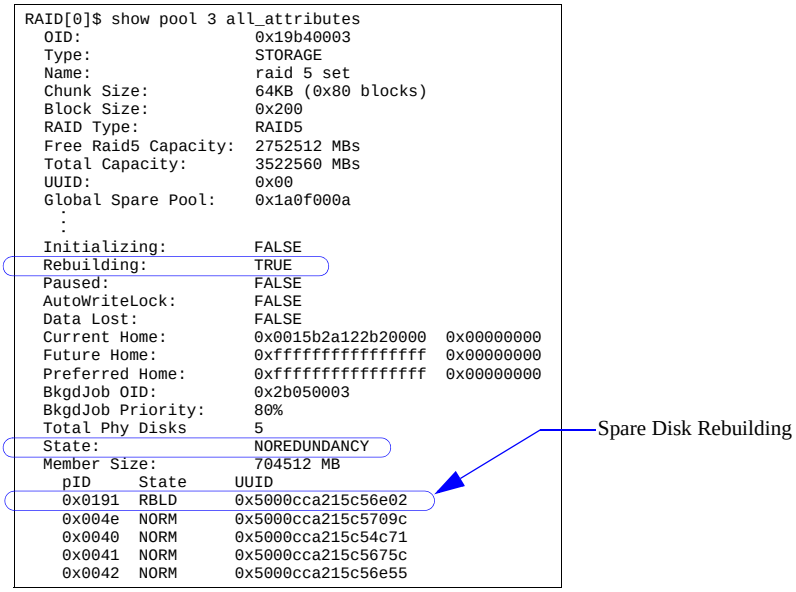

4.3.2 *Disk Module Replacement*

Upon completion of a rebuild, the spare disk becomes a member of the Storage Pool, replacing the failed disk. After you have replaced the failed disk with a new disk, the new disk is added to the "Unassigned Pool". It is recommended that you assign this new disk to the Spare Pool to "replace" the spare disk that has been used.

4.3.3 *When a Spare is not available*

When a disk is failed by the system and there is no spare disk available, you need to replace the failed disk immediately. After you have replaced the failed disk with a new disk, you can initiate a rebuild as described below:

1. Enter command **SHOW UNASSIGNED_POOL * ALL** to identify the new disk's index name.
2. Enter command **ASSIGN PHYSICAL_DISK <new-disk-id> TO_POOL <pool-id> SET_SPARE** where <new-disk-id> is the index name of the replacement disk and <pool-id> is the OID of the Storage Pool that had the failed disk.

4.3.4 Manual Rebuild

You may manually replace a failed disk using the **REPLACE** command:

```
REPLACE PHYSICAL_DISK <id> NEW_DISK <new-disk-id>
```

A Replace operation is used to replace a failed disk with a healthy spare disk. The operation can take several hours to complete depending on the size of the disk and speed of the replace operation. The rate of rebuild can be adjusted (see Section 3.3.3, "Rebuild Policy Priority" on page 78 for more information).

4.4 Controller Component Failure

The Controller Enclosure Fault LED turns amber when a fan failure, a power failure, or over temperature condition occurs (Figure 133).

Figure 133. Controller Enclosure Fault LED Indicator

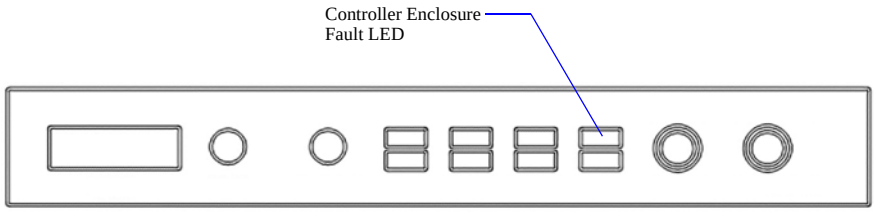

A single component failure, therefore, will not shut down the system. However, in the unlikely event of component failure, you can replace the failed component while the IS16000 is running. The replaced component will automatically be returned to service once the component has been installed and booted up.

Warning

If the IS16000 is powered up and you remove any module, replace it immediately. If the system is used with modules missing for more than a few minutes, the system can overheat, causing power failure and data loss. Such use will invalidate the warranty.

Warning

Observe all conventional ESD precautions when handling IS16000 modules and components. Avoid contact with backplane components and module connectors.

4.4.1 Replacing a Power Supply Module

If a Controller power supply fails, the Enclosure Fault LED on the front panel will turn **amber** and the green LED on the failed power supply will be **off**.

You can also determine if a power supply has failed using the CLUI command, **`SHOW POWER`**. The position of the failed power supply should match the label (PSU1 or PSU2) on the back of the Controller.

If the power supply has failed, you must replace it. Once you remove it, you must replace it within 5 minutes to prevent the system from over-heating.

NOTE : Obtain a replacement module of the same type before removing any faulty module.

Follow these steps to replace a power supply module:

1. Identify the failed power supply.
2. Disconnect its power cord.
3. Loosen the module's thumbscrew, and then slide the module out of the bay (Figure 134).

Figure 134. Controller Power Supply Module Removal

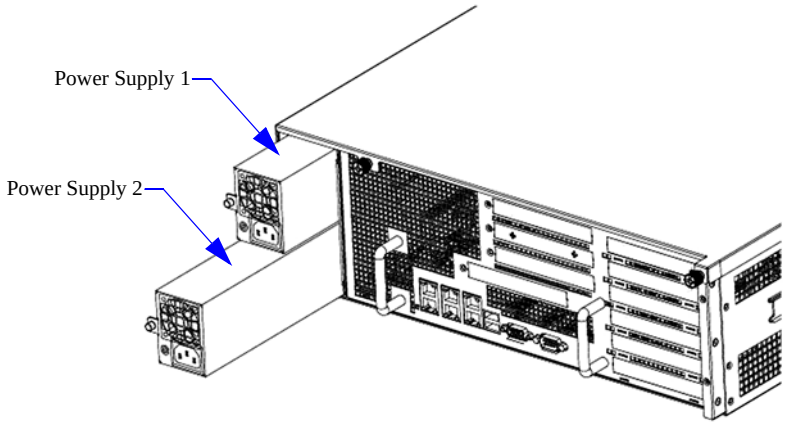

4. Slide the replacement module into the bay, making sure that it is fully inserted.
5. Tighten the thumbscrew to secure it.
6. Connect the power cord.
7. Verify that the Status LED is **green**, indicating that the module is operating normally.
8. Verify that the Enclosure Fault LED on the front panel is no longer amber.

4.4.2 Replacing a Fan

If there is any problem with the fan module, the Enclosure Fault LED on the front panel will turn **amber.**

You can determine if a fan has failed using the CLUI command, **SHOW FAN**. The position of the failed fan should match the label (FAN1, FAN2, FAN3 or FAN4) on the front of the enclosure.

If a fan has failed, you must replace it. Once you remove it, you must replace it within 5 minutes to prevent the system from over-heating.

NOTE : Obtain a replacement module of the same type before removing any faulty module.

Follow these steps to replace a fan module:

NOTE : Attempting to replace more than one fan simultaneously will cause the Controller to fail.

---

1. Identify the failed fan.
2. Loosen the fan's thumbscrew (Figure 135).
3. Slide the module out of the bay.

Figure 135. Controller Fan Removal

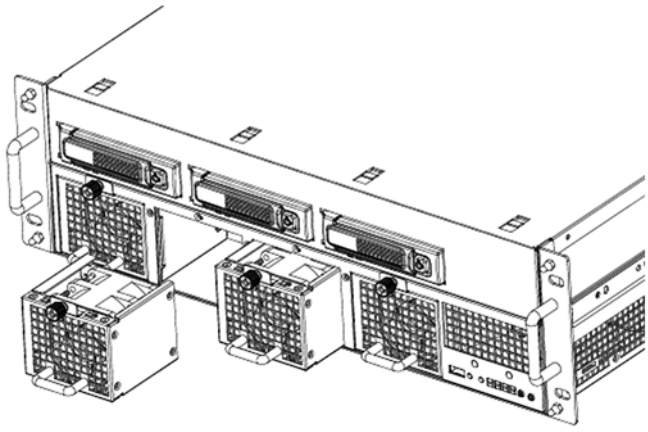

4. Slide the new fan into the bay, making sure that it is fully inserted.
5. Tighten the thumbscrew to secure it.
6. Verify that the Enclosure Fault LED is no longer amber, indicating that the new module is operating normally.

4.4.3 Replacing Internal Disk Module

Follow these steps to replace an internal disk module at the front of the Controller:

1. Disengage the lock using the supplied key.
2. Press the handle release button to release the handle (Figure 136).
3. Pull the module out of the bay.

Figure 136. Controller Internal Disk Module Replacement

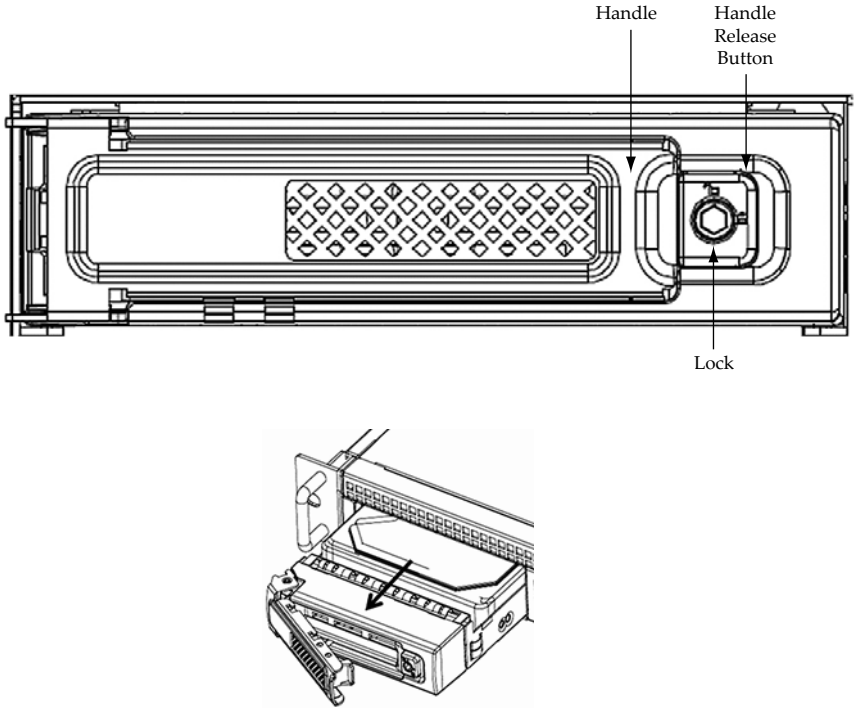

4. Slide the new module into the bay, making sure that it is fully inserted.
5. Close the handle and push the handle until it clicks, indicating that the handle is latched.

4.5 Disk Enclosure Faults

4.5.1 Front Panel Indicators

Green LEDs are always used for good or positive indication. Amber LEDs indicate there is a critical fault present within the module. The front panel indicators show the aggregated status of all the modules (Figure 137). The LEDs are defined in Figure 138. The Disk Activity LEDs indicate disk presence and flash during data I/O.

Figure 137. Disk Enclosure Front Panel LED Indicators

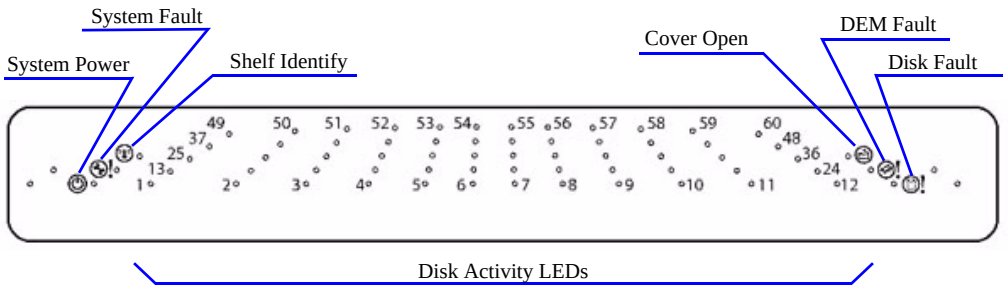

Figure 138. Front Panel LED Description

LED	Definition	Color	Status
	System Power	Green	ON - DC power is present OFF - DC power is not present LED does NOT flash under normal operating conditions
	System Fault	Amber	ON - one or more components within enclosure have failed. A service action is required. Exact failed component has its own amber fault LED lit. OFF - no detectable faults
	Shelf Identify	Blue	Flashing - receiving “Locate Enclosure” command OFF - NOT receiving “Locate Enclosure” command
	Cover Open	Amber	OFF - both cover pieces securely closed and latched in place ON - either of the cover pieces is NOT securely closed and latched in place
	DEM Fault	Amber	OFF - all DEMs operating correctly ON - at least one DEM has failed; service action required
	Disk Fault	Amber	ON - one or more HDDs are failed; SES must determine exact HDD OFF - no detectable disk faults
Individually numbered	Disk Activity	Green	ON - SAS HDD is present Flashing - indicates HDD activity OFF - no HDD activity

4.5.2 Thermal Alarm

Symptom	Cause	Action
1. Enclosure System FAULT LED is amber. **2**. An amber LED on one or more Power Cooling Modules. **3**. Air temperature exiting PSU is above 55°C.	If the internal temperature measured in the airflow through the enclosure exceeds a pre-set threshold a thermal alarm will sound.	**1**. Check local ambient environment temperature is below the upper 40°C specification. **2**. Check the installation for any airflow restrictions at the front and rear of the enclosure. A minimum gap of 1" (25mm) at the front and 2" (50mm) at the rear is recommended. **3**. Check for restrictions due to dust build-up, clean as appropriate. **4**. Check for excessive re-circulation of heated air from rear to the front, use in a fully enclosed rack installation is not recommended. **5**. If possible shutdown the enclosure and investigate the problem before continuing.

Warning

An Enclosure will shut down when a critical temperature threshold is exceeded in order to prevent permanent damage to the disks.

4.5.3 System Faults

Symptom	Cause	Action
1. The System Fault LED turns amber on front panel.	The ESI (Enclosure Services Interface) processor has detected an internal fault (e.g. failure of an internal communications path)	**1.** Check for other AMBER LED indications on the Power Cooling modules. If there is a PSU error present there may be a communications problem with that Power Cooling module. Remove and then re-fit the module, if the problem persists then replace the module. **2.** Check for other AMBER LED indications on the disk modules. If none are evident then there may either be an ESI processor problem or a backplane problem.

4.5.4 Power Supply/Cooling Faults

Symptom	Cause	Action
1. System Fault LED is amber on front panel. **2.** An amber LED on one or more Power Cooling modules.	**1.** Any power fault. **2.** A fan failure. **3.** A thermal condition which could cause PSU overheating.	**1.** Check power On/Off switch on rear of Power Cooling module is turned ON. **2.** Check AC Mains connections to Power Cooling module is live. **3**. Disconnect the Power Cooling module from mains power and remove the module from the system, re-install module. If problem persists, replace Power Cooling Module. **4**. Reduce the ambient temperature.

Figure 139 describes the PCM LED status.

- Under normal conditions, the LEDs should be illuminated constant **green**.
- If a problem is detected, the color of the relevant LED will change to **amber**.

Figure 139. Power Cooling Module LEDs

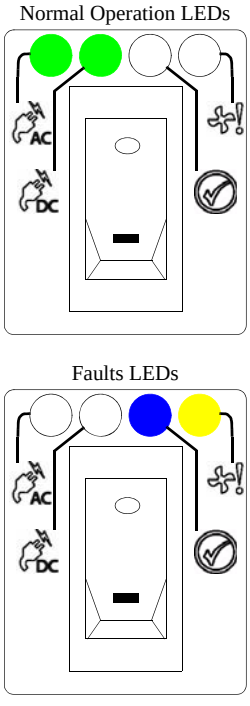

LED	Description	Color	Status
Enclosure front	Enclosure fault	Amber	ON - AC input, DC output, fan or other PCM fault detected
AC	PCM AC	Green	ON - AC input to PCM within tolerances OFF - PCM failed
DC	PCM DC	Green	ON - DC output of PCM within tolerances OFF - PCM failed
	PCM Fault	Amber	ON - AC input, DC output, fan or other PCM fault detected OFF - no fault detected
	PCM ID	Blue	Flashing - this module is receiving “Locate Power Supply” or “Locate Fan” command OFF - NOT receiving “Locate Power Supply” or “Locate Fan” command

If a power supply or fan fails, the PCM Fault LED on the module will turn **amber.**

If a power supply unit or its fan is faulty, you must replace the whole PCM. You must not take any longer than 5 minutes to replace this module to prevent the enclosure from over-heating.

NOTE : Obtain a replacement module of the same type before removing any faulty module.

To replace a module:

1. Identify the failed PCM.
2. Turn off the module's power switch and disconnect the power cord (Figure 140).
3. Loosen the module's thumbscrew. Then slide the module out of the bay.

Figure 140. Power Supply Module

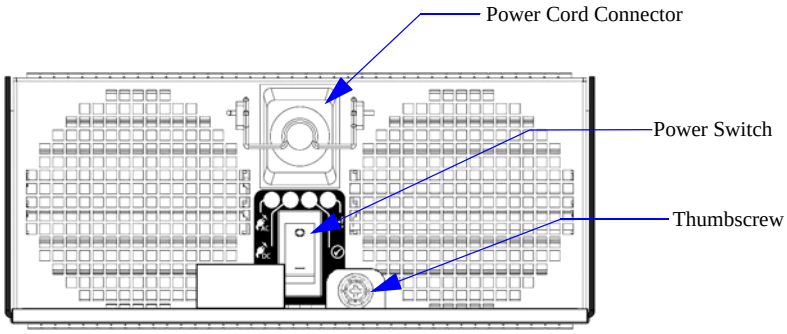

4. Ensure that the power switch on the new (replacement) module is OFF.
5. Slide the module into the bay, making sure that it is fully inserted.
6. Tighten the thumbscrew to secure it.
7. Connect the power cord.
8. Turn on the power switch.
9. Verify that the AC and DC LEDs are **green**, indicating that the module is operating normally.

4.5.5 Replacing a Disk Module

Figure 141 describes the fault LED status when an HDD has failed.

- Under normal conditions, the LEDs should be illuminated constant **green**.
- If a problem is detected, the color of the relevant LED will change to **amber**.

Figure 141. HDD Fault LED Indicators

Identifier	Location	Color	Status
Enclosure fault	Enclosure front	Amber	ON - disk fault
Disk fault	Enclosure front	Amber	ON - disk fault

NOTE : When a disk is failed by the IS16000, replace it promptly so that the operation and performance of the system are not affected.

Wear an ESD wrist strap or otherwise ground yourself when handling the disk modules and components. Electrostatic discharge can damage the circuit boards.

To remove a module:

1. Slide the enclosure out from the rack by squeezing the tabs on both rack slides. Keep pulling until the enclosure locks and you hear a clicking sound.
2. Disengage both enclosure cover latches (Figure 142) and open the covers.

Figure 142. Disk Enclosure Cover Latch

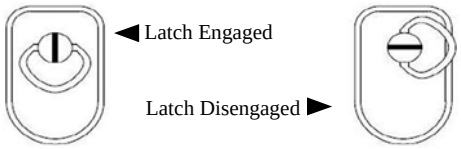

3. Open the cover.

NOTE : Opening the cover affects the system cooling. Please limit service time to 5 minutes. Reminder beeps will occur occasionally. A continuous alarm will sound after 5 minutes.

4. Locate the disk module that has the blue LED.

5. On that module, slide the latch backward to release the handle (Figure 143).

Figure 143. Release Disk Module Handle

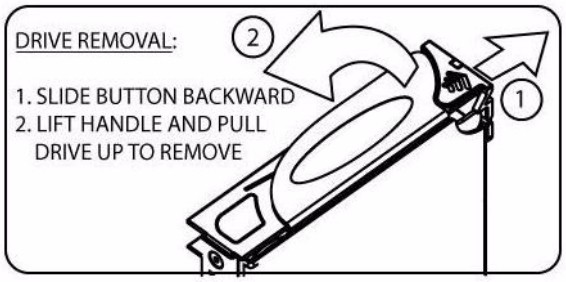

6. Lift the handle and pull the module up, just enough to disconnect the module from the backplane.

7. Wait for 30 seconds for the disk to completely spin down.

CAUTION : Disk spin down--Damage can occur to a disk if it is removed while still spinning.

---

8. Pull the module gently out of the disk bay.

To insert a module:

1. Slide the latch backward to release the handle.

2. Insert the replacement module into the disk bay. Note orientation of the module.

3. Cam the disk module home. The camming foot on the base of the module will engage into the slot in the enclosure.

4. When the module is fully inserted, close the handle. You should hear a click as the latch engages and holds the handle closed.

5. Close the enclosure covers and engage both cover locks.

NOTE : Both cover locks must be engaged before pushing the enclosure into the rack. Otherwise, open cover alarm will sound.

---

6. Squeeze the tabs on both rack slides and push the enclosure back into the rack.

4.5.6 Replacing a DEM (Drive Expander Module)

NOTE : The DEM card should only be replaced by trained personnel.

Figure 144 describes the DEM LED status.

- Under normal conditions, the LEDs should be illuminated constant **green**.
- If a problem is detected, the color of the relevant LED will change to **amber**.

Figure 144. DEM LED Indicators

Identifier	Location	Color	Status
Enclosure fault	Front enclosure	Amber	ON - DEM fault detected
DEM fault	Front enclosure	Amber	ON - DEM fault detected
DEM fault	Internal (DEM) enclosure	Amber	ON - DEM fault detected
DC OK	Internal (DEM) enclosure	Green	ON - DC voltage regulation is within limits OFF - DC output failure
Expander MIPS ready	Internal (DEM) enclosure	Green	ON - expander internal processor is booted and operating correctly OFF - expander internal processor is NOT booted or ready
DEM identify	Internal (DEM) enclosure	Blue	Flashing - receiving “Locate Expander” command OFF - NOT receiving “Locate Expander” command

A DEM failure can also be seen in the event log.

```
LOG_ES_EXPANDER_ELEMENT_ERROR     ES Enclosure 0x50001ff104521000
expander DEM1A was reported error condition. Status = 2, DISAB = 0, PF = 0.

000182 2009-10-28 15:21:58:6976333  G=62 S=1 T=1 RP=0 VP=1
LOG_ES_EXPANDER_ELEMENT_WARNING     ES Enclosure 0x50001ff104521000
expander DEM4B was reported warning condition. Status = 3, DISAB = 0, PF = 0.
```

Figure 145 shows the DEM locations in the disk enclosure.

Figure 145. DEM Locations

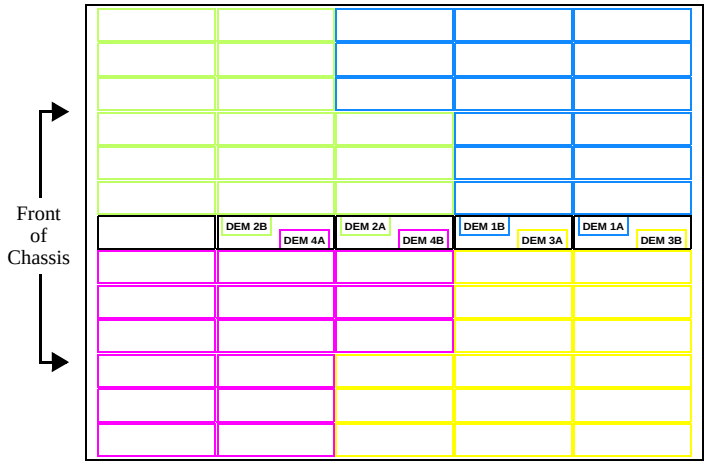

To replace a DEM:

Warning **Observe all conventional ESD precautions. Avoid contact with backplane components and module connectors.**

1. Slide the enclosure out from the rack by squeezing the tabs on both rack slides. Keep pulling until the enclosure locks and you hear a clicking sound.
2. Disengage both enclosure cover latches and open the covers (Figure 146).

Figure 146. Disk Enclosure Cover Latch

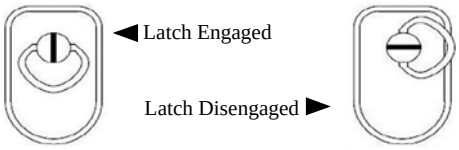

NOTE : Opening the cover affects the system cooling. Please limit service time to 5 minutes. Reminder beeps will occur occasionally. A continuous alarm will sound after 5 minutes.

3. Locate the DEM that needs to be replaced. The handle rib will light up **blue** for service.
4. Grip the top latch, rotate it until you reach the limit of rotation, and then pull up to remove the DEM (Figure 147).

Figure 147. DEM Removal

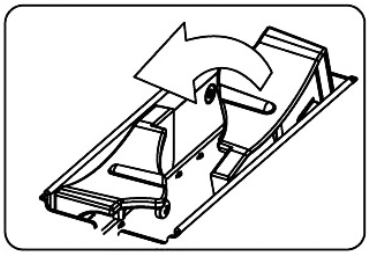

5. Insert the replacement DEM.
6. Press on the latch to rotate it down and lock it into place.
7. Close the enclosure covers and engage both cover locks.

NOTE : Both cover locks must be engaged before pushing the enclosure into the rack. Otherwise, open cover alarm will sound.

8. Squeeze the tabs on both rack slides and push the enclosure back into the rack.

4.5.7 Replacing an I/O Module

Warning

Do not remove this module unless a replacement can be immediately added. The system must not be run without all modules in place.

Figure 148 describes the Controller LED status.

- Under normal conditions, the LEDs should be illuminated constant **green**.
- If a problem is detected, the color of the relevant LED will change to **amber**.

Figure 148. I/O Module LED Indicators

Identifier	Location	Color	Status
Enclosure fault	Front enclosure	Amber	ON - I//O fault detected
I/O OK	Rear enclosure	Green	OFF - an I/O module detectable fault is present. Independent of SAS link fault condition and stays lit during a SAS link fault.
I/O Fault	Rear enclosure	Amber	ON - an I/O module detectable fault is present. Independent of SAS link fault condition and stays lit during a SAS link fault.
SAS Link Activity	Rear enclosure	Green	ON - a valid link is present on any of the 4 links of the port. OFF - none of the 4 links of the port have a valid connection.
SAS Link Fault	Rear enclosure	Amber	ON - a valid link is present on any of the 4 links of the port. OFF - none of the 4 links of the port have a valid connection.

To remove an I/O module:

1. Remove the cables from the I/O module.
2. Release the two handles on the bottom of the module by simply pulling each handle out and away from the module.
3. Pull both handles forward to cam the module out of the enclosure.
4. Grip the module securely and slide it out of the bay.

To insert a new module:

1. With the handles in the open position, slide the I/O module into the enclosure until the handles engage automatically.
2. Cam the module home by manually closing the handles. A click should be heard as the handles engage.
3. Connect the cables to the module.

CHAPTER 5

GUI Management Agent

This chapter provides information on using the IS16000 GUI (Graphic User Interface) Management Agent.

NOTE : The configuration examples provided here represent only a general guideline. These examples should not be used directly to configure your particular IS16000.

5.1 Starting the GUI Management Agent

5.1.1 Login

Using a web browser, open a link to the IP address of the IS16000 Controller. The management system supports Mozilla FireFox (version 3.0.11 and above) as well as Microsoft Internet Explorer (version 8.0 and above).

At the login screen, enter the user name **admin** with the password **password** to log into the system. User name and passwords are case-sensitive.

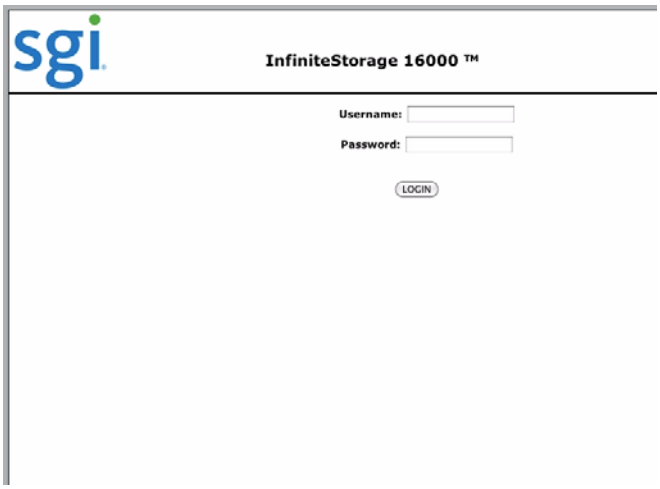

5.1.2 Initial Home Screen and Health Indicators

The home screen shows the Health Indicator. Figure 149 shows a newly configured system, the Health Indicator is, therefore, orange. An alert message also appears below the Health Indicator indicating that manual intervention is required for both Controllers.

Figure 149. Home Screen

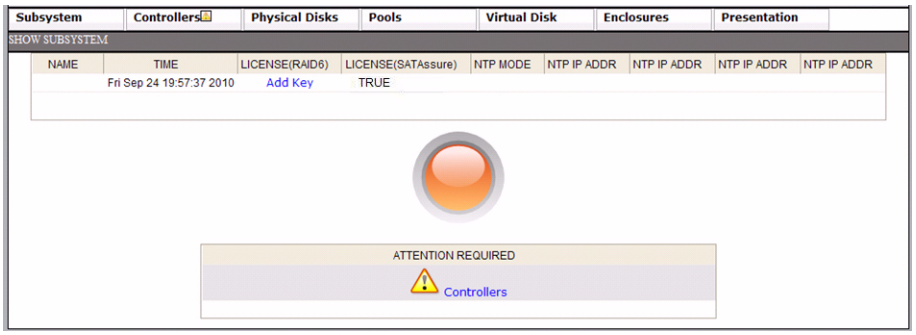

Warning

For any future encounters with MIR state issues, please refer to the Service Manual for specific instructions on how to clear them.

1. Select the Controller name. The system displays the **Controller Show All** screen (Figure 150).

Figure 150. Controller Show All Screen

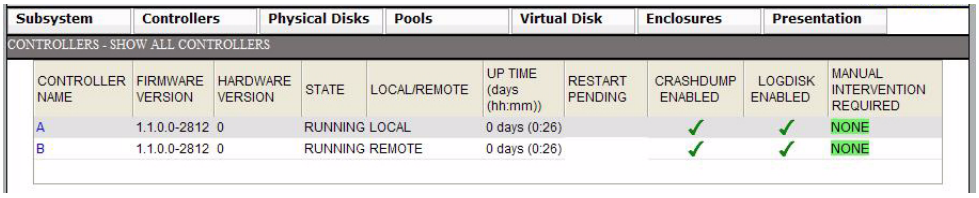

2. Select the underlined Manual Intervention Required. Details of the MIR state displays, “STATUS=MIR: no quorum disks found”.

3. Click **CLEAR MIR** to clear the MIR state. Your subsystem is now ready to be configured.

5.2 Physical Disks

5.2.1 Physical Disks Menu

The following commands are located under the Physical Disk pull down menu:

- Show Physical Disks
- Show Failed Disks
- Locate Disks
- Set Failed
- Assign to Pool
- Clear Failed

NOTE : If asterisks appear in the Physical Disk column when you attempt to view the advanced disk information or attempt to locate a disk, this shows that only one Controller sees the disk. Therefore, a Controller is down or another hardware issue exists somewhere.

5.2.2 Checking Status of Physical Disks

Before creating any Storage Pools, check the status of all the disks.

From the Physical Disks menu, select **Show All Physical Disks** to verify that all the disks are present and healthy (Figure 151).

Figure 151. Show Physical Disks Screen

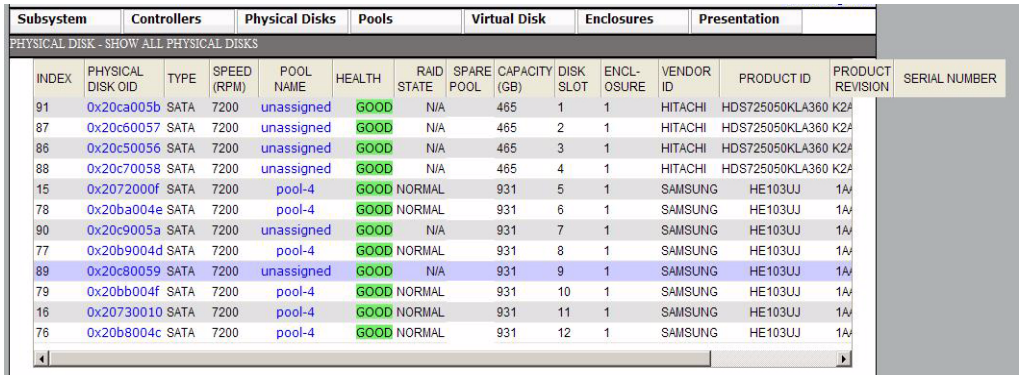

NOTE : If the disks are not visible, select **Subsystem Restart** from the Subsystem Menu to reboot the system.

5.3 Storage Pools

When creating a Storage Pool on a IS16000, the following selectable attributes are available:

- RAID Level
  Storage Pools can be configured to use either a RAID1, RAID5 or RAID6 parity scheme. In RAID1, the capacity of one disk is used for data duplication. In RAID5, the capacity of one disk is reserved for parity, allowing data recovery in the event of a single disk loss in the Storage Pool. In RAID6, the capacity of two disks is reserved for parity, allowing data recovery if either one or two disks are down in a Storage Pool. ***For maximum data protection, SGI recommends the use of RAID6***.

- Chunk Size
  The chunk size (in KiB blocks) defines the amount of data written to a single disk before proceeding to the next disk in the Storage Pool.

NOTE :RAID1 is a two member RAID set where the data is mirrored on each disk. There is no parity, hence, the chunk size is fixed.

- Disk Count
  A RAID1 Storage Pool may consists of 2 physical disks. A RAID5 Storage Pool may consist of 5 or 9 physical disks. A RAID6 Storage Pool may consist of 6 or 10 physical disks. For maximum performance, select disks with the same characteristics (such as SAS/SATA, capacity, and RPM).

- Drive Type
  Drive types can be SAS or SATA.

- Drive Size
  Drive size is the capacity of the disk.

- Spindle Speed
  You may choose (RPM): 15000, 10000, 7200, 5400, 0 or SSD.

- SATAssure
  SATAssure technology is designed to improve the reliability of enterprise SATA disks and make sure that data integrity is always mentioned for all I/O operations.

5.3.1 Storage Pool Menu

The following commands are located under the Pools pull down menu (Figure 152):

- Show Storage Pools
- Create Storage Pool
- Set Storage Pool Attributes
- Delete Storage Pool
- Locate Storage Pool

Figure 152. Pools Menu

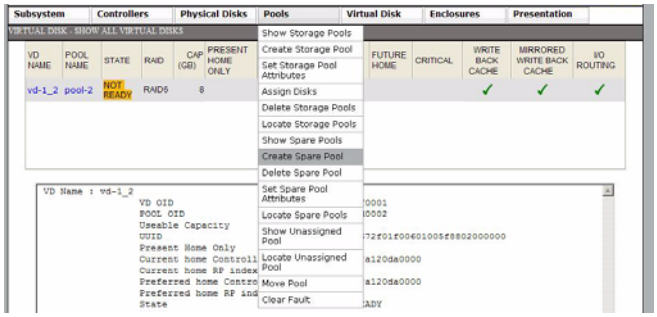

5.3.2 Create a Storage Pool

1. Select Create Storage Pool from the Pools Menu (Figure 153).

2. At the Create Storage Pool screen (Figure 153), select all the attributes for the pool:

NOTE :RAID1 does not accept a chunk size and the selection field will be disabled when RAID1 is selected.

Figure 153. Create Storage Pool

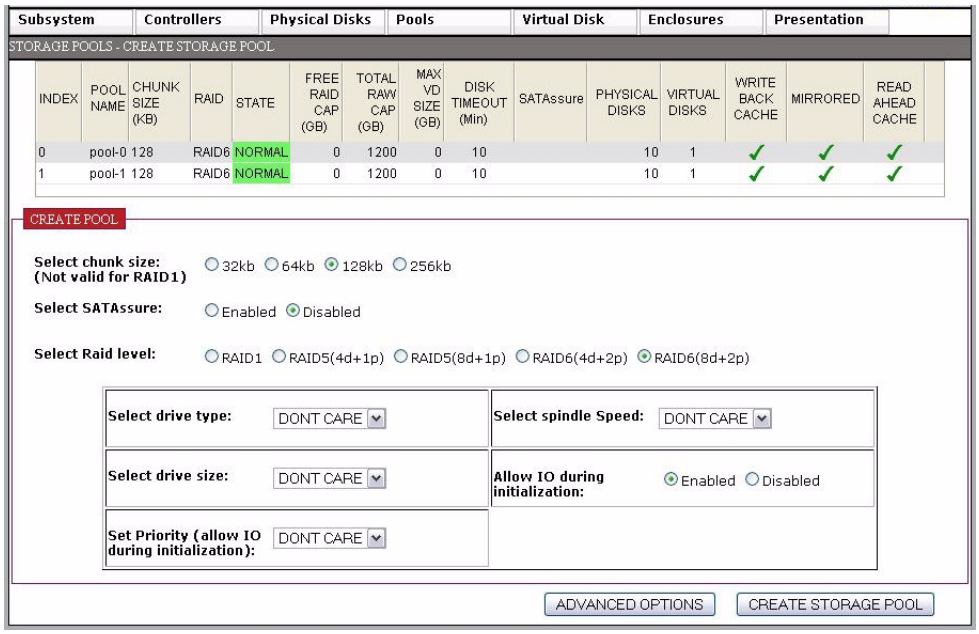

3. OPTIONAL: You may explicitly select the disks to add to the pool using the **ADVANCED OPTIONS**. At the Create Storage Pool Advanced screen (Figure 154), select the disks to add to the pool to be created. Otherwise, click **CANCEL** to cancel or click **BACK** to return to the previous screen.

Figure 154. Select Disks for Storage Pool

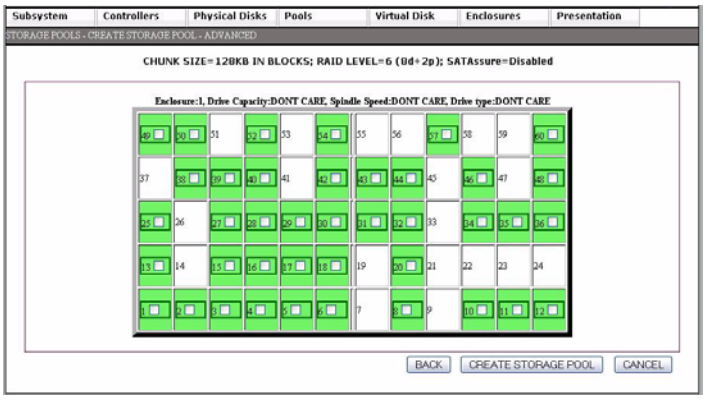

4. Click **CREATE STORAGE POOL** to create the pool. A message appears to indicate if the pool has been successfully created. The list of Storage Pools is displayed (Figure 155).

Figure 155. Storage Pool List

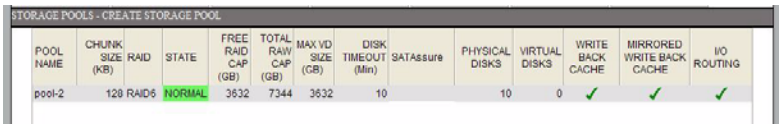

5.3.3 Configure Storage Pool Attributes

1. Select Set Storage Pool Attributes from the Pools Menu (Figure 152).

2. At the Set Attributes for Pool screen (Figure 156), select the Storage Pool whose attributes you want to change.

Figure 156. Set Storage Pool Attributes Screen

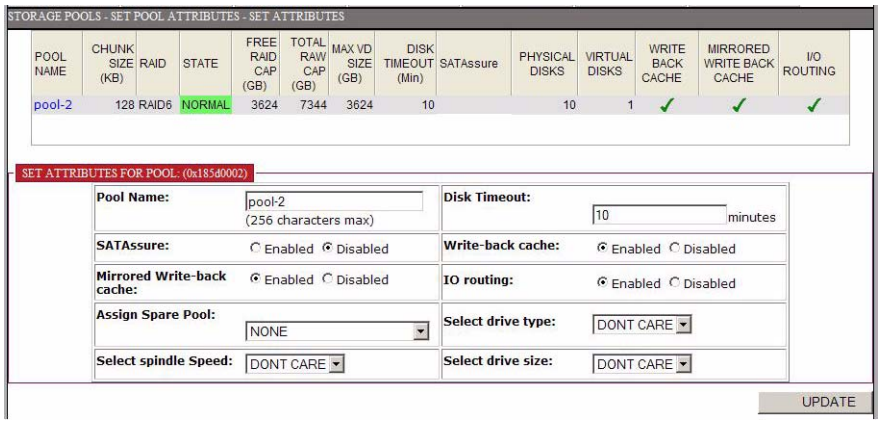

3. Make the desired changes.

4. Click **UPDATE** to save the changes. A message appears to indicate if the change was successful.

5.3.4 Locate Storage Pools

1. Select Locate Storage Pool from the Pools Menu (Figure 152).

2. At the Locate Pool screen (Figure 157), select the Storage Pool that you want to locate. The screen changes accordingly.

Figure 157. Locate Storage Pool Screen

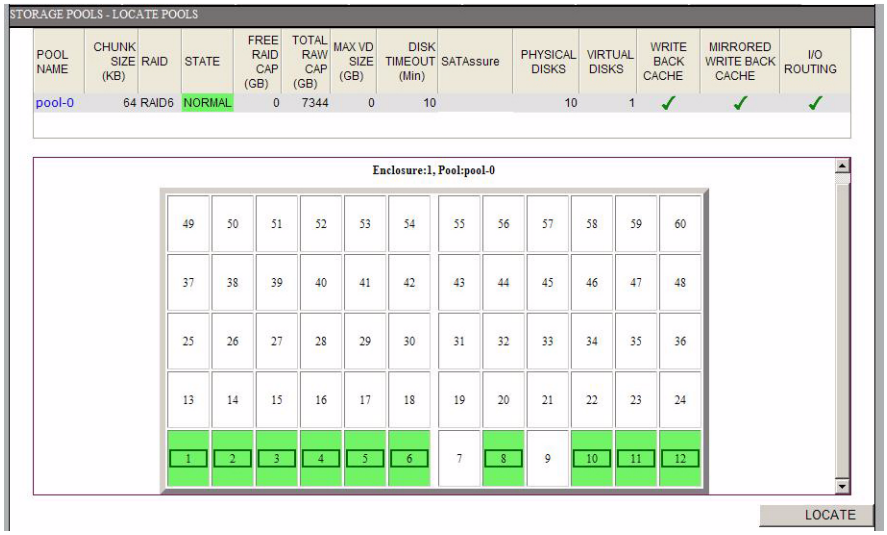

3. Click **LOCATE**. A message appears to indicate if the operation was successful.

5.3.5 Show Storage Pools

Select Show Storage Pools from the Pools Menu (Figure 152) to view the list of configured Storage Pools.

5.3.6 Show Unassigned Pools

Select Show Unassigned Pools from the Pools Menu (Figure 152). By default, all the disks or replacement disks are initially allocated to this pool. When a disk has failed or been released (from a Spare Pool, for example), it is also moved to this pool.

5.3.7 Delete a Storage Pool

1. Select Delete Storage Pool from the Pools Menu (Figure 152).

2. Select the Storage Pool that you want to delete.

3. Click **DELETE** to delete that Storage Pool or click **CANCEL** to cancel operation. A message appears to indicate if the deletion was successful.

5.3.8 Storage Pool Initialization

When a Storage Pool is created, initialization begins immediately as a background job and will continue until it is completed. Once the Storage Pool has completed its initialization, Virtual Disk(s) can be presented to host systems. However, you may immediately create the presentations.

You can monitor the progress of a job using the Show Background Jobs function under the Subsystem Menu (Figure 158).

Figure 158. Subsystem Menu

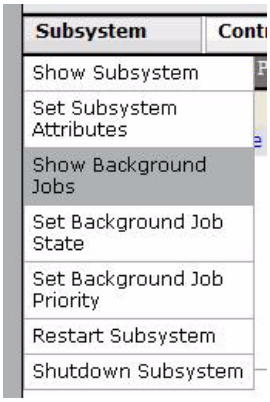

It will display the list of current jobs and the percentage of completion. Select a job to display its details (Figure 159).

Figure 159. Show Background Jobs Screen

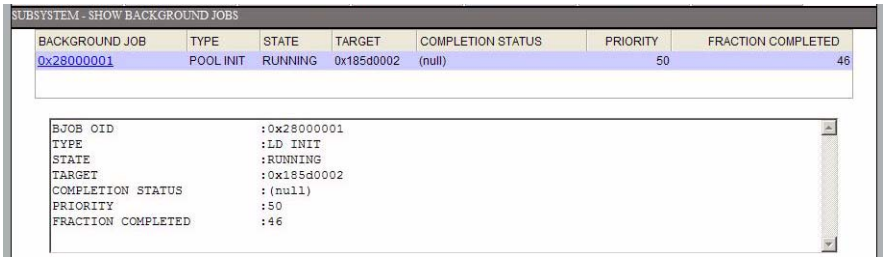

You may also check if a Virtual Disk is ready for access using the Show Virtual Disks function from the Virtual Disk Menu (Figure 160).

Figure 160. Virtual Disk Details Screen

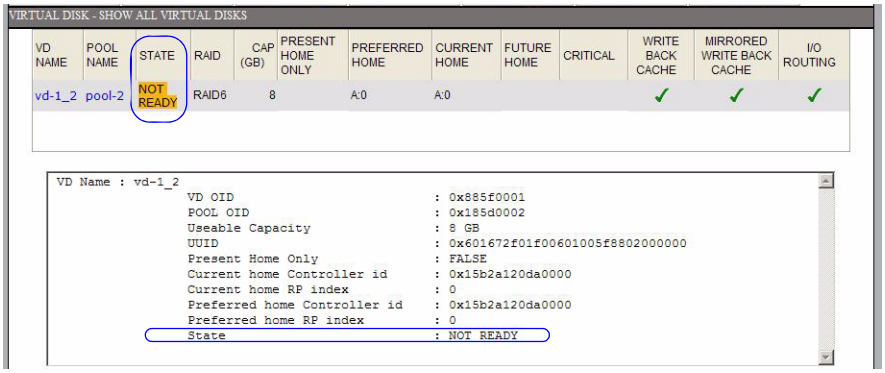

5.4 Virtual Disks

A Virtual Disk is the storage unit presented to any attached host. A Virtual Disk can be created to use all or just a part of the capacity of a single Storage Pool. VDs are created in increments of 8GiB. For example, 16GiB of storage space will be allocated when creating a VD of 10GiB.

5.4.1 Virtual Disk Menu

The following commands are located under the Virtual Disk pull down menu (Figure 161):

- Show Virtual Disks
- Create Virtual Disk
- Set Virtual Disk Attributes
- Delete Virtual Disk

Figure 161. Virtual Disk Menu

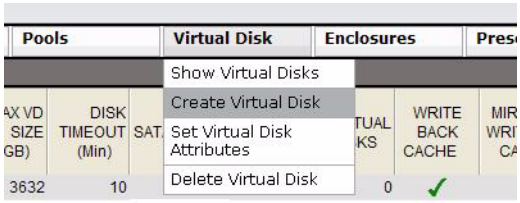

5.4.2 Show Virtual Disks

Select Show Virtual Disks from the Virtual Disk Menu (Figure 161). The list of configured Virtual Disk(s) appears (Figure 162). Select the individual VD name to display its detailed information.

Figure 162. Show Virtual Disk Details Screen

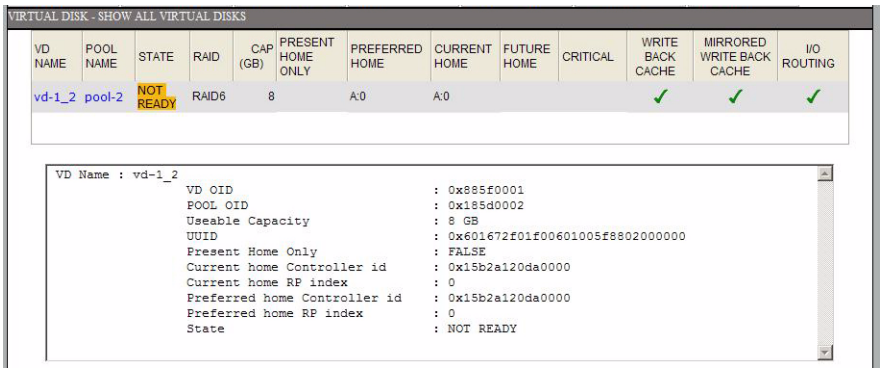

5.4.3 Create a Virtual Disk

1. Select Create Virtual Disk from the Virtual Disk Menu (Figure 161).

2. At the Create Virtual Disk screen (Figure 163), select the Storage Pool to be used.

Figure 163. Create Virtual Disk Screen

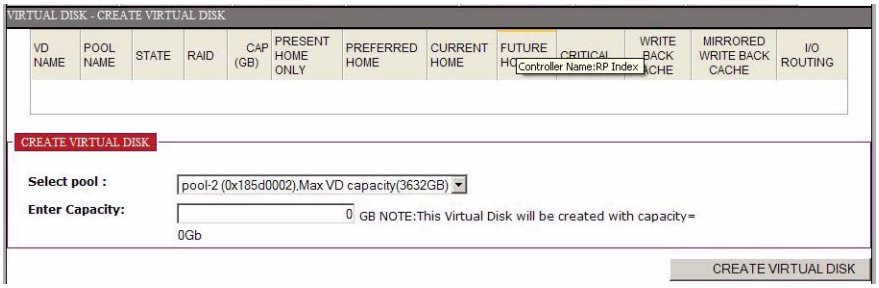

3. Enter the capacity in GiB.

4. Click **CREATE VIRTUAL DISK** to create the VD. A message appears to indicate if the VD has been successfully created. The list of Virtual Disks is displayed (Figure 164).

Figure 164. Virtual Disk List

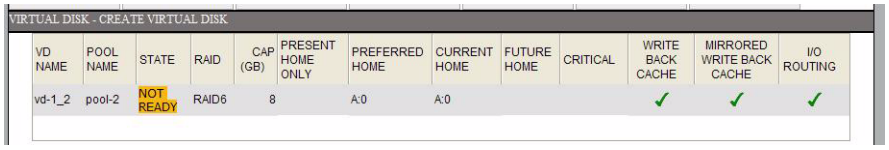

5.4.4 Configure Disk Virtual Attributes

1. Select Set Virtual Disk Attributes from the Virtual Disk Menu (Figure 161).

2. At the Set Attributes for Pool screen (Figure 165), select the VD whose attributes you want to change.

Figure 165. Set Virtual Disk Attributes Screen

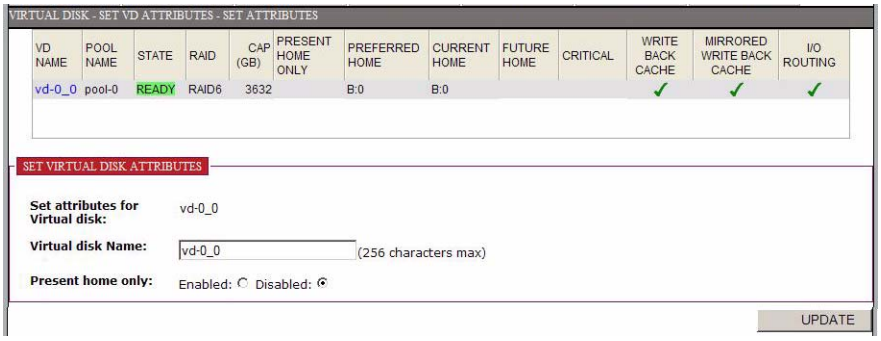

3. Make the desired changes.

4. Click **UPDATE** to save the changes.

5.4.5 Delete a Virtual Disk

When you delete a Virtual Disk, you lose access to all data that was stored on that Virtual Disk.

1. Select Delete Virtual disk from the Virtual Disk Menu (Figure 161).
2. Select the Virtual Disk that you want to delete.
3. Click **DELETE** to delete that VD or click **CANCEL** to cancel operation. A message appears to indicate if the deletion was successful.

If a presentation is configured for a Virtual Disk, deleting the Virtual Disk will return an error. You must first delete all the presentations configured on that Virtual Disk using the Delete Presentation function under the Presentation Menu.

5.5 Spare Pools

The IS16000 supports the concept of Spare Pool. A Spare Pool contains physical disks that can be used as spare disks in one or more Storage Pools. In the event of a disk failure, a disk that is assigned to a Spare Pool is automatically swapped into the Storage Pool to replace the failed disk. A rebuild then occurs immediately, reducing the possibility of data loss due to additional disk failures.

Spare Pools can be shared by Storage Pools or dedicated to a particular Storage Pool in which case the spare disks will only be used by the designated Storage Pool.

NOTE : Each Storage Pool must have a Spare Pool assigned to it.

5.5.1 Spare Pool Menu

The following commands are located under the Pools pull down menu (Figure 166):

- Show Spare Pools
- Create Spare Pool
- Delete Spare Pool
- Set Spare Pool Attributes
- Locate Spare Pool

Figure 166. Pools Menu

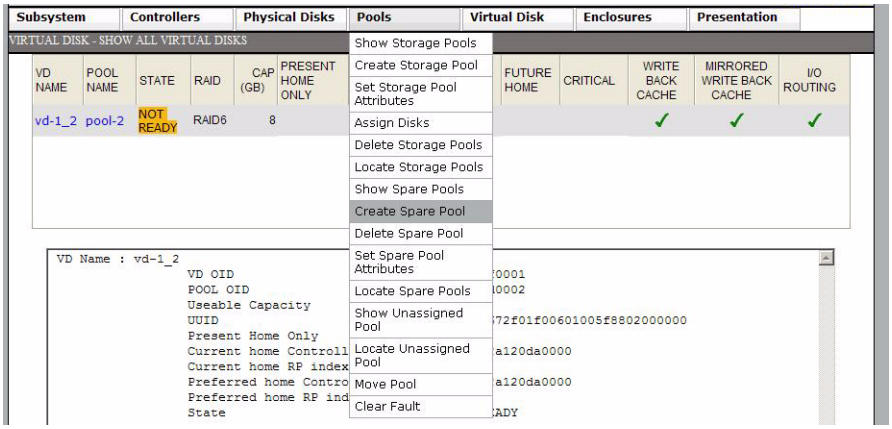

5.5.2 Show Spare Pools

Select Show Spare Pools from the Pools Menu (Figure 166) to view the list of configured Spare Pools.

5.5.3 Create a Spare Pool

1. Select Create Spare Pool from the Pools Menu (Figure 166).

2. At the Create Spare Pool screen (Figure 167), select the drive size, spindle speed, and drive type for the pool.

Figure 167. Create Spare Pool

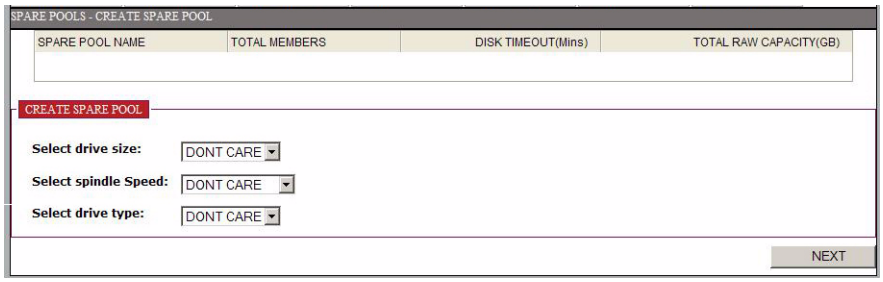

3. Click **NEXT** to bring up the Create Pool Advanced screen (Figure 168). Check to select the disks to add to the pool to be created.

NOTE : When assigning disk drives to a Spare Pool that is to be assigned to a given Storage Pool, be sure to select a disk that is as large or larger than any disks in the assigned Storage Pool.

Figure 168. Select Disks for Spare Pool

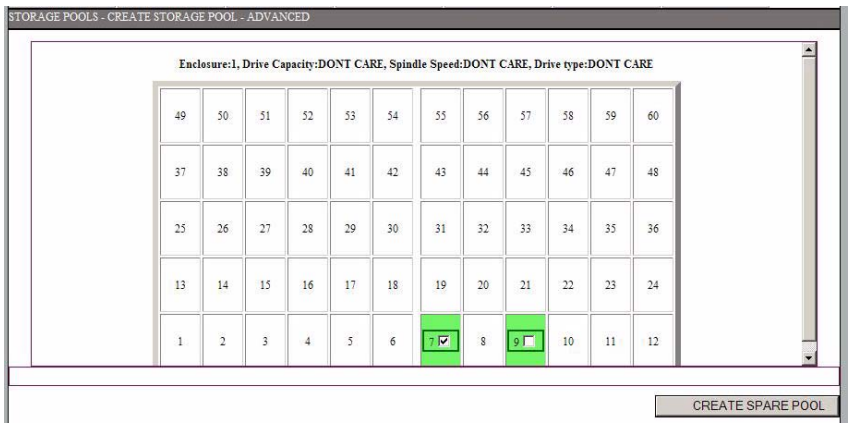

4. Click **CREATE SPARE POOL** to create the pool. A message appears to indicate if the pool has been successfully created. The list of Spare Pools is displayed (Figure 169).

Figure 169. Spare Pool List

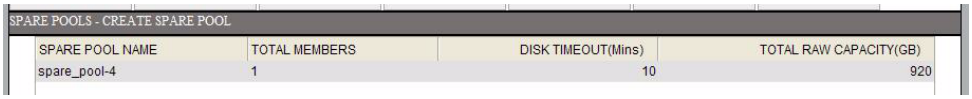

5. Assign the Spare Pool to a Storage Pool as described in Section 5.5.4, "Set Spare Pool Attributes".

5.5.4 Set Spare Pool Attributes

1. Select Set Spare Pool Attributes from the Pools Menu (Figure 166).

2. At the Set Attributes for Pool screen (Figure 156), select the Spare Pool whose attributes you want to change.

Figure 170. Set Spare Pool Attributes Screen

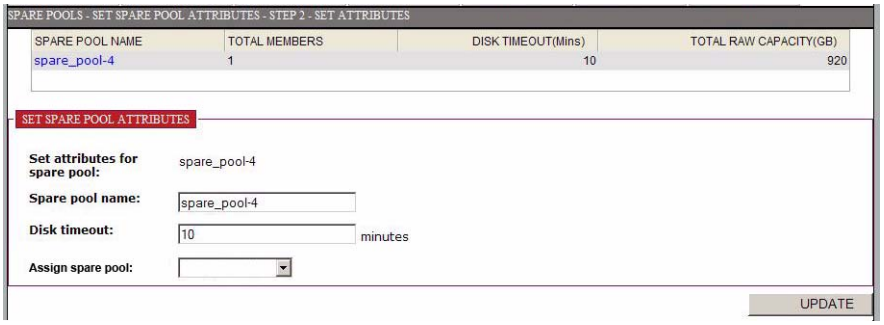

3. Make the desired changes.

   To assign the Spare Pool to a Storage Pool, select the name of the pool from the Assign Spare Pool drop down.

NOTE : When assigning a Spare Pool to a Storage Pool, be sure to select one with disks that are as large or larger than any disks in that Storage Pool.

---

4. Click **UPDATE** to save the changes. A message appears to indicate if the change was successful.

5.5.5 Locate Storage Pools

1. Select Locate Spare Pool from the Pools Menu (Figure 166).

2. Then select the Spare Pool you want to locate. The screen changes accordingly.

3. Click **LOCATE**. A message appears to indicate if the operation was successful.

5.5.6 Delete a Spare Pool

1. Select Delete Spare Pool from the Pools Menu (Figure 166).

2. Select the Spare Pool you want to delete.

3. Click **DELETE** to delete that Spare Pool or click **CANCEL** to cancel operation. A message appears to indicate if the deletion was successful.

5.6 Presentations

Virtual Disks are only presented to the hosts that have been given authorized access. A Presentation on a IS16000 has the following components:

- Discovered Initiators
  A Discovered Initiator is defined as a Fibre Channel or InfiniBand initiator that has logged into the IS16000. This is normally a single port on a Fibre Channel or InfiniBand HBA in an external computer.

- Host
  A Host is an abstract client of a Virtual Disk that has all the access to all data within the Virtual Disk. This most commonly represents an external computer with one or more Fibre Channel or InfiniBand initiators. Host attributes are OS (operating system) of which you can select GENERIC, WINDOWS, LINUX or MAC_OS. The OS attribute dictates the way the controller responds to certain Fibre Channel commands where appropriate. For example, Mac OS X expects a different response from an inquiry of LUN0 (the Controller LUN) than a Windows host.

- Channel
  A Channel is one of the Fibre Channel or InfiniBand ports on the Controller. Channel attributes are MODE of which you can select either MAC_OS or STANDARD. When MAC_OS is selected, the port's node name will be set differently in order to be visible to a Macintosh system.

- Stack
  A Stack is the input side of the subsystem. In the case of the IS16000, the stack is defined as the Fibre Channel interface. In future versions of the product, different interface stacks may exist.

- Presentation
  A Presentation is the relationship between a Host and a Virtual Disk. A Presentation implies that the related Host has some sort of access to the Virtual Disk. Attributes of a Presentation are:
    - PORT: from which the host will see the Virtual Disk
    - READ_ONLY: controls read only access
    - PRESENT_HOME_ONLY: presents the specified Virtual Disk from its designated home controller only
    - LUN: user-specified LUN number that the Virtual Disk will show to the host.

5.6.1 Presentation Menu

Figure 171 shows the Presentation pull down menu.

Figure 171. Presentation Menu

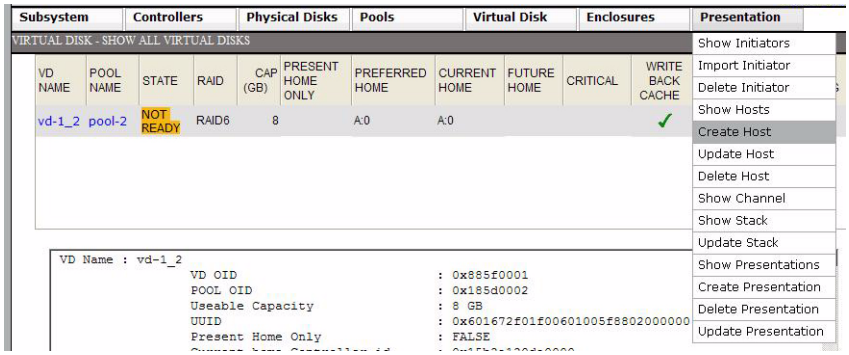

5.6.2 Show Presentations

Select Show Presentations from the Presentation Menu (Figure 171). The list of configured Presentations appears (Figure 172).

Figure 172. Show Presentations Screen

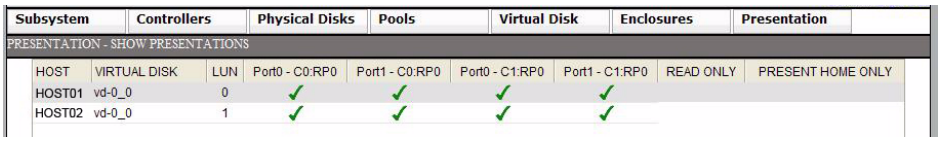

5.6.3 Set up a Presentation

Create a host:

1. Select Create Host from the Presentation Menu (Figure 171).
2. At the Create Host screen (Figure 173), enter a host name and specify the Stack and OS attributes.

Figure 173. Create Host

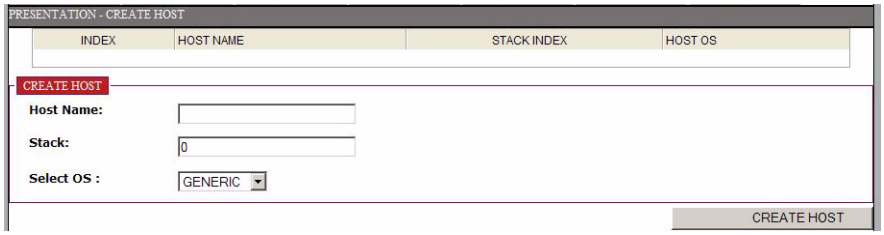

3. Click **CREATE HOST**. A message appears to indicate if the host has been successfully created. The list of created hosts is then displayed (Figure 174).

Figure 174. Created Hosts

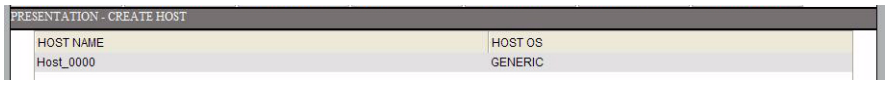

Map your newly created host to one or more Discovered Initiators:

1. Select Import Initiators from the Presentation Menu (Figure 171).

2. At the Import Initiator screen (Figure 175), select an initiator and the corresponding host(s).

Figure 175. Import Initiator Screen

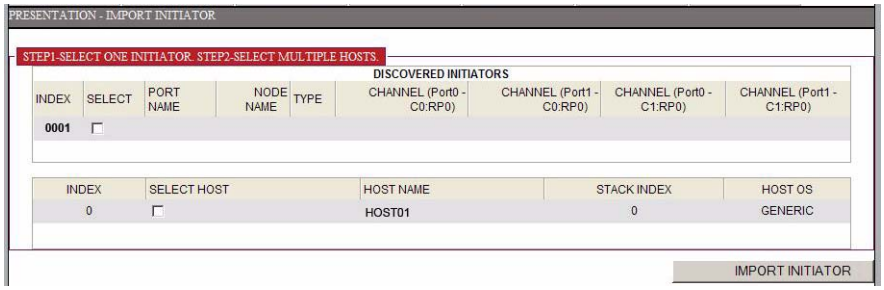

3. Click **IMPORT INITIATOR**. A message appears to indicate if the initiator has been imported successfully.

Present the Virtual Disks to the Hosts:

1. Select Create Presentation from the Presentation Menu (Figure 171).

2. At the Create Presentation screen (Figure 176), select a host and a Virtual Disk.

Figure 176. Create Presentation Screen

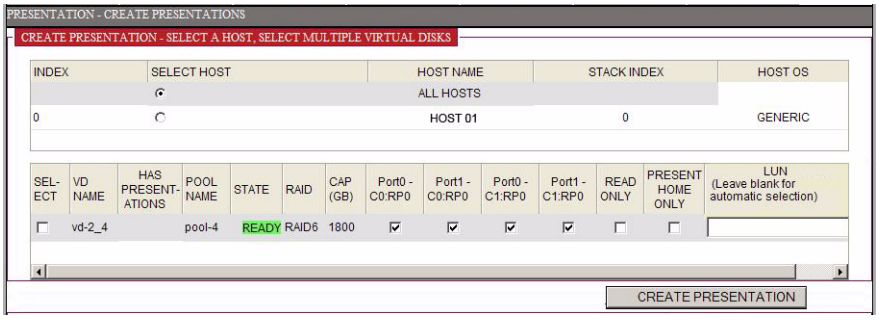

3. Then select the mask options and enter the Logical Unit Number (LUN).

   OPTIONAL: To present a Virtual Disk to all host ports on both Controllers, select “ALL HOSTS” instead of individual hosts.

4. Click **CREATE PRESENTATION**. A message appears to indicate if the presentation has been successfully created.

The external computer can access the Storage Pools once it has completed its initialization. You may also check if a Virtual Disk is ready for access using the Show Virtual Disks function under the Virtual Disk Menu.

5.7 Email and Critical Event Notification Setup

Emails will be sent automatically as notification of a selected group of warning and error events that have occurred on the IS16000.

The recipient address is user-configurable and only one address is permitted per system.

To set up or change the email address:

1. Select Update Email Settings from the Subsystem Menu.

   The upper portion of the Update Email Settings screen (Figure 177) will display the current settings, if any.

Figure 177. Update Email Settings Screen

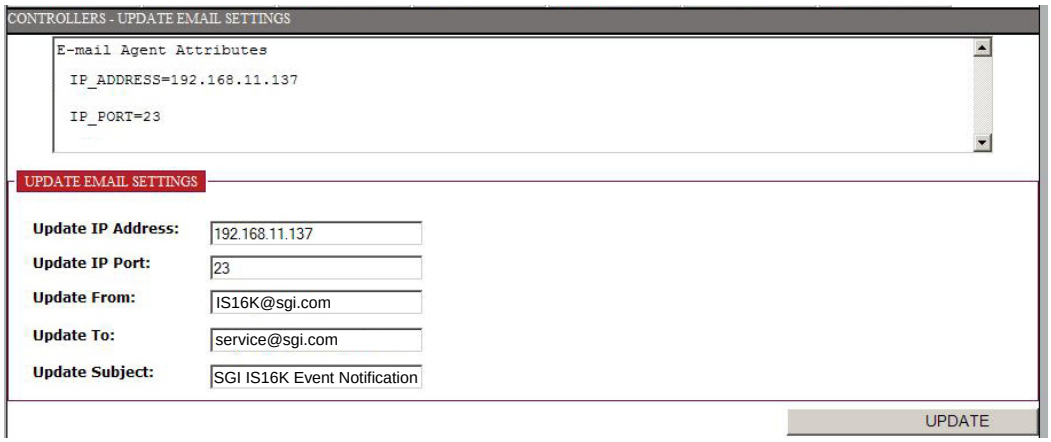

2. In the Update IP Address field, enter the IP address of the SMTP server to which the IS16000 will send email notifications.

3. In the Update IP Port field, enter the IP port of the SMTP server to which the IS16000 will send email notifications.

4. In the Update From field (maximum field size 512), enter an identification string, such as the name of the Controller. This is a string which will be included on the “From” line of the email notification message sent by the agent.

5. In the Update To field (maximum field size 512), enter an Internet Mail Format (RFC2822) email address to which IS16000 will send email notifications.

6. In the Update Subject field (maximum field size 512), enter a string to be included on the subject line of the email notification sent by the email agent.

7. Click **UPDATE** to save the changes. A message appears to indicate if the change was successful.

The tables in Section 3.7.5.3, "Inquiry Items and Events" on page 103 detail the additional inquiry items and events that will be trapped in the SNMP MIB and in the Email Agent.

5.8 NTP Mode

NTP (Network Time Protocol) mode is available on the IS16000. It provides a means for the Controllers to synchronize their time across a network, usually within a small number of milliseconds over a long period of time. You can enter up to four NTP addresses as the time servers.

To display the current settings:

1. Select Show Subsystem from the Subsystem Menu (Figure 178).

Figure 178. Subsystem Menu

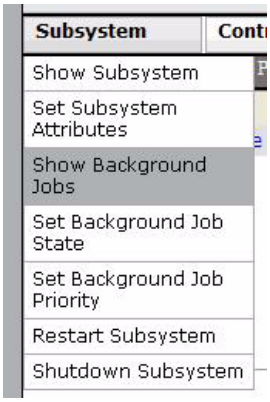

2. The Show Subsystem screen (Figure 179) displays the current setting for NTP mode and all the configured NTP IP addresses.

Figure 179. Show Subsystem Screen

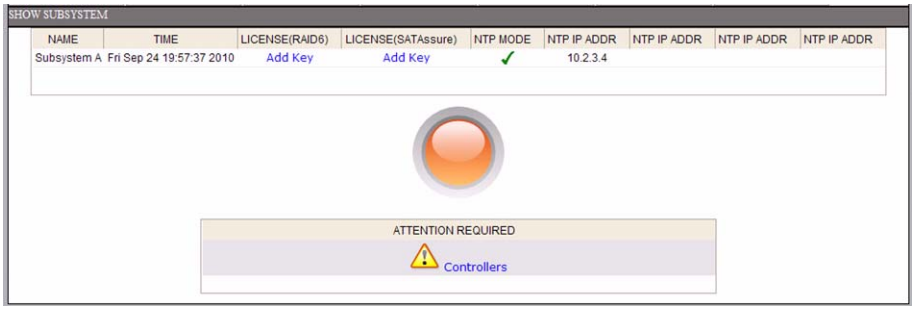

To change the settings:

1. Select Set Subsystem Attributes from the Subsystem Menu (Figure 178).

2. At the Set Subsystem Attributes screen (Figure 180), select ON/OFF to enable/disable NTP.

Figure 180. Set Subsystem Attributes Screen

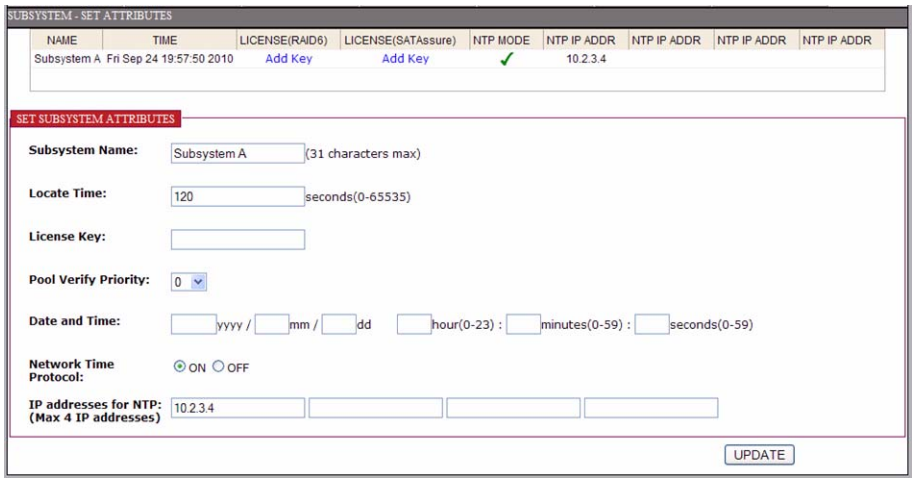

3. In the IP Addresses for NTP field, enter the IP address of the time servers.

4. Click **UPDATE** to save the changes. A message appears to indicate if the change was successful.

Please refer to page 74 for information on behavioral changes when NTP mode is enabled.

Appendices

A

Technical Specifications

Here are the technical specifications for the IS16000.

NOTE : Specifications subject to change without notice.

Configuration, Performance, & Capacity

Host side technology	(16) 8 Gb Fibre Channel or (8) 40 Gb InfiniBand
Disk side technology	(40) 4 × 3 Gb/s SAS
Supported disk technology	SAS / SATA /SLC SSD
Management interface	RS-232 and Ethernet (SSH)
Cache memory	24GB of RAM, 12GB ECC/RAID protected cache
Number of disk modules supported	Up to 600 disks, 1.2 Petabytes per array
Number of Storage Pools supported	256
Number of Virtual Disks supported	512; 64 per Storage Pool
Number of Spare Pools supported	16
Hot spare capability	Yes
RAID parity protection	1+1 RAID 1, 4+1 or 8+1 RAID 5, 4+2 or 8+2 RAID 6
Throughput	12GB/s sustained with large well aligned, sequential IO
IOPS	1M to cache; 300,000 to Disk

Reliability

SES (SCSI Enclosure Services) protocol support	Yes
Temperature monitoring	Yes
Battery-backed write-back cache	Yes
Redundant hot-swappable power modules	Dual-redundant
Redundant hot-swappable cooling modules	N+1
Redundant controllers	Dual-redundant
Redundant disk enclosures	N+1

Physical, Power & Environmental

Active/Active Dual Controllers	Dimensions	Height: 14" (356 mm) (includes 2 UPS units) Width: 17" (432 mm) Depth: 25.5" (648 mm)
	Weight	120 lbs (54.5 kg) Controllers only; 248 lbs (112.5 kg) with UPS units
	Voltage range	200-240 VAC @ 47-63 Hz
	Average power	1200 W
	Average cooling	4095 BTU/hr
Disk Enclosure	Dimensions	Height: 7" (178 mm) Width: 17" (432 mm) Depth: 36" (914 mm) without bezel; 42" (1067 mm) with cable management arms
	Weight	240 lbs (109 kg) with disk modules; 120 lbs (54.5 kg) without disk modules
	Voltage range	200-240 VAC @ 47-63 Hz
	Average power	1750 W
	Average cooling	5973 BTU/hr

Operating environment (temperature / relative humidity)	5 °C to 35 °C / 20% - 80%, non-condensing
Non-operating environment (temperature/ relative humidity)	-10 °C to 50 °C / 20% - 80%, non-condensing
Certification	UL, CE, CUL, C-Tick, FCC

NOTE : Specifications subject to change without notice.

B

Terminology

AP	Application Processor
Application Stack	A layer of the Controller OS that provides services external to the appliance through the use of Virtual Disks. For example, Fibre Channel block device.
Client Channel	A port on a Client I/O Controller that can be used by hosts to communicate with an Application Stack.
Client I/O Controller	A physical interface (FC or IB) that has a unique PCI bus number and is utilized for one or more Client Channels. A Client I/O Controller is managed by a single Application Stack.
Couplet	Dual Controller subsystem
Discovered Initiator	A volatile object that represents an FCP or SRP initiator port that is discovered on the fabric.
Disk Channel	A port on a Disk I/O Controller that can be connected to one or more Physical Disks and Disk Enclosures.
Disk I/O Controller	A physical interface (SAS) that has a unique PCI bus number and is utilized for one or more Disk Channels. A Disk I/O Controller is managed by the RAID Stack running on a single RAID Processor.
GiB (gibibyte)	A unit of digital information storage and denote 1,073,741,824 bytes (2^{30} bytes).
Host	A persistent logical object that represents an abstract client for one or more Virtual Disk's and that is given access to data on those Virtual Disk's.
ICL	Inter-Controller Link
Initiator	A persistent logical object associated with a Host that represents a FCP or SRP initiator port on a Host's HBA.
KiB (kibibyte)	A unit of digital information storage and denote 1,024 bytes (2^{10} bytes).
MiB (mebibyte)	A unit of digital information storage and denote 1,024 kibibytes (2^{20} bytes).

Physical Disk (PD)	A disk module (SAS or SATA) or solid-state device (SSD) used by the system back-end to store data and subsystem metadata.
Pool	A set of physical disks. There are three types: Storage Pool, Spare Pool, and Unassigned Pool.
Presentation	A persistent logical object that describes an association between a Host and a Virtual Disk that specifies if a Host may access the Virtual Disk and any specific constraints for access.
RAIDset	A container formed by a set of Physical Disk extents and controlled by one of the RAID algorithms, such as RAID6.
RP	RAID Processor
Singlet	Single Controller
Spare Pool	A set of physical disks that are used as spares in one or more Storage Pools.
Storage Pool	A set of physical disks whose extents may be parts of RAIDset's.
Unassigned Pool	The set of physical disks that are not assigned to any pool.
Virtual Disk (VD)	A set of disk blocks that may be presented to an operating system. It is also the primary logical object realized by the back-end.

Index

G

H

I

J

L

M

N

P

R