SGI® InfiniteStorage 6120 RAID User's Guide

007-5533-002

Preface

What is in this guide

This user guide gives you step-by-step instructions on how to install, configure and connect the SGI InfiniteStorage 6120 system to your host computer and how to use and maintain the system.

International Standards

The InfiniteStorage 6120 complies with the requirements of the following agencies and standards:

- CE
- UL

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.

Safe Handling

- Remove drives to minimize weight.
- Do not try to lift the enclosure by yourself.
- Do not lift the InfiniteStorage 6120 by the handles on the power cooling module (PCM); 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 InfiniteStorage 6120 *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 enclosure top covers closed.

- Plug-in modules are part of the enclosure 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 InfiniteStorage 6120 system must only be operated from a power supply input voltage range of 200 VAC to 240 VAC.
- The equipment is intended to operate with two (2) working PCMs. Before removal/replacement of any module, disconnect all supply power for complete isolation.
- A faulty PCM 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 PCM 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 PCM from the enclosure.
- 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.

WARNING: Do not remove covers from the PCM. Danger of electric shock inside. Return the PCM to your supplier for repair.

WARNING: Operation of the InfiniteStorage 6120 with *ANY* modules missing will disrupt the airflow and the drives 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 InfiniteStorage 6120 plug-in modules and components. Avoid contact with backplane components and module connectors.

Data Security

- Power down your host computer and all attached peripheral devices before beginning installation.
- Each enclosure contains up to 60 removable disk drive modules. Disk units are fragile. Handle them with care, and keep them away from strong magnetic fields.
- ***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 enclosure is used with modules or blanking plates missing for more than a few minutes, the enclosure can overheat, causing power failure and data loss. Such use may also invalidate the warranty.
- If you remove a drive module, replace it immediately. If it is faulty, replace it with a drive module of the same type and capacity.
- Ensure that all disk drives are removed from the enclosure before attempting to move the rack installation.
- Do not abandon your backup routines. No system is completely foolproof.

Related Publications

You should also see the following document:

SGI InfiniteStorage 6120 RAID CLUI Command Reference (007-5561-xxx)

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

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

1.1 Introduction . . . 3
1.1.1 A SAN Solution . . . 3
1.1.2 Features of the InfiniteStorage 6120 . . . 3

1.2 The InfiniteStorage 6120 System Hardware . . . 5
1.2.1 Enclosure Chassis . . . 6
1.2.2 Controller . . . 7
1.2.4 Drive Modules . . . 9
1.2.5 Drive Expander Module (DEM) Card . . . 10
1.2.6.1 Front Panel Indicators . . . 11
1.2.6.2 Internal Indicators . . . 12

2.1 Installation Overview . . . 15

2.2 Setting Up the InfiniteStorage 6120 . . . 16
2.2.1 Unpacking the InfiniteStorage 6120 . . . 16
2.2.3 Connecting the InfiniteStorage 6120 . . . 20
2.2.3.1 Host Connections . . . 20
2.2.3.2 Expansion Enclosure Connections . . . 20
2.2.3.3 Power Connections . . . 21
2.2.4 Powering On the InfiniteStorage 6120 (and Expansion Enclosure) . . . 21

2.3 Configuring the InfiniteStorage 6120 . . . 23
2.3.1 Planning Your Setup and Configuration . . . 23
2.3.2 Drive Layout . . . 24
2.3.3 Initial Configuration via RS-232 . . . 24
2.3.3.1 Serial Interface Configuration . . . 24
2.3.3.2 Configure Network Settings . . . 25
2.3.4 Setting System Time & Date . . . 25
2.3.6 Installing Licensed Features . . . 27
2.3.7 Creating Storage Pools . . . 28
2.3.8 Creating Virtual Disks . . . 31
2.3.9 Creating Spare Pools . . . 32
2.3.10 Initializing the Host System . . . 34
2.3.10.1 Present Virtual Disk to External Host . . . 35

3.1 Managing the InfiniteStorage 6120 . . . 39
3.1.1 Management Interface . . . 39
3.1.2 User Logins . . . 40
3.1.3 Available Commands . . . 40
3.1.3.1 Basic Key Operations . . . 41

3.2 Configuration Management . . . 42
3.2.1 Monitor Status of Storage Assets . . . 42
3.2.1.1 Disk Information . . . 42
3.2.1.2 Storage Pool View . . . 44
3.2.1.3 Virtual Disk View . . . 46
3.2.1.4 Spare Pool View . . . 47
3.2.1.5 Visual Indication . . . 47
3.2.2 Storage Pool and Virtual Disk Management . . . 48

3.2.2.1 Changing the Disk's Status . . . 48
3.2.2.2 Creating a Storage Pool . . . 49
3.2.2.3 Storage Pool Initialization . . . 49
3.2.2.4 Naming a Storage Pool . . . 49
3.2.2.5 Deleting a Storage Pool . . . 50
3.2.2.6 Creating a Virtual Disk . . . 50
3.2.2.7 Naming a Virtual Disk . . . 50
3.2.2.8 Deleting a Virtual Disk . . . 50
3.2.3.1 Creating a Spare Pool . . . 51
3.2.3.2 Naming a Spare Pool . . . 51
3.2.3.3 Deleting a Spare Pool . . . 51
3.2.4.1 Full and Partial Rebuilds . . . 52
3.2.4.2 Sparing Policy . . . 52
3.2.4.3 Manual Drive Replace/Rebuild . . . 53
3.2.5 SATAssure . . . 53
3.2.6 Cache Coherency . . . 54
3.2.7 Setting the System's Date and Time . . . 54

3.3 Performance Management . . . 56
3.3.1 Background Job Priority . . . 56
3.3.2 Rebuild Policy Priority . . . 57

3.4 Firmware Update Management . . . 58
3.4.1 Displaying Current Firmware Version . . . 58
3.4.2 Firmware Update Procedure . . . 58
3.4.2.1 Windows Environment (Update via USB Stick) . . . 59
3.4.2.2 Linux Environment (Update from Network) . . . 61
3.4.2.3 Windows Environment (Update from the Network) . . . 62

3.5 The System Logs . . . 64

3.6 Remote Management of InfiniteStorage 6120 . . . 65
3.6.1 Network Connection . . . 65
3.6.3 Logins . . . 66

3.7 Presentations . . . 67
3.7.1 Application Discovered Initiator Commands . . . 67
3.7.2 Application Host Commands . . . 67
3.7.3 Application Presentation Commands . . . 68

4.1 Status LED Indicators . . . 71
4.1.1 Front Panel Indicators . . . 71
4.1.2 Power Cooling Module (PCM) . . . 72
4.1.4 HDD (Hard Disk Drive) . . . 73

4.2 Troubleshooting . . . 75
4.2.1 Initial Start-up Problems . . . 75
4.2.2 RAID[0]$ Prompt Not Shown Correctly . . . 75
4.2.3 System Faults . . . 76
4.2.4 Power Supply/Cooling Faults . . . 76
4.2.5 Thermal Alarm . . . 77

4.3 Component Failure Recovery . . . 78
4.3.1 Replacing a Power/Cooling Module . . . 78

4.3.2 Replacing a Drive Module . . . 79

4.4 Recovering from a Drive Failure . . . 80
4.4.1 Automatic Rebuild . . . 80
4.4.2 Drive Replacement . . . 81
4.4.3 Spare Not Available . . . 81

4.5 Manual Intervention Required (MIR) Conditions . . . 82

5.1 The GUI Management Agent . . . 87

5.2 Starting the GUI Management Agent . . . 88
5.2.1 Login . . . 88
5.2.2 Initial Home Screen and Health Indicators . . . 88

5.3 Physical Disks . . . 90
5.3.1 Physical Disks Commands . . . 90
5.3.2 Checking Status of Physical Disks . . . 90

5.4 Storage Pools . . . 91
5.4.1 Storage Pool Commands . . . 91

5.5 Virtual Disks . . . 95
5.5.1 Virtual Disk Commands . . . 95

5.6 Spare Pools . . . 98
5.6.1 Spare Pool Commands . . . 98

5.7 Initializing the Host System . . . 102

5.8 Presentations . . . 104
5.8.1 Presentation Commands . . . 104

5.9 Email and Critical Event Notification Setup . . . 108

Appendix A. Technical Specifications . . . 111

Appendix B. Safety Requirements for Rack Installation . . . 113

Index . . . 115

SECTION 1

Introducing the

InfiniteStorage 6120

1.1 Introduction

This section describes the features of the SGI InfiniteStorage 6120 and its hardware.

1.1.1 A SAN Solution

The InfiniteStorage 6120 is a compact network storage system. It combines a controller chassis and disk enclosure chassis into a single self-contained enclosure.

The InfiniteStorage 6120 plugs seamlessly into existing network environments, thereby protecting and upgrading investments made in legacy storage and networking products to substantially improve their performance, availability and manageability.

The InfiniteStorage 6120 is designed to support mixed workload environments with industry-leading density, a large number of small random IOPS, and high sequential throughput to make it perfect shared storage for transactional, backup, or rich-media markets. It enables a multi-vendor environment comprised of standalone and clustered servers, workstations, and PCs to access and back up data stored in centralized or distributed storage devices in an easy, cost-effective and reliable manner. This modular architecture ensures high data availability and uptime along with application performance.

1.1.2 Features of the InfiniteStorage 6120

The InfiniteStorage 6120 incorporates the following features:

- **Fibre Channel connectivity**
  The InfiniteStorage 6120 provides up to 4 individual 8Gb/s Fibre Channel (FC) 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**
  Both controllers operate in active mode at all times meaning I/O requests can be serviced by either controller from any host port and failover is transparent to the host.
- **RAID 5 & RAID 6**
  The system is based on a high performance software RAID engine with default RAID5 support and an optional/licensed RAID6.
- **Configurable RAID group sizes**
  (5 or 9 drives per RAID 5 group, 6 or 10 drives per RAID 6 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 drives in the array is supported.
- **Stripe size of 64KB, 128KB, 256KB**
  Configurable stripe sizes aid in system tuning for the application's data pattern. Stripe sizes are uniquely configurable per RAID set.
- **2MB maximum request size**
  This is the maximum transfer size a host may request from the unit and covers nearly all commonly found uses.
- **SAS/SATA storage pool**
  This feature allows in-box storage pooling, where high-performance SAS drives are used for primary data and high-capacity SATA drives are used for secondary data or active archiving.

Flexible configurations along with the intermix of SATA and SAS drives within the same enclosure is supported. However, for performance and reliability reasons, there are best-practice guidelines.

- **Optional SATAssure data protection**
  SATAssure technology is designed to improve the reliability of enterprise SATA drives and to ensure that data integrity is always maintained for all I/O operations.
- **Partial rebuild**
  This feature reduces rebuild times by updating only the data that has changed while the drive was down.
- **Hot-swappable and redundant components**
  Adhering to enterprise RAS standard, almost all hardware components–such as drives, controllers, power cooling modules, and drive expander module (DEM) cards–are redundant and hot-swappable.
- **Hot spares**
  The unit supports as many hot spares as desired and they can be allocated as global hot spares or assigned to a specific RAID set.
- **Comprehensive, centralized management capability**
  The InfiniteStorage 6120 provides a wide range of management capabilities via a secured command line user interface (CLUI): configuration management, performance management, and firmware update management.
- **Management options via RS-232 and Ethernet (SSH)**
  A RS-232 port and Ethernet port are included to provide local and remote management capabilities.

1.2 The InfiniteStorage 6120 System Hardware

The InfiniteStorage 6120 is a four-unit (4U), rack-mountable enclosure that houses two power cooling modules, two controllers, and slots for up to 60 SAS or SATA drives (Figure 1). The controllers provide system management and host connectivity interfaces.

The InfiniteStorage 6120 design concept is based on a subsystem together with a set of plug-in modules and *(as supplied)* comprises:

- Enclosure chassis with front panel drive status indicator
- Two (2) power cooling plug-in modules (PCMs)
- Two plug-in controllers
- Up to 60 top loadable hard disk drives in a 5x12 matrix
- SAS and SATA drive intermix allowed
- Eight (8) SAS drive expander modules (DEMs)

Figure 1. The InfiniteStorage 6120 - Front and Rear Views

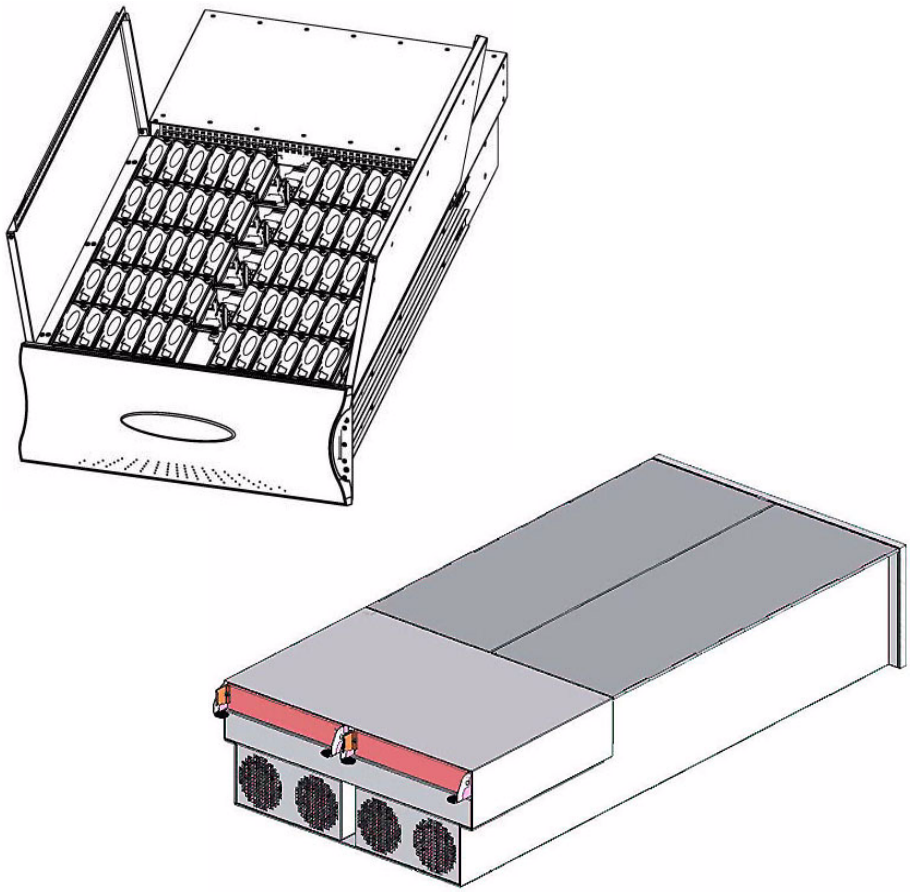

1.2.1 Enclosure Chassis

The chassis assembly contains 60 drive bays at the front, each of which accommodates a plug-in drive module capable of holding a 3.5-inch SAS or SATA hard disk drive (HDD) or solid-state disk (SDD). The 60 drive bays are arranged in five rows of twelve drives (5x12) (Figure 2). At the rear, the chassis assembly contains two (2) power cooling modules and two (2) controllers (Figure 3).

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, controllers, and the baseboard.

Figure 2. InfiniteStorage 6120 Top View

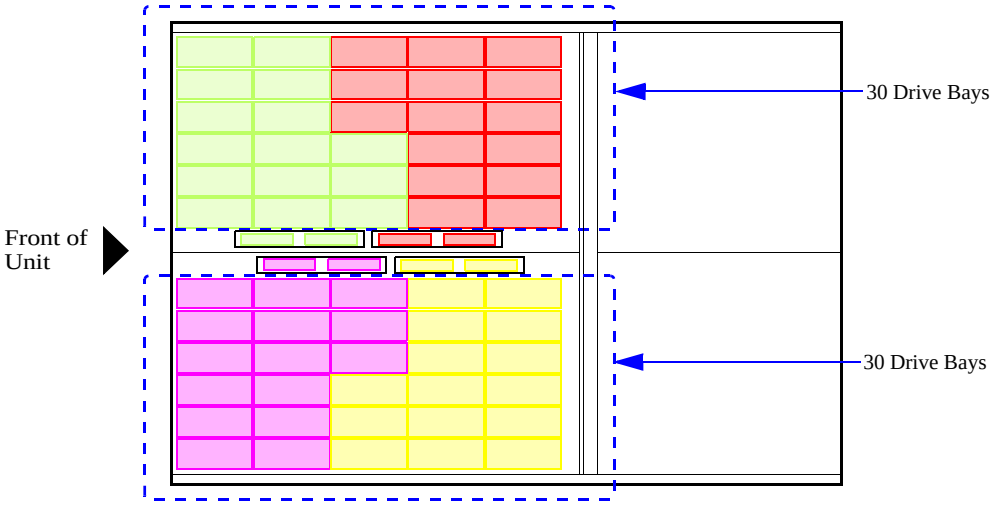

Figure 3. InfiniteStorage 6120 Rear View

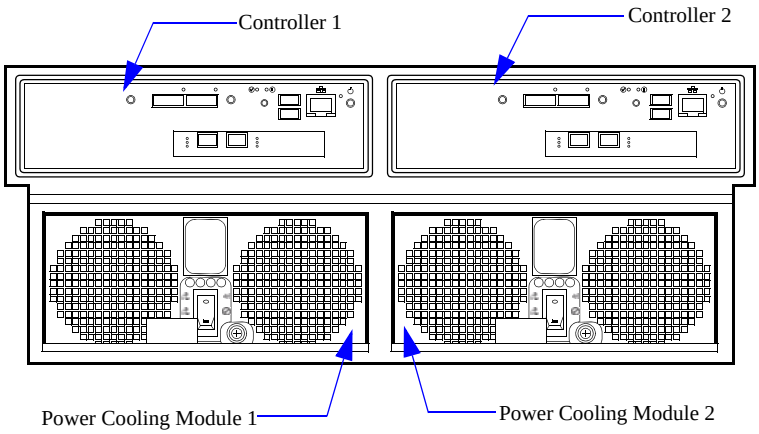

1.2.2 Controller

The two controllers are redundant and hot-swappable modules which provide the intelligence and active/active data protection features of the InfiniteStorage 6120 (Figure 4).

If a controller fails, the remaining controller will assume its functionality and continue to provide data access at a reduced performance level.

Figure 4. A Controller

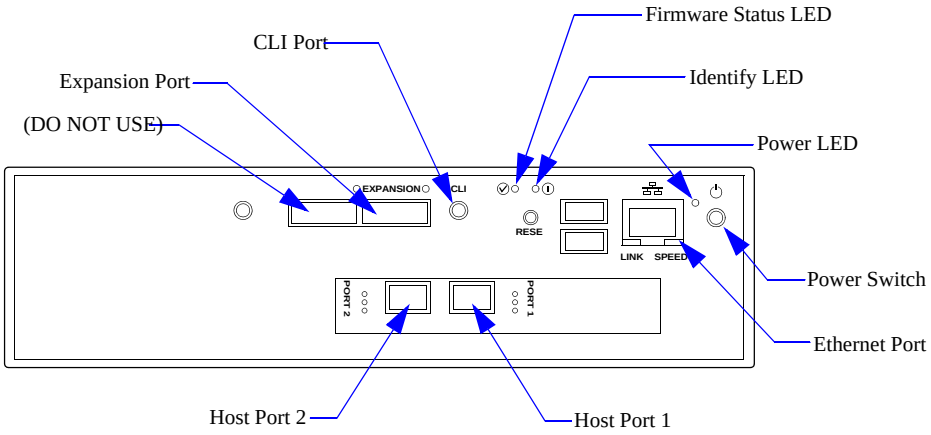

There are two 8Gb/s host ports for Fibre Channel host connections. You may connect your host systems directly to these ports or connect the ports to your Fibre Channel switches and hubs. When an FC signal is present on the port, the respective LED turns green.

The **EXPANSION** port is used to connect to a 60-drive expansion enclosure.

The **CLI** connector provides local system monitoring and configuration capabilities.

The Ethernet port provides remote monitoring and configuration capabilities. The **LINK** LED is off when there is no Ethernet link.

The Firmware Status LED blinks green when the InfiniteStorage 6120 firmware is loading. It stays solid green when the InfiniteStorage 6120 is ready to be accessed.

The Power LED turns green when there is power applied to the controller.

The Identify LED turns blue when an SES identify command is received.

1.2.3 Power Cooling Module (PCM)

Installed in the rear of the chassis (Figure 5) are two (2) power cooling modules.

Figure 5. Power Cooling Module

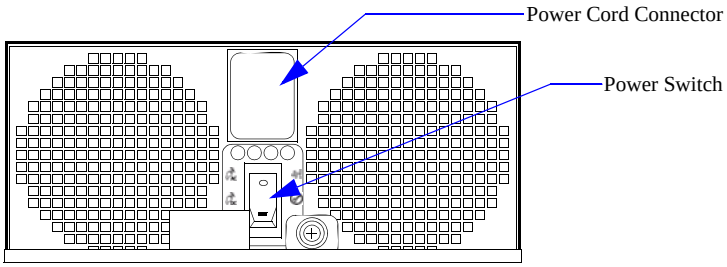

The InfiniteStorage 6120 must always be operated with two PCM installed. Module replacement should only take a few minutes to perform but must be completed with 10 minutes from removal of the failed module. Four (4) LEDs mounted on the PCM indicate the status of the module and the fans (Figure 6). Figure 7 provides a description of the color and status of the LEDs.

Figure 6. Power Cooling Module LEDs

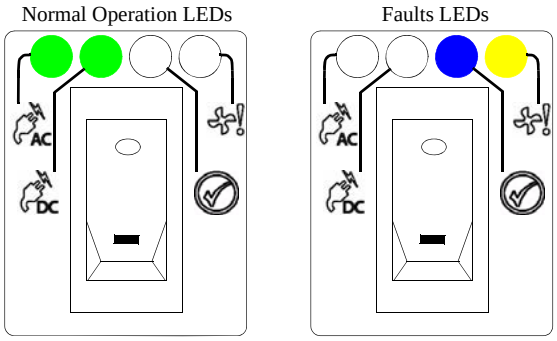

Figure 7. PCM LED Description

LED	Description	Color	Status
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 - PCM fault detected OFF - no detected PCM faults
	PCM ID	Blue	ON - receiving SES identify command OFF - NOT receiving SES identify command

1.2.4 Drive Modules

The drive module comprises a hard disk mounted in a carrier (Figure 8). Each drive 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 drive modules.

The module handle provides the following functions:

- Camming of the module into and out of drive bays
- Positive *spring loading* of the drive/baseplane connector
- Drive Status LED incorporated in handle assembly

NOTE : The InfiniteStorage 6120 design allows for drive bays to be left empty without the need for fitting dummy drive modules.

The InfiniteStorage 6120 supports a SATA interposer card which allows simultaneous access to the SATA HDD from both controllers installed in the enclosure. ***Note that a SATA interposer card, which is part of the drive carrier assembly, is required to run SATA HDDs***.

Figure 8. Drive Module

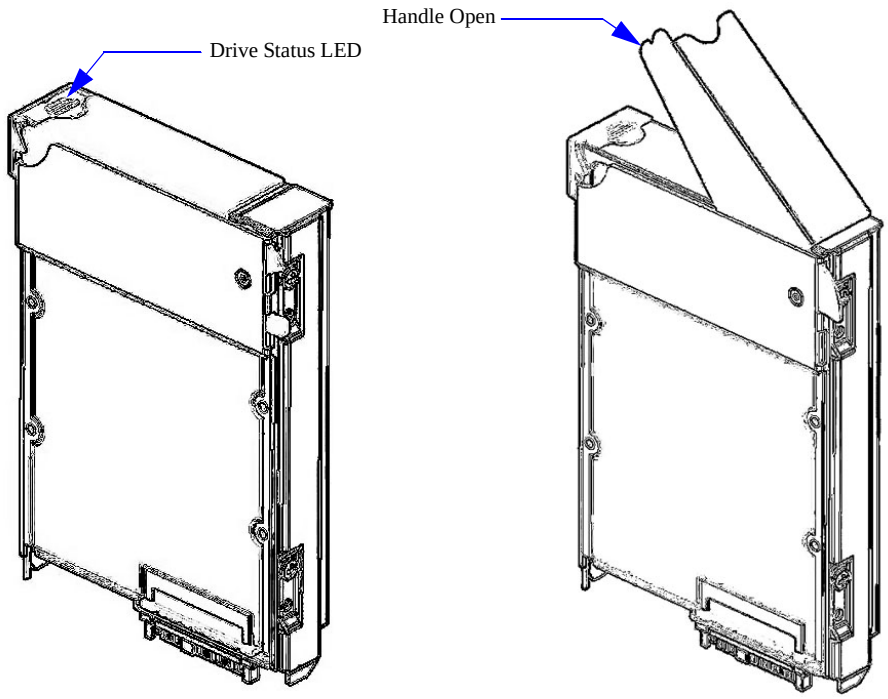

1.2.5 Drive Expander Module (DEM) Card

The InfiniteStorage 6120 contains eight (8) top-loadable, hot-swappable DEMs (Figure 9). The DEMs provide the SAS connectivity between the controllers and the HDDs located with the enclosure.

The DEMs are organized into four module pairs, with each module pair controlling a bank of 15 drives. 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 9. DEM Card Locations (Viewing from Top)

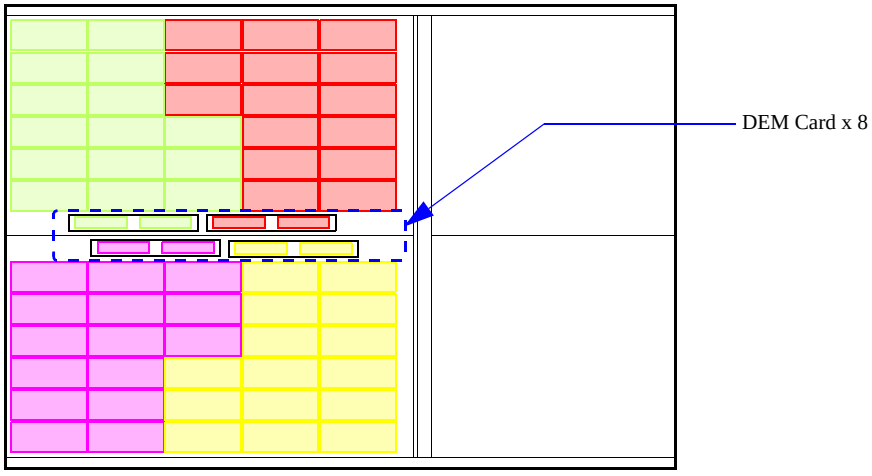

Each DEM (Figure 10) connects to a single port of the HDD based on the location within the enclosure within which it is installed. Each card provides power control signals to each drive slot. It controls HDD identify/service LEDs and monitors the status from partner DEM. Each pair of DEMs supports 15 drives.

Figure 10. DEM Card

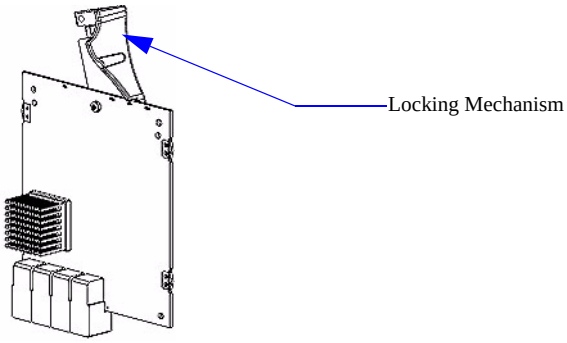

1.2.6 LED Indicators

The InfiniteStorage 6120 has both FRU and status LEDs to indicate the state of various components within the enclosure. Each FRU has a visible blue identity LED. Each FRU, except the drive module, has a visible amber LED which indicates a fault.

1.2.6.1 Front Panel Indicators

The front panel indicators show the aggregated status of all the modules (Figure 11). The LEDs are defined in Figure 12. The Drive Activity LEDs indicate drive presence and flash during data I/O.

Figure 11. Front Panel LED Indicators

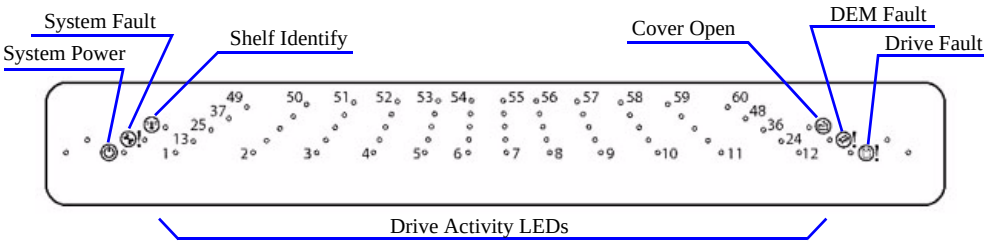

Figure 12. Front Panel LED Description

LED	Definition	Color	Status
	Shelf Identify	Blue	ON - SES is sending an identify command OFF - SES is NOT sending an identify command
	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
	System Power	Green	ON - DC power is present OFF - DC power is not present LED does NOT flash under normal operating conditions
	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
	Drive Fault	Amber	ON - one or more HDDs are failed; SES must determine exact HDD OFF - no detectable drive faults
Individually numbered	HDDs	Green	ON - when SAS HDD is present Blinks - indicates HDD activity OFF - no HDD activity

1.2.6.2 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 13.

Figure 13. Internal LED Indicators

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	ON - receiving SES identify command OFF - NOT receiving SES identify command
HDD ID	HDD internal to enclosure	Blue	ON - receiving SES identify command OFF - NOT receiving SES identify command

SECTION 2

InfiniteStorage 6120

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 InfiniteStorage 6120 system. Each step is explained in detail in Section 2.2 and Section 2.3.

1. Unpack the InfiniteStorage 6120 system.
2. Rackmount the InfiniteStorage 6120 (and expansion enclosure). (See Appendix A for physical dimensions and weight of the InfiniteStorage 6120.)
3. Connect the InfiniteStorage 6120 to your Fibre Channel (FC) switch, host computer(s), and expansion enclosure.
4. Power up the system.
5. Configure the storage array (create storage pools, virtual disks, and spare pools) via RS-232 interface or SSH.
6. Initialize the storage pools for use with your server/client systems. Partition disk space and create filesystems as needed.

2.2 Setting Up the InfiniteStorage 6120

This section details the installation of the hardware components of the InfiniteStorage 6120 system.

If your controller and disk enclosure configuration is delivered already installed in the rack, it must be removed from the shipping pallet using a minimum of 4 people. The racked unit may not be tipped more than 10 degrees, either from a level surface or rolling down an incline (such as a ramp).

NOTE : Follow the safety guidelines for rack installation given in Appendix B “Safety Requirements for Rack Installation”.

2.2.1 Unpacking the InfiniteStorage 6120

Wear an ESD wrist strap or otherwise ground yourself when handling InfiniteStorage 6120 modules and components. Electrostatic discharge can damage the circuit boards.

Before you unpack your InfiniteStorage 6120, 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 InfiniteStorage 6120 chassis and all components for signs of damage. If you detect any problems, contact SGI Technical Support immediately.

Your InfiniteStorage 6120 ships with the following:

- InfiniteStorage 6120 controller chassis with integrated controllers
- Two (2) power cables
- One (1) mini-jack serial cable for initial configuration
- One (1) enclosure faceplate
- Rack-mounting hardware

2.2.2 Rack-Mounting the InfiniteStorage 6120 Chassis

The InfiniteStorage 6120 is designed to fit within a 4U rack space. Follow these steps to mount the system in your rack:

1. Using eight (8) 10-32×0.188" Phillips panhead screws for each slider, attach the four (4) inner sliders to the chassis as shown in Figure 14 below.

Figure 14. Attach 4 Inner Sliders to Chassis

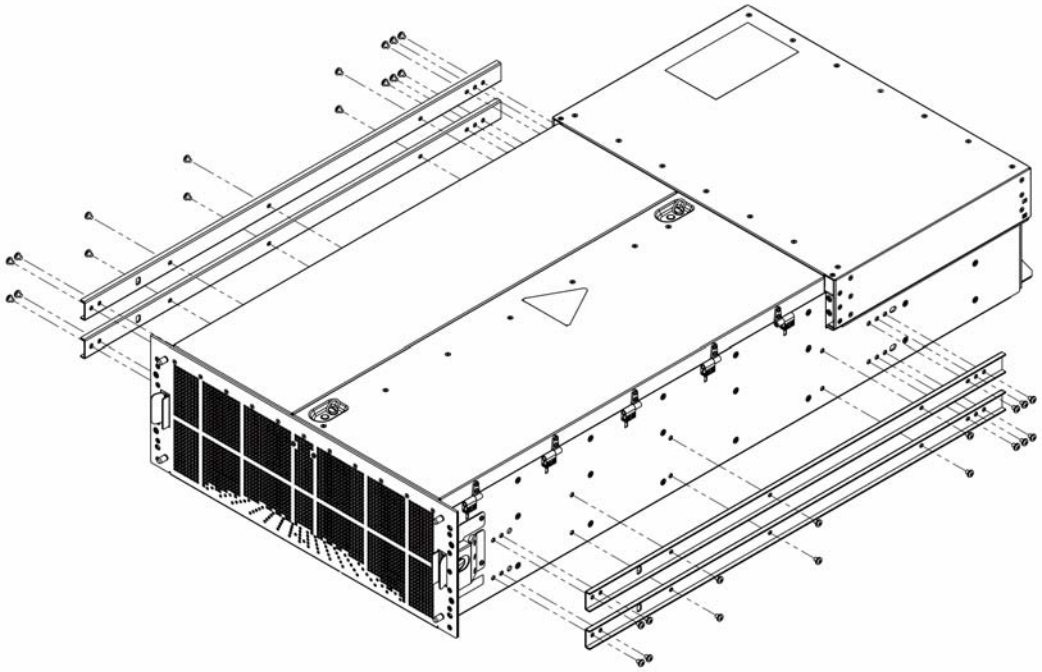

2. For expansion enclosure, skip this step and proceed to the next step.
Loosen the nine (9) screws on the rail kit and set the length to fit the depth of your rack (Figure 15). Tighten all the screws once the rail kit is set to the desired length. Repeat for the other rail kit.

Figure 15. Adjust Rail Kit Length

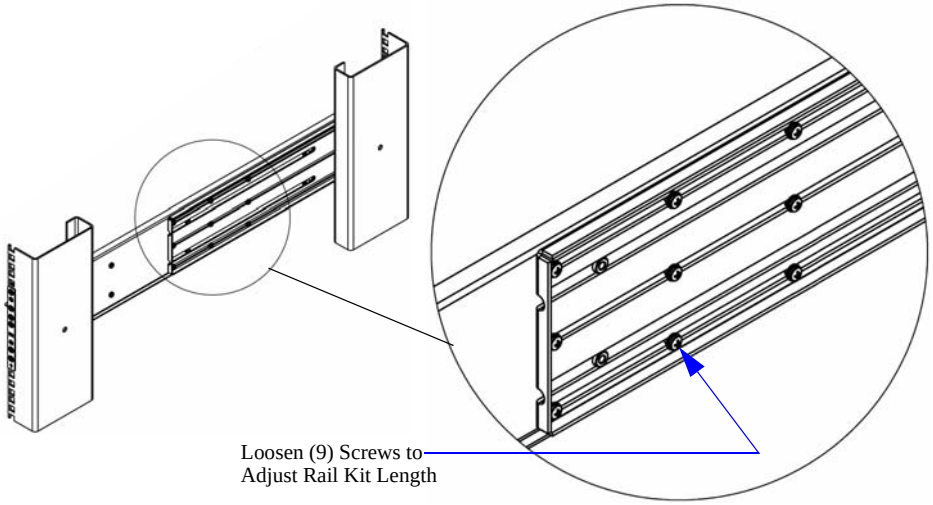

3. Attach the two (2) rail kits to your rack as shown in Figure 16. Make sure you leave enough space for the unit—bottom of screw plate aligns to bottom of InfiniteStorage 6120 chassis.
   At the front, use one (1) screw plate, one (1) shelf locking bracket (installed from behind), and six (6) #10 screws.
   For the rear end, use one (1) screw plate and four (4) #10 screws.

4. If the depth of your rack is more than 25.5 inches (647 mm), attach the two (2) cable management arm (CMA) brackets to the rear end of the rail kits (Figure 16).

Figure 16. Attach Rail Kits to Rack

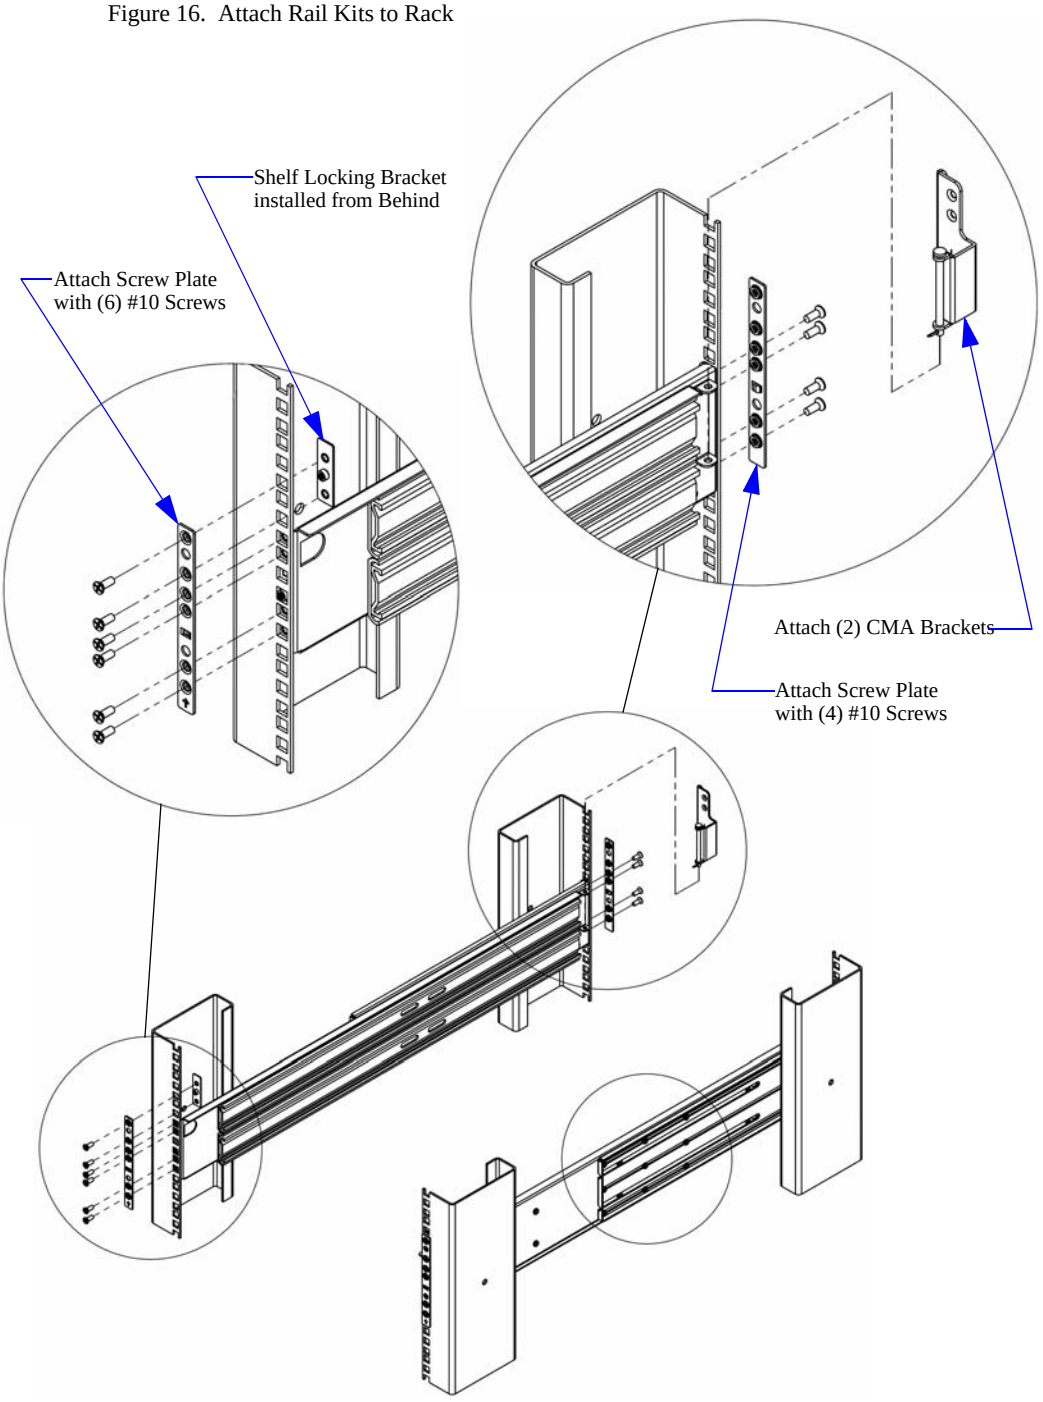

5. Using two (2) 10-32×0.50" (100d) Phillips flathead screws for each CMA, attach the CMA to the CMA brackets (Figure 17). If the depth of your rack is 25.5inches (647mm) or less, attach the CMA directly to the rack, above the rail kits. For square-hole rack, add 100d #10 shoulder washers.

6. With the help of a partner, hold the chassis close to the rail kits. Align the rear flanges on chassis with the rail kit flanges (Figure 17) and slide the chassis into rack until the sliders are engaged (slide in about 12 inches until stopped). Then slide the chassis out until locked. Depress the slider locks and proceed to slide the chassis all the way in.

7. Secure the two CMA to rear of chassis with two (2) 8-32×0.25" Phillips panhead screws.

Figure 17. Install Chassis to Rack

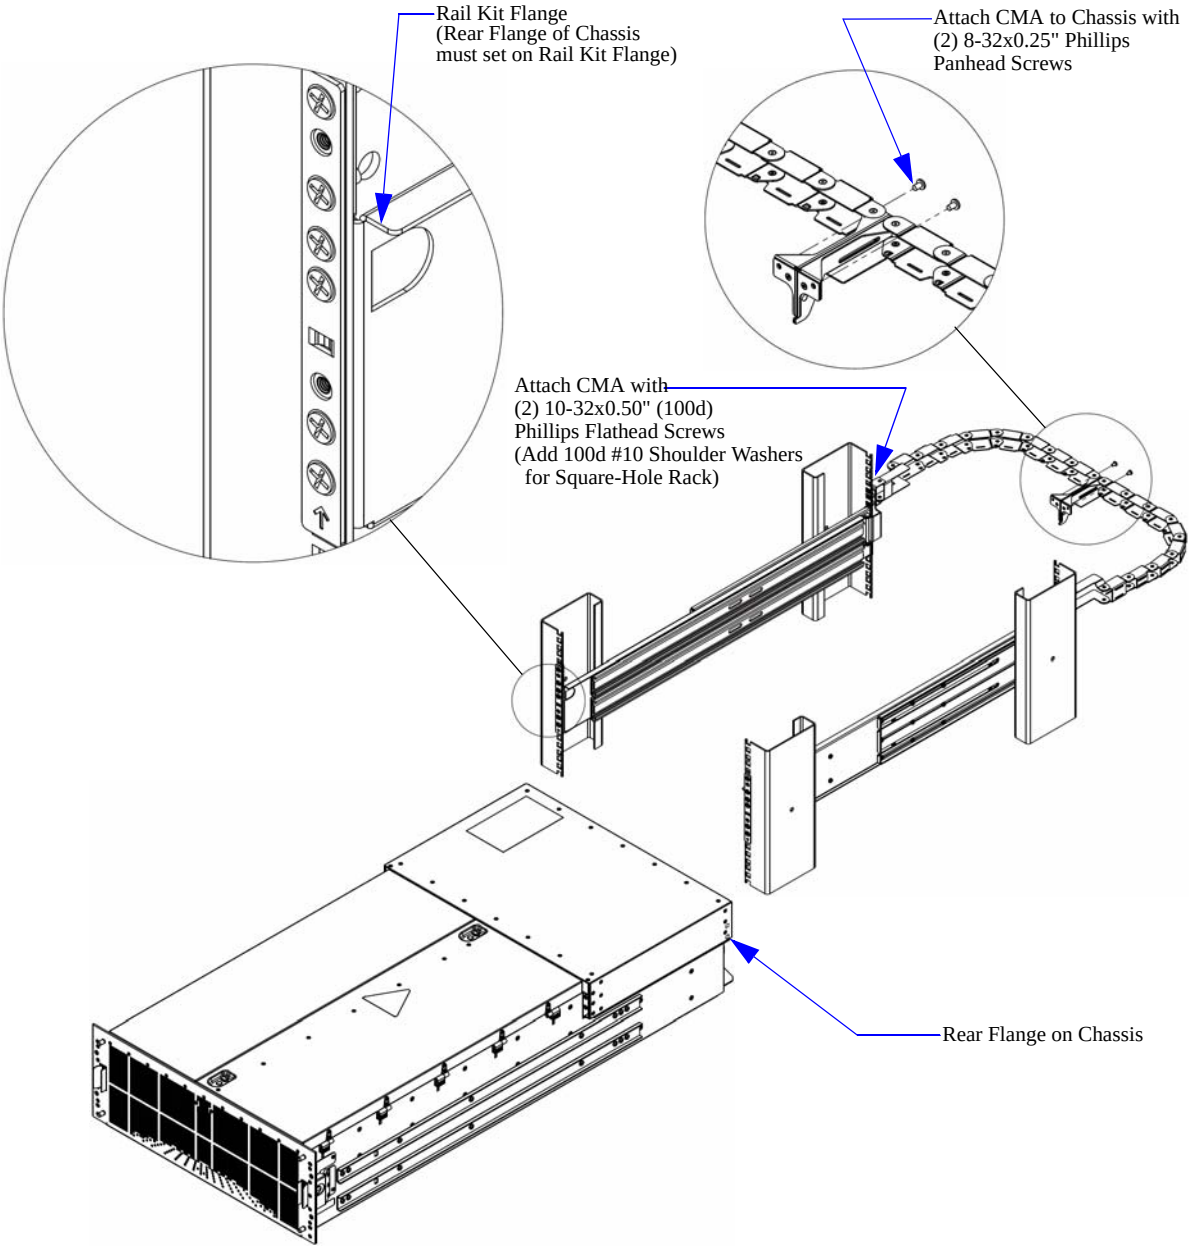

2.2.3 Connecting the InfiniteStorage 6120

2.2.3.1 Host Connections

Each InfiniteStorage 6120 supports up to four Fibre Channel host connections, two per controller. Each connection is auto-sensing and supports 2Gb/s, 4Gb/s or 8Gb/s. You may connect more than four hosts to the unit using a Fibre Channel switch (not provided).

The host ports are numbered as shown in Figure 18 below. Connect your switch or host system(s) to these ports.

Figure 18. A Controller

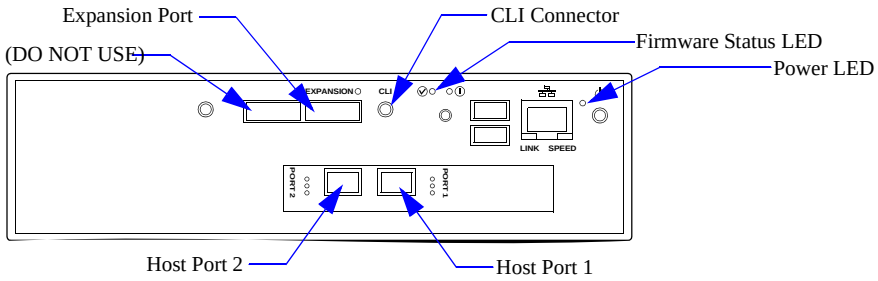

2.2.3.2 Expansion Enclosure Connections

A expansion enclosure can be daisy-chained to the InfiniteStorage 6120 to increase storage capacity. Follow these steps to connect an expansion enclosure:

NOTE : Only a 60-drive expansion enclosure can be used with the InfiniteStorage 6120 for storage expansion.

---

1. Using the supplied mini-SAS connector, connect one end of the cable to the port labeled **EXPANSION** on the InfiniteStorage 6120 controller (Figure 18).
2. Connect the other end of the cable to the port labeled **EXP2** on the expansion enclosure (Figure 19).

Figure 19. I/O Module on Expansion Enclosure

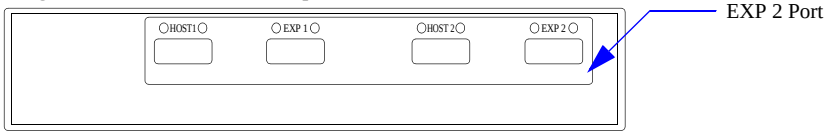

3. Repeat for the second controller (Figure 20).

Figure 20. Daisy-Chaining

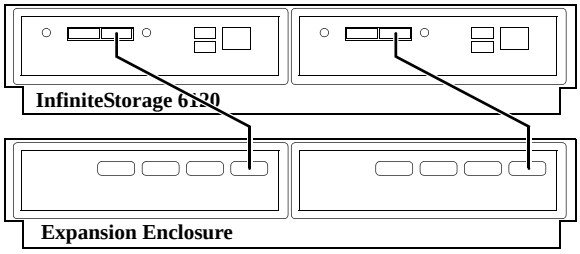

2.2.3.3 Power Connections

Each InfiniteStorage 6120 (and expansion enclosure) ships with two fully redundant power supplies that require 220 V connections. Only the power cables supplied with the system should be used. The power supplies are accessible from the rear of the unit.

1. Verify that the power switch on both supplies is in the OFF position (Figure 21).
2. Connect the power cables from the InfiniteStorage 6120 to the appropriate 220 V outlets. For maximum redundancy, connect the two power supplies to different AC circuits.
3. Similarly, connect the power cables to the expansion enclosure.

Figure 21. Power Cooling Module

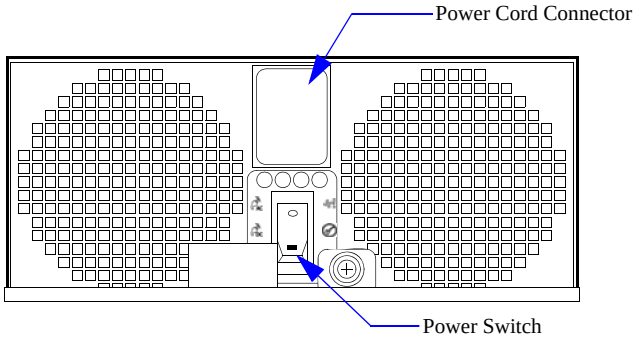

2.2.4 Powering On the InfiniteStorage 6120 (and Expansion Enclosure)

1. On the InfiniteStorage 6120, turn the switch on each power supply to the ON position.
2. Verify that the Power LED on each controller turns green within a few seconds (Figure 22), indicating a successful power application.

Figure 22. A Controller

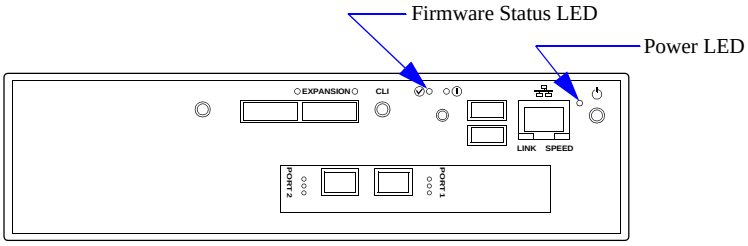

3. Verify that the Firmware Status LED blinks green within one minute of power application (Figure 22), indicating that the InfiniteStorage 6120 firmware is currently loading.
4. Verify that the Firmware Status LED stays solid green within two minutes of power application, indicating that the InfiniteStorage 6120 is ready to be accessed.
5. If an expansion enclosure is connected, turn the switch on each power supply to the ON position.

You may now configure the system as described in Section 2.3 "Configuring the InfiniteStorage 6120".

2.3 Configuring the InfiniteStorage 6120

This section provides information on configuring your InfiniteStorage 6120.

NOTE : The configuration examples provided here represent only a general guideline. These examples should not be used directly to configure your particular InfiniteStorage 6120.

The CLUI (command line user interface) commands used in these examples are fully documented in Sections 3.1 through 3.5—though exact commands may change depending on your firmware version. To access the most up-to-date commands, use the CLUI's online HELP feature.

2.3.1 Planning Your Setup and Configuration

Before proceeding to configure the storage settings for the InfiniteStorage 6120, it is necessary to understand the basic organization of the system.

The InfiniteStorage 6120 uses storage pools and virtual disks to configure drive storage for use by host systems. A storage pool (sometimes referred to as a RAID group) is a collection of 5 to 10 disk drives, ideally of the same capacity and type. In RAID5, the capacity of one drive is reserved for parity, allowing data recovery in the event of a single drive loss in the storage pool. In RAID6, the capacity of two drives is reserved for parity, allowing data recovery if either one or two drives are down in a storage pool.

A single RAID5 storage pool can be configured using 5 or 9 drives. A single RAID6 storage pool can be configured using 6 or 10 drives. The InfiniteStorage 6120 supports up to 24 storage pools in a single system. A storage pool can be configured with drives of different capacity, in which case the InfiniteStorage 6120 will use the lowest capacity for all drives, but this is not recommended.

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. You can create up to 64 virtual disks per storage pool.

NOTE : In this document, the terms *GB* and *GiB* are used interchangeably and denote 1,073,741 bytes (2^{30} bytes).

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

- Unassigned pool - By default, all the drives or replacement drives are initially allocated to this pool. When a drive 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 5 to 10 disk drives, ideally of the same capacity and type.
- Spare pool - This type of pool contains drive(s) that can be used as spare drives in one or more storage pools.

2.3.2 *Drive Layout*

Figure 23 shows the layout of drives in an InfiniteStorage 6120. Drives are always referenced by enclosure ID and slot number as illustrated.

The InfiniteStorage 6120 supports mixing SAS and SATA drives within the same enclosure. For most efficient airflow and reduced vibrations, SAS drives should be placed in the front of the enclosure (lower slot numbers) while SATA drives should be placed in the rear of the enclosure (higher slot numbers).

Figure 23. InfiniteStorage 6120 Drive Layout

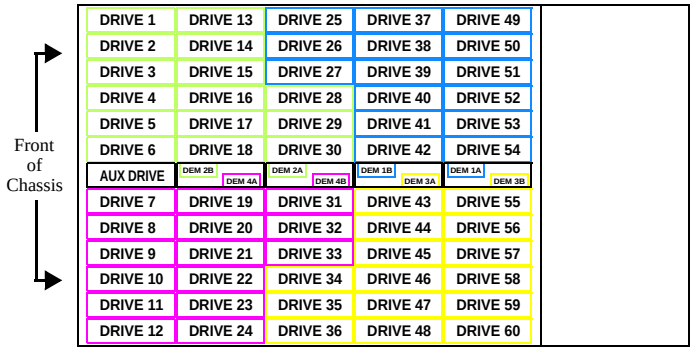

2.3.3 *Initial Configuration via RS-232*

The InfiniteStorage 6120 can be configured and administered either via serial connection (using the supplied mini-jack 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 in the next subsection.

2.3.3.1 Serial Interface Configuration

1. Connect the supplied mini-jack cable to the port labeled **CLI** on Controller 1 (Figure 24). Connect the other end of the cable to a serial port on a standard PC.

Figure 24. CLI Port on Controller

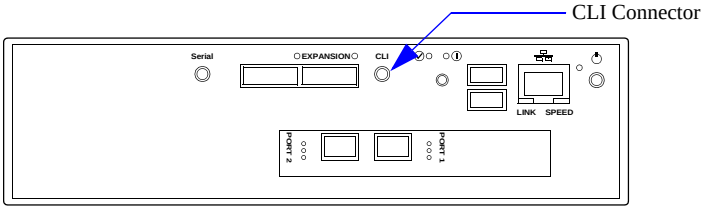

2. Load a serial console program (such as HyperTerminal, minicom, or screen) and use the following settings for the serial connection:
   - Bits per second: 115,200
   - Data bits: 8
   - Parity: None
   - Stop bits: 1
   - Flow control: None
3. Once connected, hit the <Enter> key to bring up the password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

2.3.3.2 Configure Network Settings

To enable SSH access to each controller, the network settings must be configured as follows:

1. Enter command: **`UI SET NETWORK_INTERFACE LOCAL 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 25).

   Figure 25. Set Network Interface Example

```
RAID[0]$ ui set network_interface local 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)
```

2. To verify the new settings, enter command (Figure 26):
   `UI SHOW NETWORK_INTERFACE LOCAL *`

   Figure 26. Show Network Interface Settings

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

3. Switch the serial cable to the **CLI** port on Controller 2 and repeat the above steps to enter a different IP address for Controller 2.

4. Enter command **`SHUTDOWN SUBSYSTEM RESTART`** to restart the controllers and make the changes take effect.

2.3.4 Setting System Time & Date

The system time and date for the InfiniteStorage 6120 are factory-configured for the U.S. Pacific Standard Time (PST) zone. If you are located in a different time zone, you need to change the settings so that the time stamps for all events are correct.

To display the current system time (Figure 27), enter command **`SHOW SUBSYSTEM ALL_ATTRIBUTES`**

Figure 27. Show Current System Time

```
RAID[0]$ show subsystem all_attributes
RP Subsystem Name:
UID:                    0x0000000000000000
Subsystem Time:         Fri Jan 30 16:48:55  2009
Time zone:              0 minutes
DST offset:             0 minutes
Locate Dwell Time:      120 seconds
Mirror Synch:           NA
```

To change the date and time to March 1, 2009 2:15:32 pm, for example, type:
`SET SUBSYSTEM DATE_AND_TIME=2009:3:1:14:15:32`

2.3.5 Checking Status of Physical Disks

Before creating any storage pools, check the status of all the drives.

1. You can use the following commands:
SHOW UNASSIGNED_POOL -ALL (Figure 28)
SHOW PHYSICAL_DISK * (Figure 29)

Figure 28. Show Unassigned Pool Example Screen

```
RAID$ show unassigned_pool -all
  OID:              0x1800ffff
  Type:             UNASSIGNED
  Capacity:         42270720 MBs
  Failed Capacity:  0 MBs
  Total Phy Disks   60
```

Figure 29. Show Disk Information Screen

```
RAID[0]$ show physical_disk *

                                                                                                                  |Member|
Encl|Slot| Vendor |     Product ID     |Type|Cap GB |RPM|Revision|    Serial Number   |Pool|Status| Idx|Status|       WWN       |
-----------------------------------------------------------------------------------------------------------------------------------
   1    1  HITACHI      HDS725050KLA360 SATA     465  7K K2AOAD1A       KRVN65ZAKAZLKH UNAS  GOOD    57         5000cca20def4516
   1    2  HITACHI      HDS725050KLA360 SATA     465  7K K2AOAD1A       KRVN65ZAKNL79H UNAS  GOOD    32  NORM   5000cca20df3a263
   1    3  HITACHI      HDS725050KLA360 SATA     465  7K K2AOAD1A       KRVP65ZAKD8Y5F UNAS  GOOD    42  NORM   5000cca20defdc87
   1    4  HITACHI      HDS725050KLA360 SATA     465  7K K2AOAD1A       KRVN65ZAK72BLH UNAS  GOOD     1  NORM   5000cca20ded7e05
   1    5  HITACHI      HDS725050KLA360 SATA     465  7K K2AOAD1A       KRVN65ZAKB269H UNAS  GOOD    38  NORM   5000cca20def4edd
   1    6  HITACHI      HDS725050KLA360 SATA     465  7K K2AOAD1A       KRVN65ZAKB1LHH UNAS  GOOD    23  NORM   5000cca20def4c96
   1    7  Hitachi      HUA721075KLA330 SATA     698  7K GK8OA70M       GTF200P8G6K2YF UNAS  GOOD    40  NORM   5000cca215c2fa66
   1    8  HITACHI      HDS725050KLA360 SATA     465  7K K2AOAD1A       KRVN65ZAKA6P9H UNAS  GOOD    41  NORM   5000cca20deeeb53
   1    9  HITACHI      HDS725050KLA360 SATA     465  7K K2AOAD1A       KRVN65ZAKAGVYF UNAS  GOOD    47         5000cca20def0a0a
   1   10  HITACHI      HDS725050KLA360 SATA     465  7K K2AOAD1A       KRVN65ZAKA5SDH    2  GOOD    45  NORM   5000cca20deee7d3
   1   11  HITACHI      HDS725050KLA360 SATA     465  7K K2AOAD1A       KRVN65ZAK8VH8F    2  GOOD    39  NORM   5000cca20dee4cee
   1   12  HITACHI      HDS725050KLA360 SATA     465  7K K2AOAD1A       KRVN65ZAKAZ75H    1  GOOD    26  NORM   5000cca20def43b5
   1   13  Hitachi      HUA721075KLA330 SATA     698  7K GK8OA70M       GTF202P8GTMNRF    4  GOOD    20  NORM   5000cca215cb32bd
   1   14  Hitachi      HUA721075KLA330 SATA     698  7K GK8OA70M       GTF202P8GTLDJF    4  GOOD    33  NORM   5000cca215cb2dfe
   1   15  Hitachi      HUA721075KLA330 SATA     698  7K GK8OA70M       GTF202P8GTPR7F    3  GOOD     3  NORM   5000cca215cb3a6e
   1   16  Hitachi      HUA721075KLA330 SATA     698  7K GK8OA70M       GTF202P8GTPW9F    5  GOOD    50  NORM   5000cca215cb3b0b
   1   17  Hitachi      HUA721050KLA330 SATA     465  7K GK6OAB0A       GTE402P6G94E8E    1  GOOD    27  NORM   5000cca214c427f6
   1   18  Hitachi      HUA721075KLA330 SATA     698  7K GK8OA70M       GTF202P8GTPNXF UNAS  GOOD    58         5000cca215cb3a45
   1   19  Hitachi      HUA721075KLA330 SATA     698  7K GK8OA70M       GTF202P8GTPTBF    4  GOOD    18  NORM   5000cca215cb3ab0
   1   20  HITACHI      HDS725050KLA360 SATA     465  7K K2AOAD1A       KRVN63ZAKSMX8D    0  GOOD    21  NORM   5000cca20df505cb
   1   21  HITACHI      HDS725050KLA360 SATA     465  7K K2AOAD1A       KRVP65ZAKERMGF    0  GOOD    17  NORM   5000cca20df08447
   1   22  Hitachi      HUA721050KLA330 SATA     465  7K GK6OAB0A       GTE402P6G9M0SE    1  GOOD    28  NORM   5000cca214c45ec1
   1   23  HITACHI      HDS725050KLA360 SATA     465  7K K2AOAD1A       KRVN63ZAKNTM6D    1  GOOD    35  NORM   5000cca20df3ba5a
   1   24  HITACHI      HDS725050KLA360 SATA     465  7K K2AOAD1A       KRVN65ZAKKVV3F    1  GOOD    30  NORM   5000cca20df26595
   1   25  HITACHI      HDS725050KLA360 SATA     465  7K K2AOAD1A       KRVN35ZAKE8HZH    0  GOOD     5  NORM   5000cca20df04f4d
   1   26  HITACHI      HDS725050KLA360 SATA     465  7K K2AOAD1A       KRVP65ZAKDW9PH    0  GOOD     6  NORM   5000cca20df0217e
   1   27  HITACHI      HDS725050KLA360 SATA     465  7K K2AOAD1A       KRVN63ZAKSNNDD    0  GOOD     8  NORM   5000cca20df50898
   1   28  HITACHI      HDS725050KLA360 SATA     465  7K K2AOAD1A       KRVN65ZAKAZ4LH    0  GOOD     9  NORM   5000cca20def4365
   1   29  HITACHI      HDS725050KLA360 SATA     465  7K K2AOAD1A       KRVN65ZAKB2GEH    2  GOOD    37  NORM   5000cca20def4fd9
   1   30  HITACHI      HDS725050KLA360 SATA     465  7K K2AOAD1A       KRVP65ZAKESH1F    1  GOOD    31  NORM   5000cca20df0877f
   1   31  HITACHI      HDS725050KLA360 SATA     465  7K K2AOAD1A       KRVN03ZAGMUMUD UNAS  GOOD    49         5000cca20dc9034c
   1   32  HITACHI      HDS725050KLA360 SATA     465  7K K2AOAD1A       KRVN35ZAK6VLBF    1  GOOD    25  NORM   5000cca20ded6490
   1   33  HITACHI      HDS725050KLA360 SATA     465  7K K2AOAD1A       KRVN65ZAKN0S7F    1  GOOD    36  NORM   5000cca20df360bf
   1   34  HITACHI      HDS725050KLA360 SATA     465  7K K2AOAD1A       KRVN65ZAK9VNRH    2  GOOD    44  NORM   5000cca20deec1f6
   1   35  HITACHI      HDS725050KLA360 SATA     465  7K K2AOAD1A       KRVN65ZAKL1MYH    0  GOOD    22  NORM   5000cca20df27b5c
   1   36  HITACHI      HDS725050KLA360 SATA     465  7K K2AOAD1A       KRVN03ZAGH9Y4D    2  GOOD    43  NORM   5000cca20dc6f6e1
   1   37  Hitachi      HUA721075KLA330 SATA     698  7K GK8OA70M       GTF202P8GTSLNF    3  GOOD     7  NORM   5000cca215cb4181
   1   38  Hitachi      HUA721075KLA330 SATA     698  7K GK8OA70M       GTF202P8GTRPVF    5  GOOD    54  NORM   5000cca215cb3e23
   1   39  Hitachi      HUA721075KLA330 SATA     698  7K GK8OA70M       GTF202P8GTSMGF    3  GOOD    14  NORM   5000cca215cb419a
   1   40  Hitachi      HUA721075KLA330 SATA     698  7K GK8OA70M       GTF202P8GTK4HF    5  GOOD    56  NORM   5000cca215cb2944
   1   41  Hitachi      HUA721075KLA330 SATA     698  7K GK8OA70M       GTF202P8GTRPDF    5  GOOD    52  NORM   5000cca215cb3e15
   1   42  Hitachi      HUA721075KLA330 SATA     698  7K GK8OA70M       GTF202P8GTRLVF    4  GOOD    24  NORM   5000cca215cb3dc6
   1   43  Hitachi      HUA721075KLA330 SATA     698  7K GK8OA70M       GTF202P8GTNPYF    3  GOOD    13  NORM   5000cca215cb36a4
   1   44  Hitachi      HUA721075KLA330 SATA     698  7K GK8OA70M       GTF202P8GTPWUF    5  GOOD    53  NORM   5000cca215cb3b1b
   1   45  Hitachi      HUA721075KLA330 SATA     698  7K GK8OA70M       GTF202P8GTPSHF    3  GOOD    12  NORM   5000cca215cb3a95
   1   46  Hitachi      HUA721075KLA330 SATA     698  7K GK8OA70M       GTF202P8GTPKNF    4  GOOD    29  NORM   5000cca215cb39e0
   1   47  Hitachi      HUA721050KLA330 SATA     465  7K GK6OAB0A       GTE402P6G94EYE    2  GOOD    46  NORM   5000cca214c4280b
   1   48  Hitachi      HUA721075KLA330 SATA     698  7K GK8OA70M       GTF202P8GTRP6F    5  GOOD    51  NORM   5000cca215cb3e0f
   1   49  Hitachi      HUA721075KLA330 SATA     698  7K GK8OA70M       GTF202P8GTV0KF    4  GOOD    15  NORM   5000cca215cb4a93
   1   50  Hitachi      HUA721075KLA330 SATA     698  7K GK8OA70M       GTF202P8GTUYLF    3  GOOD     2  NORM   5000cca215cb4a56
   1   51  Hitachi      HUA721075KLA330 SATA     698  7K GK8OA70M       GTF202P8GTSJZF UNAS  GOOD    59         5000cca215cb414d
   1   52  Hitachi      HUA721075KLA330 SATA     698  7K GK8OA70M       GTF202P8GTNV9F    5  GOOD    55  NORM   5000cca215cb372b
   1   53  Hitachi      HUA721075KLA330 SATA     698  7K GK8OA70M       GTF202P8GTTZ5F    4  GOOD    19  NORM   5000cca215cb46a7
   1   54  Hitachi      HUA721075KLA330 SATA     698  7K GK8OA70M       GTF202P8GTR3ZF    4  GOOD    34  NORM   5000cca215cb3bf9
   1   55  Hitachi      HUA721075KLA330 SATA     698  7K GK8OA70M       GTF202P8GTT3DF    3  GOOD     4  NORM   5000cca215cb4369
   1   56  Hitachi      HUA721075KLA330 SATA     698  7K GK8OA70M       GTF202P8GTPGWF    3  GOOD    10  NORM   5000cca215cb398a
   1   57  Hitachi      HUA721075KLA330 SATA     698  7K GK8OA70M       GTF202P8GTPH3F    3  GOOD    11  NORM   5000cca215cb3991
   1   58  Hitachi      HUA721075KLA330 SATA     698  7K GK8OA70M       GTF202P8GTSMLF    4  GOOD    16  NORM   5000cca215cb419e
   1   59  Hitachi      HUA721075KLA330 SATA     698  7K GK8OA70M       GTF202P8GTSK6F UNAS  GOOD    60         5000cca215cb4154
   1   60  Hitachi      HUA721075KLA330 SATA     698  7K GK8OA70M       GTF202P8GTPATF    5  GOOD    48  NORM   5000cca215cb390b

Total Physical Disks:      60
Total Assigned Disks:      54
Total Unassigned Disks:    6
  Total SATA Disks:        60
  Total Member State Normal:   54
```

2. Verify that all the drives are present and healthy. If the drives are not visible, type **SHUTDOWN SUBSYSTEM RESTART** at the CLUI prompt to reboot both controllers.

   If there is a drive missing, enter command **SHOW PHYSICAL_DISK * ALL_ATTRIBUTES**. Then search through the list to find out which slot is not reporting (Figure 30).

Figure 30. Physical Disk List Example Screen

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

  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:     GOOD
  Rotation Speed:   7200 RPM
  Device Type:      SATA
  Member State:     UNASSIGNED
  Spare:            FALSE
  Failed:           FALSE
  UUID:             0x5000cca215c564560
   .
   .
   .
```

2.3.6 Installing Licensed Features

Before creating storage pools and virtual disks, you should install any licensed features (such as RAID 6 and SATAssure) and activate the associated licenses.

License key generation requires that you have your feature license number, JBOD enclosure serial number, and controller ID information (Figure 31). To complete license key generation, go to the SGI licensing site, http://www.sgi.com/support/licensing/special.html.

See the SGI customer letter that accompanies your feature license order for licensing instructions.

Figure 31. Installing Licensed Features

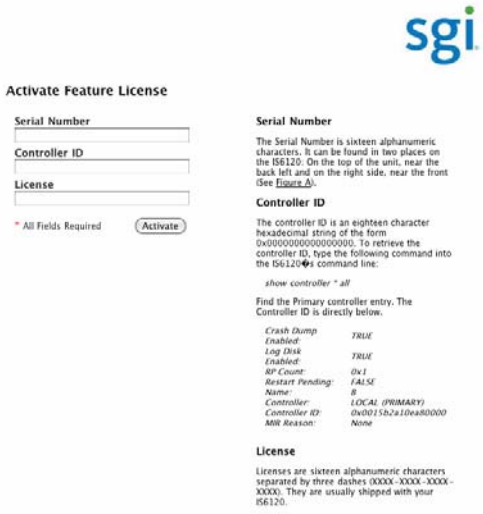

2.3.7 Creating Storage Pools

A storage pool on a InfiniteStorage 6120 has the following attributes:

- RAID level
  Storage pools can be configured to use either a RAID5 or RAID6 parity scheme. In RAID5, the capacity of one drive is reserved for parity, allowing data recovery in the event of a single drive loss in the storage pool. In RAID6, the capacity of two drives is reserved for parity, allowing data recovery if either one or two drives are down in a storage pool. ***For maximum data protection, use RAID6***.

- Chunk size
  The chunk size defines the amount of data written to a single disk before proceeding to the next disk in the storage pool. For optimal performance, the chunk size should be proportional to the expected host request size. For example, a large chunk size should be used when transferring large blocks of data while a small chunk size should be used when transferring multiple small blocks of data. The default chunk size is 128K.

- Assignment policy
  The assignment policy defines whether drives with the same interface type (SAS or SATA), spindle speed, or capacity are placed within the same pool. If there are not enough drives of the defined type, then the **CREATE POOL** command fails with a not-enough-drives message.

  Once an assignment policy is used to create a pool, that policy also governs how spares are chosen for that pool. Only drives of the same defined type will be used for sparing on that pool. If there are no drives of the defined type available in the spare pool, none will be used.

  Note : To change the assignment policy after the fact, use the **SET POOL** *command to change the desired attributes.*

- Drive count
  A RAID5 storage pool may consist of 5 or 9 drives. A RAID6 storage pool may consist of 6 or 10 drives. For maximum performance, storage pools should be created using drives 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 [CHUNK_SIZE=<value>][ASSIGN_POLICY=([SAS|SATA], 15000RPM|10000RPM|7200RPM|5400RPM|0RPM|SSD,[<*capacity-raw*>GB])] [NUMBER=<*n*>] or **PHYSICAL_DISK=<*id1*>,...,<*idn*>] [RAID_LEVEL=RAID5|RAID6|RAID-SH]**

   You may either explicitly select the disks for the pool by specifying their index names with the **PHYSICAL_DISK** parameter or specify the number of disks in the pool by using **NUMBER** parameter, in which case the next available disks will be selected (Figure 32).

   Figure 32. CREATE POOL Example (1)

```
RAID[0]$ create pool raid_level=raid5 chunk_size=64kb physical_disk=0x003e 0x003f
0x0040 0x0041 0x0042
POOL 0 OID=0x19b60000 create STATUS='Success' (0x0)

RAID$ 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 33), only the selected type of drives will be used for the storage pool and you will not need to enter the individual disk name.

   Figure 33. CREATE 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.

---

3. To view the list of configured storage pools, enter command **SHOW POOL *** (Figure 34).

Figure 34. Show Pool Information Screen

```
                                          |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
  4 pool-4         NORMAL       128   5             6192      0      0   W   R F I          10  UNASSIGNED  AUTO
  5 pool-5         NORMAL       128   5             6192      0      0   W   R F I I        10  UNASSIGNED  AUTO

Total Storage Pools: 6
```

Legend for the `Settings` field (Figure 34):
`S` - SATAssure
`W` - Write-back cache
`M` - Mirrored write-back cache
`R` - Read ahead
`F` - Full-stripe cache
`I` - IO routing

SHOW POOL <id> ALL_ATTRIBUTES will display detailed information of the storage pool (Figure 35).

Figure 35. SHOW POOL ALL_ATTRIBUTES Example Screen

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

  OID:                  0x19b40000
  Type:                 STORAGE
  Name:                 0x19b40000
  Chunk Size:           64KB (0x80 blocks)
  Block Size:           512
  RAID Type:            RAID5
  Free Raid5 Capacity:  2801664 MBs
  Max VD Size:          2801664 MBs
  Total Capacity:       3522560 MBs
  UUID:                 0x00
  Global Spare Pool:    UNASSIGNED
  DiskTimeout(FRT):     10 minutes
  Init Policy:          NO IO permitted
  Init Priority:        50%
  Full Rebuild Priority:        80%
  Fractional Rebuild Priority:  90%
  Sparing Policy:       AUTOMATIC
  Assign Policy:
    Device Type:        NA
    Rotation Speed:     NA
    Raw Capacity:       NA
  SATAssure:            FALSE
  IO ROUTING:           TRUE
  WBC:                  TRUE
  MWBC:                 FALSE
  Initializing:         TRUE
  Rebuilding:           FALSE
  Paused:               FALSE
  AutoWriteLock:        FALSE
  Critical:             FALSE
  Current Home:         0x0015b2a122b20000  0x00000000
  Future Home:          0xffffffffffffffff  0x00000000
  Preferred Home:       0x0015b2a122b20000  0x00000000
  BkgdJob OID:          0x28000003
  Total Phy Disks       5
  State:                NORMAL
  Member Size:          704512 MB
    pID     State     UUID
    0x003e  NORM      0x5000cca215c56e47
    0x003f  NORM      0x5000cca215c56456
    0x0040  NORM      0x5000cca215c54c71
    0x0041  NORM      0x5000cca215c5675c
    0x0042  NORM      0x5000cca215c56e55
```

2.3.8 Creating Virtual Disks

A virtual disk can be created on just a part of a storage pool or uses all of the capacity of a single storage pool.

To create a virtual disk, use the **CREATE VIRTUAL_DISK** command:

1. At the CLUI prompt, type:
   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 36).

   A message is displayed to indicate whether the virtual disk creation was successful.

   Figure 36. 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)
```

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.

NOTE : Virtual disk(s) on a storage pool are NOT accessible until its initialization is complete.

---

3. To view the list of configured virtual disks, type **SHOW VIRTUAL_DISK *** (Figure 37).

Figure 37. 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
  4 vd-4_4          READY     4  5     5480   W   I          0(L) 0   0(L) 0   None   INACTIVE
  5 vd-5_5         NOT RDY    5  5     5480   W   I   I      0(L) 0   1(R) 0  1(R) 0 0x28000005

Total Virtual Disks: 6
```

SHOW VIRTUAL_DISK <id> ALL_ATTRIBUTES will display detailed information of the virtual disk (Figure 38). In contrast, **SHOW POOL <id> ALL_ATTRIBUTES** will display detailed information of the storage pool (Figure 35).

Figure 38. SHOW VIRTUAL_DISK ALL_ATTRIBUTES Example

```
RAID[0]$ show virtual_disk 0 all_attributes

  OID:              0x89ba0000
  Name:             89ba0000
  POOL OID:         0x19b40000
  Capacity:         16384 MBs
  Offset:           0x0
  State:            NOT RDY
  Raidlevel:        RAID5
  IO ROUTING:       TRUE
  WBC:              TRUE
  MWBC:             FALSE
  Initializing:     TRUE
  Paused:           FALSE
  AutoWriteLock:    FALSE
  Critical:         FALSE
  Present Home Only:FALSE
  Current Home:     0x0015b2a122b20000  0x00000000
  Future Home:      0xffffffffffffffff  0x00000000
  Preferred Home:   0x0015b2a122b20000  0x00000000
  BkgdJob OID:      0x28000000
  UUID:             60001ff0722b20000000000089ba0000
```

2.3.9 Creating Spare Pools

The InfiniteStorage 6120 supports the concept of a 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 drive failure, a drive that is assigned to a spare pool is automatically swapped into the storage pool to replace the failed drive. A rebuild then occurs immediately, reducing the possibility of data loss due to additional drive failures.

Spare pools can be shared by storage pools or dedicated to a particular , 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 drive(s) from the unassigned pool and assign it to the spare pool.

NOTE : For maximum performance, create spare pools using drives with the same characteristics as those drives within the storage pool(s) for which they are designated to spare. The characteristics include items like SAS/SATA, capacity, and RPM.

To create a spare pool:

1. At the CLUI prompt, type: **`CREATE SPARE_POOL`** (Figure 39)

   A message is displayed to indicate whether the new spare pool creation was successful.

   Figure 39. CREATE SPARE POOL Example Screen

   ```
   RAID[0]$ create spare_pool

   SPARE POOL 6 OID=0x19bc0006 create STATUS='Success' (0x0)
   ```

2. Assign disk(s) to the spare pool, enter command:
   `ASSIGN PHYSICAL_DISK <disk-id> TO_POOL <pool-id>`
   where <disk-id> is the index name of the disk to be added.

   To view the list of unassigned disks, enter command **`SHOW UNASSIGNED_POOL PHYSICAL_DISK`**

   A message is displayed to indicate whether the disk has been added to the spare pool (Figure 40).

Figure 40. Assign Physical Disk to Spare Pool Example Screen

```
RAID[0]$ assign physical_disk 0x004e to_pool 6

PHYSICAL_DISK 78 OID=0x20b6004e assigned to POOL 6 OID=0x19bc0006STATUS='Success' (0x0)
```

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 command **ASSIGN PHYSICAL_DISK <disk-id> TO_POOL 0x1800ffff**

   To delete a spare pool, enter command **DELETE SPARE_POOL <pool-id>**

4. To view the list of configured spare pool(s), enter command **SHOW SPARE_POOL *** (Figure 41).

Figure 41. Show Spare Pool Information Screen

```
RAID[0]$ show spare *

                              |Disk |Total |Total|Storage |
Idx|Name                  |Blocks| T/O |Cap GB| PDs |Pool Idx|
-------------------------------------------------------------
  6 spare_pool-6             512    10       0    0
  7 spare_pool-7             512    10       0    0

Total Spare Pools: 2
```

SHOW SPARE_POOL <id> ALL_ATTRIBUTES will display detailed information of the spare pool (Figure 42). **SHOW SPARE_POOL <id> PHYSICAL_DISKS** will display the list of disks in the pool.

Figure 42. Show Spare Pool Example Screen (2)

```
RAID[0]$ show spare_pool 6 all_attributes
  OID:              0x19bc0006
  Type:             GLOBAL SPARE
  Name:             19bc0006
  Block Size:       0x200
  DiskTimeout(FRT): 10 minutes
  Total Capacity:   1409024 MBs
  UUID:             0x00
  Total Phy Disks   2

RAID[0]$ show spare_pool 6 physical_disk
OID: 0x20b6004e Index: 0x004e
OID: 0x20b7004f Index: 0x004f
```

5. Assign spare pools to storage pools. Enter 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.

NOTE : Each storage pool must 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.

A message is displayed to indicate whether the spare pool has been designated (Figure 43). Repeat this step for each storage pool in the system.

Figure 43. 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)
```

6. Use the **SHOW POOL** command to view the new storage pool information (Figure 44).

Figure 44. Show Pool Example Screen

```
RAID[0]$ show pool 0 all_attributes
  OID:                  0x19b40000
  Type:                 STORAGE
  Name:                 0x19b40000
  Chunk Size:           64KB (0x80 blocks)
  Block Size:           0x200
  RAID Type:            RAID5
  Free Raid5 Capacity:  2752512 MBs
  Total Capacity:       3522560 MBs
  UUID:                 0x00
  Global Spare Pool:    0x19bc0006
  DiskTimeout(FRT):     10 minutes
  Init Policy:          NO IO permitted
   .
   .
   .
```

2.3.10 Initializing the Host System

When a storage pool is created, initialization begins automatically and will continue in the background until it is completed. Once the storage pool has completed its initialization, its virtual disk(s) can be accessed immediately.

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 45).

Figure 45. Show Job Example Screen

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

  OID:              0x28000000
  Target:           0x19b40000
  Type:             INITIALIZE
  Status:           RUNNING
  Priority:         50
  Fraction Complete:71%
```

You may also check if a virtual disk is ready for access using the **SHOW VIRTUAL_DISK <id> ALL_ATTRIBUTES** command (Figure 46).

Figure 46. 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
   .
   .
   .
```

2.3.10.1 Present Virtual Disk to External Host

A *presentation* on an InfiniteStorage 6120 has the following components:

- Discovered Initiators
  A *discovered initiator* is defined as a Fibre Channel initiator that has logged into the InfiniteStorage 6120. This is normally a single port on a Fibre Channel 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 initiators. The host attribute OS (operating system) allows you to 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 ports on the controller. The channel attribute MODE allows you to 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 InfiniteStorage 6120, 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. The following are attributes of a presentation:
  - PORT – 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 – LUN number (user-specified) that the virtual disk will show to the host.

To create a host:

Use the **`APPLICATION CREATE HOST [OSTYPE=GENERIC|LINUX|MAC_OS|WINDOWS]`** command where OSTYPE is the operating system. GENERIC is the default.

To delete a host:

Use the **`APPLICATION DELETE HOST=<object-id>`** command.

To display presentations associated with the specified application host:

Use the **APPLICATION SHOW HOST=<object-id>[PRESENTATIONS]** command.

To display the presentations associated with the specified application host:

Use the **APPLICATION SHOW HOST=<object-id>[INITIATORS]** command.

To create a presentation to a virtual disk for the specified host:

Use the **APPLICATION CREATE PRESENTATION VIRTUAL_DISK=<object-id> HOST=<object-id> | ALL** command.

To create an initiator object that is associated with the specified host:

Use the **APPLICATION IMPORT DISCOVERED_INITIATOR=<object-id> HOST=<object-id>** command.

To present the virtual disk to all controller ports for both controllers (promiscuous mode):

Use the **APPLICATION CREATE PREENTATION VIRTUAL_DISK=<object-id> HOST=ALL** command.

SECTION 3

InfiniteStorage 6120

Management

3.1 Managing the InfiniteStorage 6120

The InfiniteStorage 6120 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 only apply to firmware version 1.1.0.0. 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 InfiniteStorage 6120 can be accessed locally through a serial interface or remotely through SSH.

NOTE : An InfiniteStorage 6120 may only have one active login (serial or SSH) per controller at any given time.

Locally via Serial Interface

Any RS-232 terminal or terminal emulator (such as Hyperterminal) can be used to configure and monitor the InfiniteStorage 6120.

1. Connect the supplied mini-jack cable to the port labeled **CLI** on Controller 1 (Figure 47). Connect the other end of the cable to a serial port on a standard PC.

Figure 47. CLI Port on Controller

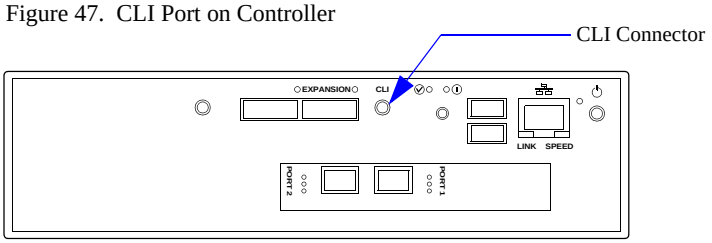

2. Load a serial console program (such as HyperTerminal, minicom, or screen) and use the following settings for the serial connection:
   - Bits per second: 115,200
   - Data bits: 8
   - Parity: None
   - Stop bits: 1
   - Flow control: None
3. Once connected, hit the <Enter> key to bring up the password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

Remotely via SSH

To configure and monitor the InfiniteStorage 6120 remotely, connect the controller to your Ethernet network. Please refer to Section 3.6, "Remote Management of InfiniteStorage 6120" 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, hit the <Enter> key to bring up the password prompt. Enter the password: **`user`**

Login via SSH

For a SSH session, enter the user name **`USER`** and password **`USER`** (Figure 48).

Figure 48. SSH Login Screen

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

Linux (none) 2.6.25-sgi-016620-3 #2 SMP Wed Jan 14 10:38:28 MST 2009 x
InfiniteStorage 6120
```

Logout

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

3.1.3 Available Commands

Entering the **`?`** character at the CLUI prompt will display the available commands within the InfiniteStorage 6120 CLUI (Figure 49). To get help information on a command, type the command followed by a question mark.

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

Figure 49. 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
   UPDATE           Update raid objects
```

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 50.

Simple, not full-command, line editing only is supported under the following conditions:

- CLUI prompts the user for more information.
- The 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 50.

Figure 50. 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~	Gets the oldest command in the history buffer.
End	Esc [K or Esc [4~	Gets 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	Deletes to beginning of line.
Ctrl-K	Ctrl-K	Deletes to end of line.
Ctrl-A	Ctrl-A	Moves cursor to beginning of line.
Ctrl-E	Ctrl-E	Moves cursor to end of line.

3.2 Configuration Management

The InfiniteStorage 6120 provides uniform configuration management across a heterogeneous SAN.

3.2.1 Monitor Status of Storage Assets

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

- Unassigned pool - By default, all the drives or replacement drives are initially allocated to this pool. When a drive 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 5 to 10 disk drives, 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.

3.2.1.1 Disk Information

`SHOW PHYSICAL_DISK *` displays a list of all the disks installed in the system (Figure 51).

`SHOW UNASSIGNED_POOL PHYSICAL_DISK` displays the list of unassigned disks.

`SHOW PHYSICAL_DISK=<disk-id> ALL` retrieves information about a specific disk (Figure 52).

Figure 51. Show Physical Disk List Example Screen

```
RAID[0]$ show physical_disk *

                                                                                        |Member|
Encl|Slot| Vendor |     Product ID     |Type|Cap GB |RPM|Revision|    Serial Number   |Pool|Status| Idx|Status|       WWN       |
----------------------------------------------------------------------------------------------------------------------------------
   1    1  HITACHI      HDS725050KLA360 SATA     465   7K K2AOAD1A       KRVN65ZAKAZLKH UNAS  GOOD    57        5000cca20def4516
   1    2  HITACHI      HDS725050KLA360 SATA     465   7K K2AOAD1A       KRVN65ZAKNL79H UNAS  GOOD    32  NORM  5000cca20df3a263
   1    3  HITACHI      HDS725050KLA360 SATA     465   7K K2AOAD1A       KRVP65ZAKD8Y5F UNAS  GOOD    42  NORM  5000cca20defdc87
   1    4  HITACHI      HDS725050KLA360 SATA     465   7K K2AOAD1A       KRVN65ZAK72BLH UNAS  GOOD     1  NORM  5000cca20ded7e05
   1    5  HITACHI      HDS725050KLA360 SATA     465   7K K2AOAD1A       KRVN65ZAKB269H UNAS  GOOD    38  NORM  5000cca20def4edd
   1    6  HITACHI      HDS725050KLA360 SATA     465   7K K2AOAD1A       KRVN65ZAKB1LHH UNAS  GOOD    23  NORM  5000cca20def4c96
   1    7  Hitachi      HUA721075KLA330 SATA     698   7K GK8OA70M       GTF200P8G6K2YF UNAS  GOOD    40  NORM  5000cca215c2fa66
   1    8  HITACHI      HDS725050KLA360 SATA     465   7K K2AOAD1A       KRVN65ZAKA6P9H UNAS  GOOD    41  NORM  5000cca20deeeb53
   1    9  HITACHI      HDS725050KLA360 SATA     465   7K K2AOAD1A       KRVN65ZAKAGVYF UNAS  GOOD    47        5000cca20def0a0a
   1   10  HITACHI      HDS725050KLA360 SATA     465   7K K2AOAD1A       KRVN65ZAKA5SDH    2  GOOD    45  NORM  5000cca20deee7d3
   1   11  HITACHI      HDS725050KLA360 SATA     465   7K K2AOAD1A       KRVN65ZAK8VH8F    2  GOOD    39  NORM  5000cca20dee4cee
   1   12  HITACHI      HDS725050KLA360 SATA     465   7K K2AOAD1A       KRVN65ZAKAZ75H    1  GOOD    26  NORM  5000cca20def43b5
   1   13  Hitachi      HUA721075KLA330 SATA     698   7K GK8OA70M       GTF202P8GTMNRF    4  GOOD    20  NORM  5000cca215cb32bd
   1   14  Hitachi      HUA721075KLA330 SATA     698   7K GK8OA70M       GTF202P8GTLDJF    4  GOOD    33  NORM  5000cca215cb2dfe
   1   15  Hitachi      HUA721075KLA330 SATA     698   7K GK8OA70M       GTF202P8GTPR7F    3  GOOD     3  NORM  5000cca215cb3a6e
   1   16  Hitachi      HUA721075KLA330 SATA     698   7K GK8OA70M       GTF202P8GTPW9F    5  GOOD    50  NORM  5000cca215cb3b0b
   1   17  Hitachi      HUA721050KLA330 SATA     465   7K GK6OAB0A       GTE402P6G94E8E    1  GOOD    27  NORM  5000cca214c427f6
   1   18  Hitachi      HUA721075KLA330 SATA     698   7K GK8OA70M       GTF202P8GTPNXF UNAS  GOOD    58        5000cca215cb3a45
   1   19  Hitachi      HUA721075KLA330 SATA     698   7K GK8OA70M       GTF202P8GTPTBF    4  GOOD    18  NORM  5000cca215cb3ab0
   1   20  HITACHI      HDS725050KLA360 SATA     465   7K K2AOAD1A       KRVN63ZAKSMX8D    0  GOOD    21  NORM  5000cca20df505cb
   1   21  HITACHI      HDS725050KLA360 SATA     465   7K K2AOAD1A       KRVP65ZAKERMGF    0  GOOD    17  NORM  5000cca20df08447
   1   22  Hitachi      HUA721050KLA330 SATA     465   7K GK6OAB0A       GTE402P6G9M0SE    1  GOOD    28  NORM  5000cca214c45ec1
   1   23  HITACHI      HDS725050KLA360 SATA     465   7K K2AOAD1A       KRVN63ZAKNTM6D    1  GOOD    35  NORM  5000cca20df3ba5a
   1   24  HITACHI      HDS725050KLA360 SATA     465   7K K2AOAD1A       KRVN65ZAKKVV3F    1  GOOD    30  NORM  5000cca20df26595
   1   25  HITACHI      HDS725050KLA360 SATA     465   7K K2AOAD1A       KRVN35ZAKE8HZH    0  GOOD     5  NORM  5000cca20df04f4d
   1   26  HITACHI      HDS725050KLA360 SATA     465   7K K2AOAD1A       KRVP65ZAKDW9PH    0  GOOD     6  NORM  5000cca20df0217e
   1   27  HITACHI      HDS725050KLA360 SATA     465   7K K2AOAD1A       KRVN63ZAKSNNDD    0  GOOD     8  NORM  5000cca20df50898
   1   28  HITACHI      HDS725050KLA360 SATA     465   7K K2AOAD1A       KRVN65ZAKAZ4LH    0  GOOD     9  NORM  5000cca20def4365
   1   29  HITACHI      HDS725050KLA360 SATA     465   7K K2AOAD1A       KRVN65ZAKB2GEH    2  GOOD    37  NORM  5000cca20def4fd9
   1   30  HITACHI      HDS725050KLA360 SATA     465   7K K2AOAD1A       KRVP65ZAKESH1F    1  GOOD    31  NORM  5000cca20df0877f
   1   31  HITACHI      HDS725050KLA360 SATA     465   7K K2AOAD1A       KRVN03ZAGMUMUD UNAS  GOOD    49        5000cca20dc9034c
   1   32  HITACHI      HDS725050KLA360 SATA     465   7K K2AOAD1A       KRVN35ZAK6VLBF    1  GOOD    25  NORM  5000cca20ded6490
   1   33  HITACHI      HDS725050KLA360 SATA     465   7K K2AOAD1A       KRVN65ZAKN0S7F    1  GOOD    36  NORM  5000cca20df360bf
   1   34  HITACHI      HDS725050KLA360 SATA     465   7K K2AOAD1A       KRVN65ZAK9VNRH    2  GOOD    44  NORM  5000cca20deec1f6
   1   35  HITACHI      HDS725050KLA360 SATA     465   7K K2AOAD1A       KRVN65ZAKL1MYH    0  GOOD    22  NORM  5000cca20df27b5c
   1   36  HITACHI      HDS725050KLA360 SATA     465   7K K2AOAD1A       KRVN03ZAGH9Y4D    2  GOOD    43  NORM  5000cca20dc6f6e1
   1   37  Hitachi      HUA721075KLA330 SATA     698   7K GK8OA70M       GTF202P8GTSLNF    3  GOOD     7  NORM  5000cca215cb4181
   1   38  Hitachi      HUA721075KLA330 SATA     698   7K GK8OA70M       GTF202P8GTRPVF    5  GOOD    54  NORM  5000cca215cb3e23
   1   39  Hitachi      HUA721075KLA330 SATA     698   7K GK8OA70M       GTF202P8GTSMGF    3  GOOD    14  NORM  5000cca215cb419a
   1   40  Hitachi      HUA721075KLA330 SATA     698   7K GK8OA70M       GTF202P8GTK4HF    5  GOOD    56  NORM  5000cca215cb2944
   1   41  Hitachi      HUA721075KLA330 SATA     698   7K GK8OA70M       GTF202P8GTRPDF    5  GOOD    52  NORM  5000cca215cb3e15
   1   42  Hitachi      HUA721075KLA330 SATA     698   7K GK8OA70M       GTF202P8GTRLVF    4  GOOD    24  NORM  5000cca215cb3dc6
   1   43  Hitachi      HUA721075KLA330 SATA     698   7K GK8OA70M       GTF202P8GTNPYF    3  GOOD    13  NORM  5000cca215cb36a4
   1   44  Hitachi      HUA721075KLA330 SATA     698   7K GK8OA70M       GTF202P8GTPWUF    5  GOOD    53  NORM  5000cca215cb3b1b
   1   45  Hitachi      HUA721075KLA330 SATA     698   7K GK8OA70M       GTF202P8GTPSHF    3  GOOD    12  NORM  5000cca215cb3a95
   1   46  Hitachi      HUA721075KLA330 SATA     698   7K GK8OA70M       GTF202P8GTPKNF    4  GOOD    29  NORM  5000cca215cb39e0
   1   47  Hitachi      HUA721050KLA330 SATA     465   7K GK6OAB0A       GTE402P6G94EYE    2  GOOD    46  NORM  5000cca214c4280b
   1   48  Hitachi      HUA721075KLA330 SATA     698   7K GK8OA70M       GTF202P8GTRP6F    5  GOOD    51  NORM  5000cca215cb3e0f
   1   49  Hitachi      HUA721075KLA330 SATA     698   7K GK8OA70M       GTF202P8GTV0KF    4  GOOD    15  NORM  5000cca215cb4a93
   1   50  Hitachi      HUA721075KLA330 SATA     698   7K GK8OA70M       GTF202P8GTUYLF    3  GOOD     2  NORM  5000cca215cb4a56
   1   51  Hitachi      HUA721075KLA330 SATA     698   7K GK8OA70M       GTF202P8GTSJZF UNAS  GOOD    59        5000cca215cb414d
   1   52  Hitachi      HUA721075KLA330 SATA     698   7K GK8OA70M       GTF202P8GTNV9F    5  GOOD    55  NORM  5000cca215cb372b
   1   53  Hitachi      HUA721075KLA330 SATA     698   7K GK8OA70M       GTF202P8GTTZ5F    4  GOOD    19  NORM  5000cca215cb46a7
   1   54  Hitachi      HUA721075KLA330 SATA     698   7K GK8OA70M       GTF202P8GTR3ZF    4  GOOD    34  NORM  5000cca215cb3bf9
   1   55  Hitachi      HUA721075KLA330 SATA     698   7K GK8OA70M       GTF202P8GTT3DF    3  GOOD     4  NORM  5000cca215cb4369
   1   56  Hitachi      HUA721075KLA330 SATA     698   7K GK8OA70M       GTF202P8GTPGWF    3  GOOD    10  NORM  5000cca215cb398a
   1   57  Hitachi      HUA721075KLA330 SATA     698   7K GK8OA70M       GTF202P8GTPH3F    3  GOOD    11  NORM  5000cca215cb3991
   1   58  Hitachi      HUA721075KLA330 SATA     698   7K GK8OA70M       GTF202P8GTSMLF    4  GOOD    16  NORM  5000cca215cb419e
   1   59  Hitachi      HUA721075KLA330 SATA     698   7K GK8OA70M       GTF202P8GTSK6F UNAS  GOOD    60        5000cca215cb4154
   1   60  Hitachi      HUA721075KLA330 SATA     698   7K GK8OA70M       GTF202P8GTPATF    5  GOOD    48  NORM  5000cca215cb390b

Total Physical Disks:      60
Total Assigned Disks:      54
Total Unassigned Disks:    6
  Total SATA Disks:        60
  Total Member State Normal:   54
```

Figure 52. Physical Disk Information Example Screen

```
RAID[0]$ show physical_disk 0x003f all_attributes

  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:     GOOD
  Rotation Speed:   7200 RPM
  Device Type:      SATA
  Member State:     UNASSIGNED
  Spare:            FALSE
  Failed:           FALSE
  UUID:             0x5000cca215c564560
```

If there is a failed drive, **SHOW UNASSIGNED_POOL FAILED ALL** will display the failed drive's information (Figure 53).

Figure 53. Failed Drive 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
  Spare:            FALSE
  Failed:           TRUE
  UUID:             0x5000cca215c564560
```

3.2.1.2 Storage Pool View

Storage pools are the basic building blocks of the InfiniteStorage 6120.

SHOW POOL * displays the list of configured storage pools in the system (Figure 54).

Figure 54. Show Configured 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   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          10  UNASSIGNED  AUTO
  4 pool-4          NORMAL       128   5             6192      0      0   W   R F I          10  UNASSIGNED  AUTO
  5 pool-5          NORMAL       128   5             6192      0      0   W   R F I          10  UNASSIGNED  AUTO

Total Storage Pools: 6
```

SHOW POOL <id> PHYSICAL_DISK displays the list of disks associated with the specified storage pool (Figure 55).

Figure 55. Show Storage Pool's Physical Disks Example Screen

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

                                                                                                       |Member|
Encl|Slot| Vendor |     Product ID     |Type|Cap GB |RPM|Revision|    Serial Number   |Pool|Status| Idx|Status|       WWN       |
----------------------------------------------------------------------------------------------------------------------------------
-
   1    4  HITACHI      HDS725050KLA360 SATA     465  7K K2AOAD1A      KRVN65ZAK72BLH    0  GOOD     1  NORM   5000cca20ded7e05
   1    6  HITACHI      HDS725050KLA360 SATA     465  7K K2A0AD1A      KRVN65ZAKB1LHH    0  GOOD    23  NORM   5000cca20def4c96
   1   20  HITACHI      HDS725050KLA360 SATA     465  7K K2AOAD1A      KRVN63ZAKSMX8D    0  GOOD    21  NORM   5000cca20df505cb
   1   21  HITACHI      HDS725050KLA360 SATA     465  7K K2AOAD1A      KRVP65ZAKERMGF    0  GOOD    17  NORM   5000cca20df08447
   1   25  HITACHI      HDS725050KLA360 SATA     465  7K K2AOAD1A      KRVN35ZAKE8HZH    0  GOOD     5  NORM   5000cca20df04f4d
   1   26  HITACHI      HDS725050KLA360 SATA     465  7K K2AOAD1A      KRVP65ZAKDW9PH    0  GOOD     6  NORM   5000cca20df0217e
   1   27  HITACHI      HDS725050KLA360 SATA     465  7K K2AOAD1A      KRVN63ZAKSNNDD    0  GOOD     8  NORM   5000cca20df50898
   1   28  HITACHI      HDS725050KLA360 SATA     465  7K K2AOAD1A      KRVN65ZAKAZ4LH    0  GOOD     9  NORM   5000cca20def4365
   1   35  HITACHI      HDS725050KLA360 SATA     465  7K K2AOAD1A      KRVN65ZAKL1MYH    0  GOOD    22  NORM   5000cca20df27b5c

Total Physical Disks:      9
Total Assigned Disks:      9
Total Unassigned Disks:    0
  Total SATA Disks:        9
  Total Member State Normal:   9
```

SHOW POOL <id> VIRTUAL_DISK displays the list of virtual disks associated with the specified storage pool (Figure 56).

Figure 56. Show storage pool's Virtual Disks Example Screen

```
RAID[0]$ show pool 0 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

Total Virtual Disks: 1
```

SHOW POOL <id> ALL displays detailed information of the storage pool (Figure 57).

Figure 57. Storage Pool Details Example Screen

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

  OID:                  0x19b40000
  Type:                 STORAGE
  Name:                 0x19b40000
  Chunk Size:           64KB (0x80 blocks)
  Block Size:           0x200
  RAID Type:            RAID5
  Free Raid5 Capacity:  2801664 MBs
  Max VD Size:          2801664 MBs
  Total Capacity:       3522560 MBs
  UUID:                 0x00
  Global Spare Pool:    UNASSIGNED
  DiskTimeout(FRT):     10 minutes
  Init Policy:          NO IO permitted
  Init Priority:        50%
  Full Rebuild Priority:        80%
  Fractional Rebuild Priority:  90%
  Sparing Policy:       AUTOMATIC
  Assign Policy:
    Device Type:        NA
    Rotation Speed:     NA
    Raw Capacity:       NA
  SATAssure:            FALSE
  IO ROUTING:           TRUE
  WBC:                  TRUE
  MWBC:                 FALSE
  Initializing:         TRUE
  Rebuilding:           FALSE
  Paused:               FALSE
  AutoWriteLock:        FALSE
  Critical:             FALSE
  Current Home:         0x0015b2a122b20000  0x00000000
  Future Home:          0xffffffffffffffff  0x00000000
  Preferred Home:       0x0015b2a122b20000  0x00000000
  BkgdJob OID:          0x28000003
  Total Phy Disks       5
  State:                NORMAL
  Member Size:          704512 MB
    pID     State     UUID
    0x003e  NORM      0x5000cca215c56e47
    0x003f  NORM      0x5000cca215c56456
    0x0040  NORM      0x5000cca215c54c71
    0x0041  NORM      0x5000cca215c5675c
    0x0042  NORM      0x5000cca215c56e55
```

3.2.1.3 Virtual Disk View

SHOW VIRTUAL_DISK * displays the list of configured virtual disks in the system (Figure 58).

Figure 58. List of Virtual Disks Example Screens

```
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           READY     3  5     5480   W   I           0(L) 0   1(R) 0  1(R) 0  INACTIVE
  4 vd-4_4           READY     4  5     5480   W   I           0(L) 0   0(L) 0   None   INACTIVE
  6 vd-6_5           READY     5  5     5480   W   I           0(L) 0   1(R) 0  1(R) 0  INACTIVE

Total Virtual Disks: 6
```

SHOW VIRTUAL_DISK <id> ALL_ATTRIBUTES will display detailed information of the virtual disk (Figure 59).

Figure 59. Virtual Disk Details Example Screen

```
RAID[0]$ show virtual_disk 0 all_attributes

  OID:              0x89ba0000
  Name:             89ba0000
  POOL OID:         0x19b40000
  Capacity:         16384 MBs
  Offset:           0x0
  State:            NOT RDY
  Raidlevel:        RAID5
  IO ROUTING:       TRUE
  WBC:              TRUE
  MWBC:             FALSE
  Initializing:     TRUE
  Paused:           FALSE
  AutoWriteLock:    FALSE
  Data Lost:        FALSE
  Present Home Only:FALSE
  Current Home:     0x0015b2a122b20000  0x00000000
  Future Home:      0xffffffffffffffff  0x00000000
  Preferred Home:   0x0015b2a122b20000  0x00000000
  BkgdJob OID:      0x28000000
  UUID:             60001ff0722b20000000000089ba0000
```

3.2.1.4 Spare Pool View

SHOW SPARE_POOL * will display the list of configured spare pool(s) (Figure 60).

Figure 60. List of Spare Pools Example Screens

```
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       0    0
  7 spare_pool-7            512    10       0    0

Total Spare Pools: 2
```

SHOW SPARE_POOL <id> ALL_ATTRIBUTES will display detailed information of the spare pool (Figure 61). **SHOW SPARE_POOL <id> PHYSICAL_DISKS** will display the list of disks in the pool.

Figure 61. Spare Pool Example Screen

```
RAID[0]$ show spare_pool 6 all_attributes
  OID:              0x19bc0006
  Type:             GLOBAL SPARE
  Name:             19bc0006
  Block Size:       0x200
  DiskTimeout(FRT): 10 minutes
  Total Capacity:   1409024 MBs
  UUID:             0x00
  Total Phy Disks   2

RAID$ show spare_pool 6 physical_disk
OID: 0x20b6004e Index: 0x004e
OID: 0x20b7004f Index: 0x004f
```

3.2.1.5 Visual Indication

Locate Physical Disk

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

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

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

Locate Storage Pool

LOCATE POOL <id> provides a visual indication of the specified storage pool. The status LED of the drives in the specified storage pool will blink.

Locate Enclosure

LOCATE ENCLOSURE <id> provides a visual indication of the specified enclosure. The Shelf Identify LED of the enclosure will blink.

3.2.2 Storage Pool and Virtual Disk Management

The InfiniteStorage 6120 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.

The storage array may consist of up to 24 storage pools. The storage pools can be used individually or split into multiple virtual disks. A virtual disk can be as small as part of a storage pool or uses all 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. Up to 64 virtual disks are supported in total.

NOTE : In this document, the terms "GB" and "GiB" are used interchangeably, and denote 1,073,741 bytes (2^{30} bytes).

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 62).

Figure 62. Storage Pool List Example Screen

```
RAID[0]$ show pool *

OID: 0x19b40000 Index: 0x0000 Name: 19b40000
OID: 0x19b60001 Index: 0x0001 Name: 19b60001
```

3.2.2.1 Changing the Disk's Status

SET PHYSICAL_DISK <disk-id> FAILED 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 failure 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 where the disk failed. This command will also initiate a rebuild if a spare has not already been assigned to the storage pool.

3.2.2.2 Creating a Storage Pool

To create a storage pool, enter command:
CREATE POOL RAID_LEVEL=<RAID level> CHUNK_SIZE=<chunk size> {ASSIGN_POLICY=[SAS|SATA] NUMBER=<number of disks> or PHYSICAL_DISK=<list of disks>}

where:

- RAID level = raid5 or raid6
- chunk size = 64kb, 128kb, or 256kb
- number of disks = 5 or 9 drives for RAID 5, 6 or 10 drives for RAID 6
- list of disks = index name of disks

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.

Examples:

- To use only the specified disks:
  CREATE POOL RAID_LEVEL=RAID5 CHUNK_SIZE=64KB PHYSICAL_DISK=0x003e 0x003f 0x0040 0x0041 0x0042
- 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. Once the storage pool has completed its initialization, its virtual disk(s) can be accessed immediately.

You may use the **SHOW POOL** command to check if the storage pool is being initialized (see Section 3.2.1.2, "Storage Pool View" 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 Naming a Storage Pool

SET POOL <pool-id> NAME="<name>" lets you specify a name to identify the storage pool (Figure 63).

Figure 63. Set Pool Name Example Screen

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

3.2.2.5 Deleting a Storage Pool

DELETE POOL=<id> deletes the specified storage pool from the system.

The DELETE POOL command erases all the data on the storage pool.

3.2.2.6 Creating a Virtual Disk

To create a virtual disk, enter 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 64).

A message is displayed to indicate whether the virtual disk creation was successful.

Figure 64. 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)
```

3.2.2.7 Naming a Virtual Disk

SET VIRTUAL_DISK <id> NAME="<name>" lets you specify a name to identify the virtual disk (Figure 65).

Figure 65. Set Virtual Disk Name Example Screen

```
RAID[0]$ set virtual_disk 0 name="LUN 0"
VIRTUAL_DISK 0 OID=0x89ba000 attributes set with STATUS='Success' (0x0)
```

3.2.2.8 Deleting a Virtual Disk

DELETE VIRTUAL_DISK=<id> deletes the specified virtual disk from the system.

When you delete a virtual disk, you lose access to all data that was stored on that virtual disk.

3.2.3 Spare Pool Management

The InfiniteStorage 6120 supports the concept of a 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 drives from the unassigned pool and assign them to the spare pool.

In the event of drive failure, a drive that is assigned to a spare pool is automatically swapped into the storage pool to replace the failed drive. A rebuild then occurs immediately, reducing the possibility of data loss due to additional drive failures.

NOTE : Each storage pool must 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.

3.2.3.1 Creating a Spare Pool

Use these commands to create a spare pool, add disks 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.

3.2.3.2 Naming a Spare Pool

`SET SPARE_POOL <pool-id> NAME="<name>"` lets you specify a name to identify the spare pool.

3.2.3.3 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>`** .

3.2.4 Drive Rebuild

The InfiniteStorage 6120's automatic drive failure recovery procedures ensure that absolute data integrity is maintained while operating in degraded mode.

3.2.4.1 Full and Partial Rebuilds

Both full and partial (fractional) rebuilds are supported in the InfiniteStorage 6120. 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 66).

Figure 66. Show Pool All Attributes

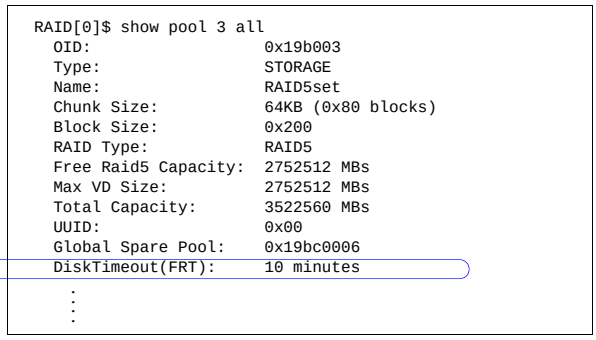

To change the disk timeout setting, use this command:
SET POOL <id> DISK_TIMEOUT=<timeout>
where <timeout> is in the range of <0..240> minutes. The default setting is 10.

3.2.4.2 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 drive failure, the InfiniteStorage 6120 will automatically initiate a drive 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 67).

Figure 67. Show Pool All Attributes

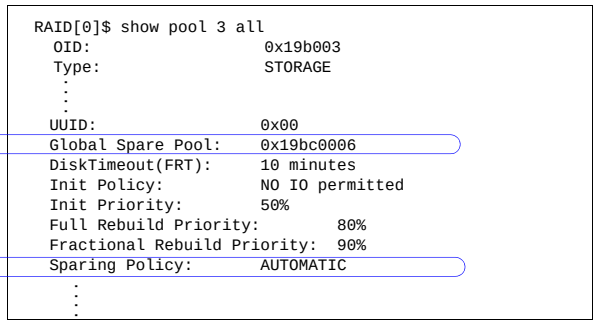

To change the sparing policy setting, use this command:
SET POOL <id> SPARING_POLICY=[AUTOMATIC|MANUAL]
The default setting is automatic which is recommended.

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 56 for information on how to adjust the rate of rebuild.

3.2.4.3 Manual Drive Replace/Rebuild

If no spare pool is assigned to the storage pool, you must manually add a drive to the pool to replace a failed drive.

To add a drive to the pool to replace a failed drive, use this command:
ASSIGN PD=<id of the disk to be added> STORAGE_POOL=<id of the storage pool> SET_SPARE

The disk will be put into the storage pool as the spare drive and it will be used in the storage pool to replace the failed drive.

3.2.5 SATAssure

If you have purchased the optional SATAssure license, you may enable it for the storage pool. To display the current setting, enterthe command
SHOW POOL <pool-id> ALL_ATTRIBUTES .

Figure 68. Show Storage Pool Details Example Screen

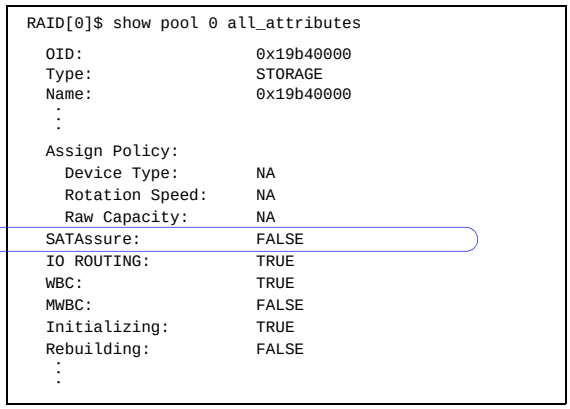

To enable SATAssure, enter the command **SET POOL <pool-id> SATASSURE=TRUE** .

To disable SATAssure, enter the command **SET POOL <pool-id> SATASSURE=FALSE** .

3.2.6 Cache Coherency

By default, the InfiniteStorage 6120 runs in couplet mode, where both controllers are running simultaneously, communicating through internal Ethernet connections. This means that the InfiniteStorage 6120 is always running with cache coherency enabled.

The write back cache settings are configurable for each storage pool. To display the current settings, enter command **SHOW POOL <id> ALL_ATTRIBUTES** (Figure 69).

Figure 69. Show Write Back Cache Setting

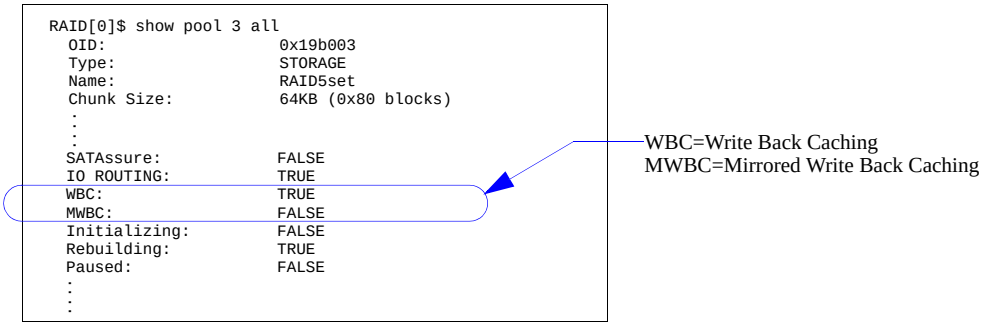

To change the settings, use these commands:
SET POOL <id> MIRRORED_WRITE_BACK_CACHE=[TRUE|FALSE]
SET POOL <id> WRITE_BACK_CACHING=[TRUE|FALSE]

3.2.7 Setting the System's Date and Time

SHOW SUBSYSTEM ALL_ATTRIBUTES displays the current system date and time (Figure 70).

Figure 70. Show System Time

```
RAID[0]$ show subsystem all_attributes
RP Subsystem Name:
UID:                    0x0000000000000000
Subsystem Time:         Fri Jan 30 16:48:55  2009
Time zone:              0 minutes
DST offset:             0 minutes
Locate Dwell Time:      120 seconds
Mirror Synch:           NA
```

The system time and date for the InfiniteStorage 6120 are factory-configured for the U.S. Pacific Standard Time (PST) zone.

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, 2009 2:15:32 pm, for example, type:
SET SUBSYSTEM DATE_AND_TIME=2009:3:1:14:15:32

3.2.8 Restarting the InfiniteStorage 6120

System Restart

SHUTDOWN CONTROLLER [LOCAL|REMOTE|0|1] RESTART performs a restart on the specified controller.

SHUTDOWN SUBSYSTEM RESTART 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

SHUTDOWN CONTROLLER [LOCAL|REMOTE|0|1] shuts down the specified controller.

SHUTDOWN SUBSYSTEM shuts down both controllers.

If you need to power down the InfiniteStorage 6120, use **SHUTDOWN** prior to shutting off the power. This will cause the InfiniteStorage 6120 to immediately flush its cache, abort all initialization and rebuild operations, and proceed with an orderly shutdown.

All hosts actively using the InfiniteStorage 6120 should be safely shutdown and all users logged out before using this command. The InfiniteStorage 6120 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 InfiniteStorage 6120 for maintenance. **SHUTDOWN** flushes any data left in the cache and prepares the InfiniteStorage 6120 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 71), enter this command:
SHOW JOB * ALL_ATTRIBUTES
The type of job and percentage of completion are displayed.

Figure 71. Show Job Example Screen

```
RAID[0]$ show job * all
  OID:              0x2b010003
  Target:           0x19b40003
  Type:             REBUILD
  Status:           RUNNING
  Priority:         80
  Fraction Complete:5%
  OID:              0x28000004
  Target:           0x19b60004
  Type:             INITIALIZE
  Status:           RUNNING
  Priority:         50
  Fraction Complete:10%
```

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 this 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 InfiniteStorage 6120 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 72).
To resume the job, enter command **RESUME JOB=<id>**

Figure 72. Pause/Resume a Job

```
RAID[0]$ pause job 0x28000003
JOB 3 OID=0x28000003 paused with STATUS='Janus Success' (0x0)

RAID[0]$ show job * all_attributes
  OID:              0x28000003
  Target:           0x19b40003
  Type:             INITIALIZE
  Status:           PAUSED
  Priority:         50
  Fraction Complete:11%

RAID[0]$ resume job 0x28000003
JOB 3 OID=0x28000003 resumed with STATUS='Janus Success' (0x0)

RAID[0]$ show job * all_attributes
  OID:              0x28000003
  Target:           0x19b40003
  Type:             INITIALIZE
  Status:           RUNNING
  Priority:         50
  Fraction Complete:11%
```

3.3.2 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 command **SHOW POOL <id> ALL_ATTRIBUTES** (Figure 73).

Figure 73. Show Pool Information

```
RAID[0]$ show pool 3 all_attributes
  OID:                  0x19b40003
  Type:                 STORAGE
  Name:                 raid 5 set
  Chunk Size:           64KB (0x80 blocks)
  Block Size:           0x200
  RAID Type:            RAID5
  Free Raid5 Capacity:  2752512 MBs
  Total Capacity:       3522560 MBs
  UUID:                 0x00
  Global Spare Pool:    0x1a0f000a
  DiskTimeout(FRT):     10 minutes
  Init Policy:          NO IO permitted
  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 74), enter this 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 74. 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 75), enter this 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 75. Set Partial Rebuild Priority

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

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

3.4 Firmware Update Management

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

3.4.1 Displaying Current Firmware Version

The **SHOW CONTROLLER <id> ALL_ATTRIBUTES** command displays version information of the InfiniteStorage 6120's hardware and firmware (Figure 76).

Figure 76. Controller Information Screen

```
RAID[0]$ show controller local all_attributes
OID:                    0x38000001
Firmware Version:
  Release:              1.0.3.3
  Source Version:       1877
  Fully Checked In:     Yes
  Private Build:        Yes
  Build Type:           Production
  Build Date and Time:  2009-03-02-09:49:EST
  Builder Username:     compiler
  Builder Hostname:     ereo-debian
  Build for CPU Type:   AMD-64-bit
Hardware Version:       0x0
State:                  RUNNING
Local AP OID:           0x00000000
Memory Size:            0x0
Max Q of S ID:          0x0
Up Time (seconds):      0x7fe3
Last Event Sequence #:  0x27b
Log Facility Mask:
  :
  :
  :
Crash Dump Enabled:     TRUE
Log Disk Enabled:       TRUE
RP Count:               0x1
Restart Pending:        FALSE
Name:                   B
Controller:             LOCAL        (SECONDARY)
Controller ID:          0x0015b2a1214a0000
MIR Reason:             None
```

3.4.2 Firmware Update Procedure

NOTE : The firmware update procedure described here only applies to upgrades from firmware version 1.0.4.1.xxxx or later to a new version. Upgrades from an older version of firmware must be done by a trained technician.

Before you begin upgrading the firmware, make sure all access to the InfiniteStorage 6120 is stopped, all volumes on the storage array are unmounted, and allow sufficient time for the InfiniteStorage 6120 to flush all cached data.

If you are upgrading from v1.0.3.3.xxx or later, there are two methods that can be utilized:

- Create an upgrade package on a USB flash drive.
- For Linux or Windows, copy the new firmware image over to the controller via the network instead of using a USB flash drive.

Warning

Firmware update should only be performed when both controllers are powered on and running.

NOTE : In the examples that follow, the filename of the firmware image is `flash-6620-1.1.0.0.2890-opt.sgi.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 .

3.4.2.1 Windows Environment (Update via USB Stick)

Contact SGI Technical Support to obtain a USB stick with the new firmware. Then perform the following steps to update the firmware:

Software Prerequisite: This firmware update method requires software that can extract (uncompress) a TGZ file. By TGZ file, we mean a TAR file that was compressed by GZIP. Microsoft Windows XP and Vista normally include zip extraction; just right-click on the file in Internet Explorer. If you do not have a Zip decompression program, download or purchase a software tool that will allow you to perform this operation.

1. Create a USB stick for the code update:
   - Insert your 4GB USB stick into your computer.
   - In Windows Explorer, right-click on the device with the Volume Name `JANUS_UP` and select **Format** from the pop-up menu (Figure 77). Then click the **Start** button in the Format screen.

Figure 77. Firmware Update Example Screen (1)

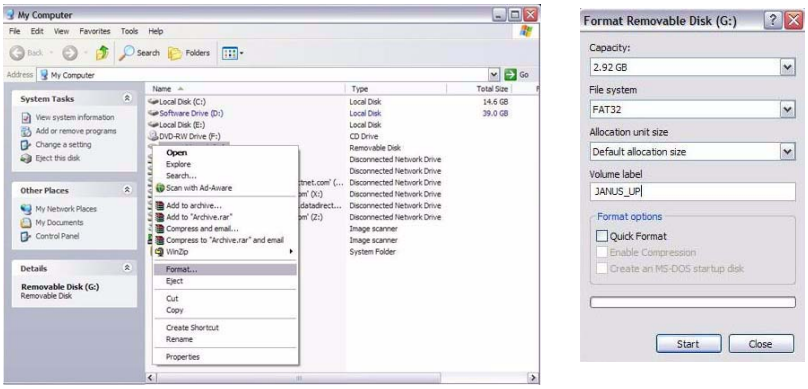

   - To copy and extract the code on the USB stick, right-click on the filename and select **Extract Files Here** (Figure 78). Extract the firmware file to the `JANUS_UP` USB volume on your USB stick.

Figure 78. Firmware Update Example Screen (2)

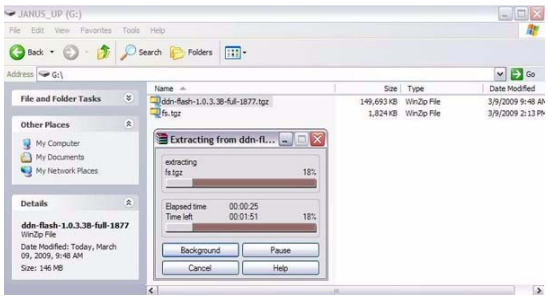

2. Shut down one of the controllers:
   At the CLUI prompt, enter command **SHUTDOWN REMOTE.**

3. Insert the USB stick into the back of **the controller that was not shut down in step 2**.

4. At the CLUI prompt, enter command **UPDATE CONTROLLER LOCAL** (Figure 79)

   Figure 79. Update Firmware Screen

```
RAID[0]$ update controller local
Starting upgrade...
Current Janus Linux, skipping Linux install
Installing JANUS code into sgi-v2
Upgrading Janus sgi-v2 opt...
Changing GRUB boot from sgi-v1 to sgi-v2
Done upgrading Janus
unmount: /tmp/usb: device is busy
Broadcast message from root@(none) (Sat Mar 14 13:25:52 2009):
The system is going down for reboot NOW!
```

5. Upon completion of reboot, login and enter command **SHOW CONTROLLER LOCAL ALL**. to verify that the firmware version is correct.

NOTE : Any virtual disks that exist will be in a CRITICAL state. This is expected behavior and has no ill effect on the system.

---

6. Shut down the upgraded controller:
   At the CLUI prompt, enter command **SHUTDOWN LOCAL.**

7. Upgrade the other controller:

   a)Restart the controller.

   b)Repeat the steps 3-5 on the other controller. You do not need to back out either controller during this procedure.

   c)Once the controller has been upgraded, restart the first controller by pressing the power switch on the back of the controller.

8. Once both controllers are running, all virtual disks in the system will be reported as **CRITICAL**. This state is due to the new Cache Safe feature and is expected. Clear this critical state:
   At the CLUI prompt, enter command **CLEAR POOL * CRITICAL.**

3.4.2.2 Linux Environment (Update from Network)

Follow these steps to update the firmware:

1. Copy the new firmware to the controller using an scp or sftp client. The username is `firmware` and the password is `Firmware` .(Note that entries are case-sensitive.)
   For example:
   `scp flash-6120-1.1.0.0.2890-opt.sgi.tgz firmware@10.32.31.240:`

2. Shut down one of the controllers:
   At the CLUI prompt, enter command **`SHUTDOWN REMOTE.`**

3. At the CLUI prompt, enter this command :
   `UPDATE_FIRMWARE CONTROLLER=<controller-id> FILE="<file-specification>"`

   For example:
   `UPDATE_FIRMWARE CONTROLLER=LOCAL`
   `FILE="flash-6120-1.1.0.0.2890-opt.sgi.tgz"`
   (Note that the filename must be enclosed within double quotation marks.)

4. Upon completion of reboot, log in and enter command **`SHOW CONTROLLER LOCAL ALL`**.
   Verify that the firmware version is correct.

NOTE : Any virtual disks that exist will be in a CRITICAL state. This is expected behavior and has no ill effect on the system.

---

5. Shutdown the upgraded controller:
   At the CLUI prompt, enter command **`SHUTDOWN LOCAL.`**

6. Upgrade the other controller:

   a)Restart the controller.

   b)Repeat the steps 3-5 on the other controller. You do not need to back out either controller during this procedure.

   c)Once the controller has been upgraded, restart the first controller by pressing the power switch on the back of the controller.

7. Once both controllers are running, all virtual disks in the system will be reported as CRITICAL. This state is due to the new Cache Safe feature and is expected. Clear this critical state:
   At the CLUI prompt, enter command **`CLEAR POOL * CRITICAL.`**

3.4.2.3 Windows Environment (Update from the Network)

This method requires a SFTP or SCP client, and involves two steps – copying the firmware image to the controller and launching the update. This is an offline upgrade so all I/O load from hosts must be stopped prior to 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** (Figure 80).

At the login screen, use the user name `admin` with the password `password` to log into the system. User name and passwords are case-sensitive..

Figure 80. Login Screen

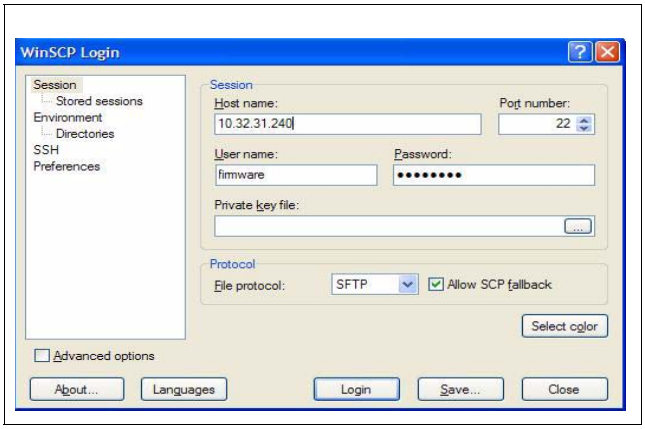

4. Find the image file, drag it to the destination window, and select **Copy**.
   The copy operation commences.

Launch the Update

1. Log in to the controller via an SSH session with the username of `user` and password of `user`.

2. Shut down one of the controllers:
   At the CLUI prompt, enter command **`SHUTDOWN REMOTE.`**

3. At the CLI prompt, enter:
   RAID[0]$ update_firmware controller=local
   file="`flash-6120-1.1.0.0.2890-opt.sgi.tgz`"

NOTE : The quotes around the filename of the firmware image are required. The firmware update utility will be launched and the controller will be automatically rebooted to use the new firmware image.

---

4. Upon completion of reboot, log in and enter command **`SHOW CONTROLLER LOCAL ALL`** to verify that the firmware version is correct.

NOTE : Any virtual disks that exist will be in a CRITICAL state. This is expected behavior and has no ill effect on the system.

---

5. Shut down the upgraded controller:
At the CLUI prompt, enter command **`SHUTDOWN LOCAL.`**

6. Upgrade the other controller:

a)Restart the controller.

b)Repeat the steps 3-5 on the other controller. You do not need to back out either controller during this procedure.

c)Once the controller has been upgraded, restart the first controller by pressing the power switch on the back of the controller.

7. Once both controllers are running, all virtual disks in the system will be reported as CRITICAL. This state is due to the new Cache Safe feature and is expected. Clear this critical state:
At the CLUI prompt, enter command **`CLEAR POOL * CRITICAL.`**

3.5 The System Logs

To display the event log starting at the start sequence number in ascending sequence number order (Figure 81), enter this command:

SHOW CONTROLLER <id> LOG ASCEND

Figure 81. Controller Log Example Screen (1)

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

000001 2009-02-05 12:13:31:3387789 G=0 S=0 T=1 RP=0 VP=63
LOG_JOI_BUILD_INFO1 JOI FW VERSION ON PROCESSOR 0X40 = (1.0.3.1-0).

000002 2009-02-05 12:13:31:3387798 G=0 S=0 T=1 RP=0 VP=63
LOG_JOI_BUILD_INFO2 JOI FW WAS BUILT ON JJDEBIAN AT 2009-02-05-14:17:MST (PRODUCT).

000003 2009-02-05 12:13:52:8320854 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 82), enter this command:

SHOW CONTROLLER <id> LOG DESCEND

Figure 82. Controller Log Example Screen (2)

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

000024 2009-02-11 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 2009-2-11 5:8:48; NEARBY LOG ENTRIES MAY APPEAR OUT
OF TIME ORDER. OFFSET= 0X1C987C765CD3B2B.

000023 2009-02-06 04:49:20:5069068 G=3 S=1 T=1 RP=0 VP=1 LOG_LOGDISK_ENABLE_RECEIVED
_FROM_STATE LOGRECEIVED FROM STATE

000022 2009-02-06 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 2009-2-6 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 83), enter this command:

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

Figure 83. Controller Log Example Screen (3)

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

000021 2009-02-06 04:49:20:4731906 G=4 S=2 T=1 RP=0 VP=1 LOG_ST_MIR_STATE STATE
 MIR STATE STATE:000A

000022 2009-02-06 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 2009-2-6 4:49:20; NEARBY LOG ENTRIES MAY APPEAR
OUT OF TIME ORDER.OFFSET = 0X1C987C7A302483D.

000023 2009-02-06 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 2009-02-11 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 2009-2-11 5:8:48; NEARBY LOG ENTRIES MAY APPEAR
OUT OF TIME ORDER. OFFSET  = 0X1C987C765CD3B2B.
```

3.6 Remote Management of InfiniteStorage 6120

The InfiniteStorage 6120 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).

3.6.1 Network Connection

Connect the Telnet port on the back of the controllers to your Ethernet network (Figure 84). Then configure the network interface as described below.

Figure 84. Connectors on Controller

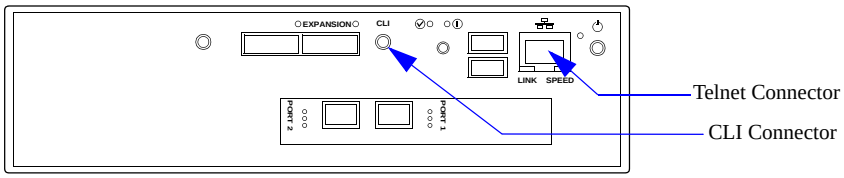

NOTE : Currently, the InfiniteStorage 6120 does not support network configuration protocols such as DHCP or BOOTP.

3.6.2 Network Interface Setup

UI SHOW NETWORK_INTERFACE [LOCAL|REMOTE] * displays the list of the network interfaces with their associated controller IDs and devices (Figure 85). **LOCAL** gives you information on the controller that you are currently logged into. **REMOTE** gives you information from the other controller in the pair.

Figure 85. Current Network Interface Settings Screen

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

To change the network interface settings on the controller to which you are connected, enter this 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 86).

Figure 86. 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=10.32.1(
NETWORK_INTERFACE 0 set with STATUS='Success' (0x0)
```

Enter command **SHUTDOWN SUBSYSTEM RESTART** to restart the controllers and make the changes take effect.

NOTE : This network interface setup must be performed separately on each controller. Since you cannot use a remote command currently, you must plug the serial port into each controller to issue the commands to each controller.

3.6.3 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.

3.7 Presentations

A *presentation* on an InfiniteStorage 6120 has the following components:

- Discovered Initiators
  A *discovered initiator* is defined as a Fibre Channel initiator that has logged into the InfiniteStorage 6120. This is normally a single port on a Fibre Channel 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 initiators. The host attribute OS (operating system) allows you to 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 ports on the controller. The channel attribute MODE allows you to 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 InfiniteStorage 6120, 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. The following are attributes of a presentation:
  - PORT – 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 – LUN number (user-specified) that the virtual disk will show to the host.

3.7.1 Application Discovered Initiator Commands

To import a discovered initiator object that is associated with the specified host, use the **APPLICATION IMPORT DISCOVERED_INITIATOR=<object-id> HOST=<object-id>** command.

To display the attributes of a specified initiator, use the **APPLICATION SHOW DISCOVERED_INITIATOR=<object-id> [ALL_ATTRIBUTES]** command.

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

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

3.7.2 Application Host Commands

To create a host stack, use the **APPLICATION CREATE HOST [OSTYPE=GENERIC|LINUX|MAC_OS|WINDOWS]** command. In the case of a FCP target

application stack, the ID string is not used. In the case of the Lustre OSS, the ID string is the name of a block device. The default operating system type is GENERIC.

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

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

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

3.7.3 Application Presentation Commands

To create a presentation to a virtual disk for the specified host, use the **APPLICATION CREATE PRESENTATION VIRTUAL_DISK=<object-id> HOST=<object-id>** command.

To delete a presentation to a virtual disk for the specified host, use the **APPLICATION DELETE PRESENTATION VIRTUAL_DISK=<object-id> HOST=<object-id> [FORCE]** command. Note the optional parameter of FORCE deletes without asking questions.

To delete the specified presentation, use the **APPLICATION DELETE PRESENTATION=<object-id> [FORCE]** command. Note the optional parameter of FORCE deletes without asking questions.

The wildcard * deletes all presentations.

To set the specified attribute to the specified value, use the **APPLICATION SET PRESENTATION=<object-id> <attribute-name>=<value> [<attribute-name>=<value>…]** command. 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 read_only.

SECTION 4

Supporting the

InfiniteStorage 6120

The InfiniteStorage 6120 system includes an enclosure services processor and associated monitoring and control logic to enable it to diagnose problems within the enclosure's power, cooling, and drive systems.

The sensors for power and cooling conditions are housed within the power supply/cooling modules. Each module is independently monitored.

4.1 Status LED Indicators

Green LEDs are always used for good or positive indication. Amber LEDs indicate there is a critical fault present within the module.

4.1.1 Front Panel Indicators

The front panel indicators show the aggregated status of all the modules (Figure 87). The LEDs are defined in Figure 88.

Figure 87. Front Panel LED Indicators

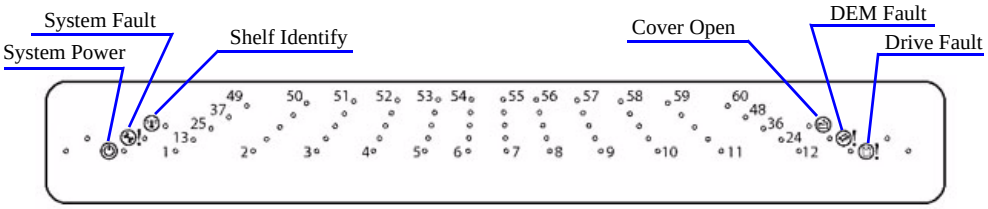

Figure 88. Front Panel LED Description

LED	Definition	Color	Status
	Shelf Identify	Blue	ON - SES is sending an identify command . OFF - SES is NOT sending an identify command .
	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
	System Power	Green	ON - DC power is present. OFF - DC power is not present. LED does NOT flash under normal operating conditions.
	Cover Open	Amber	OFF - both cover pieces is securely closed and latched in place . ON - Either of the cover pieces is NOT securely closed and latched in place.

Figure 88. Front Panel LED Description

LED	Definition	Color	Status
	DEM Fault	Amber	OFF - all DEMs operating correctly. ON - at least one DEM has failed; service action required.
	Drive Fault	Amber	ON - One or more HDDs are failed; SES must determine exact HDD. OFF - No detectable drive faults.
Individually numbered	HDDs	Green	ON - SAS HDD is present. Blinks - Indicates HDD activity. OFF - No HDD activity.

4.1.2 Power Cooling Module (PCM)

Figure 89 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 89. Power Cooling Module LEDs

Normal Operation LEDs

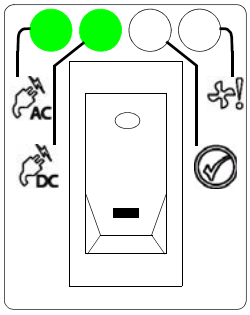

Faults LEDs

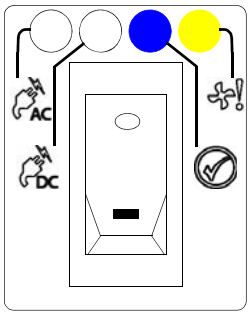

LED	Description	Color	Status
Enclosure front	Enclosure fault	Amber	ON - AC input, DC output, fan or other PCM fault detected
	PCM AC	Green	ON - AC input to PCM within tolerances OFF - PCM failed
	PCM DC	Green	ON - DC output of PCM within tolerances OFF - PCM failed
	PCM Fault	Amber	ON - AC input, DC output, fan or other PCM fault detected OFF - no fault detected
	PCM ID	Blue	ON - receiving SES identify command OFF - NOT receiving SES identify command

4.1.3 Controller

Figure 91 describes the controller LED (Figure 90) 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 90. LEDs on Controller

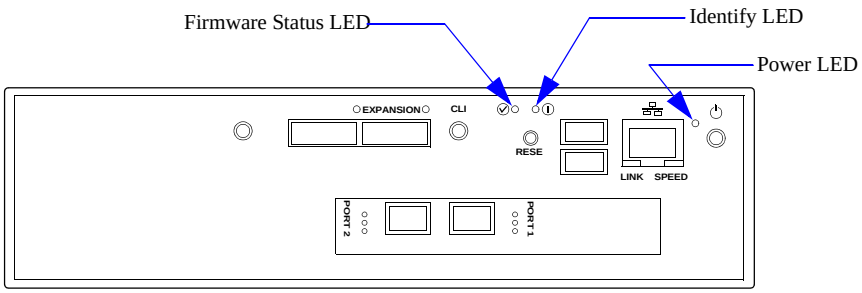

Figure 91. Controller LED Indicators

Identifier	Location	Color	Status
Enclosure fault	Front enclosure	Amber	ON - IO fault detected.
Firmware status	Rear enclosure	Green	ON - InfiniteStorage 6120 is ready to be accessed. Blinks - Firmware is loading OFF - InfiniteStorage 6120 is NOT ready to be accessed.
Power	Rear enclosure	Green	ON - Power is applied to controller. OFF - Power is NOT applied to controller.
Identify	Rear enclosure	Blue	ON - SES identify command is received. OFF - SES identify command is NOT received.

4.1.4 HDD (Hard Disk Drive)

Figure 92 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 92. HDD LED Indicators

Identifier	Location	Color	Status
Enclosure fault	Enclosure front	Amber	ON - drive fault
Drive fault	Enclosure front	Amber	ON - drive fault

4.1.5 DEM (Drive Expander Module)

Figure 93 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 93. 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	ON- This DEM is being sent an identify command, the SEP.

4.2 Troubleshooting

This section describes common problems, with possible solutions, which can occur with the InfiniteStorage 6120 system.

4.2.1 Initial Start-up Problems

Faulty Cords

First, check that you have wired up the enclosure correctly. Then if any of the following is true, call SGI Technical Support:

- Cords are missing or damaged .
- Plugs are incorrect.
- Cords are too short.

Call Customer Service for a replacement.

Computer Does Not Recognize the InfiniteStorage 6120

1. Ensure that the host interface cables from the InfiniteStorage 6120 to the host computer are connected correctly.
2. Ensure that all drive modules have been correctly installed.
3. Ensure that the controller is set up as follows:
   - The controller has been correctly installed and all external links and cables are securely fitted.
   - The maximum cable length has not been exceeded.

4.2.2 `RAID[0]$` Prompt Not Shown Correctly

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

- The firmware on this controller is in a *MIR* state and the MIR state needs to be resolved. Refer to Section 4.5, "Manual Intervention Required (MIR) Conditions" on page 82 for troubleshooting information.
- 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.

4.2.3 System Faults

Symptom	Cause	Action
1. The System Fault LED turns amber on front panel.	The ESI processor has detected an internal fault (for example, 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 drive modules. If none are evident, then there may either be an ESI processor problem or a backplane problem.

4.2.4 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.** Ensure that the power On/Off switch on rear of power cooling module is turned ON. **2.** Ensure that the AC main connection to the power cooling module is live. **3**. Disconnect the power cooling module from the main power, remove the module from the system, and re-install the module. If the problem persists, replace the power cooling module. **4**. Reduce the ambient temperature.

4.2.5 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 the PSU is above 55°C.	If the air temperature through the enclosure exceeds a pre-set threshold, a thermal alarm will sound.	**1**. Ensure that the 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 fron. Use in a fully enclosed rack installation is not recommended. **5**. If possible, shut down 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 disk drives.

4.3 Component Failure Recovery

The InfiniteStorage 6120 contains redundant and hot-swappable power cooling modules. 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 controller is running. The replaced component will automatically be returned to service once the component has been installed and booted up.

If the InfiniteStorage 6120 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.

Observe all conventional ESD precautions when handling InfiniteStorage 6120 modules and components. Avoid contact with backplane components and module connectors.

4.3.1 Replacing a Power/Cooling Module

If a power supply or fan fail, the PCM Fault LED on the power cooling 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 five 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.

As there should always be two (2) power cooling modules installed, the system can continue running while you replace the faulty module.

1. Identify the failed PCM with an amber Fault LED.
2. Turn off the module's power switch and disconnect the power cord.
3. Remove the module's thumbscrew, then slide the module out of the bay.
4. Ensure that the power switch on the new (replacement) module is OFF.
5. Slide the module into the bay making sure it is fully inserted. Install the thumbscrew to secure it.
6. Connect the power cord. Turn on the power switch. Verify that the AC and DC LEDs are **green**, indicating that the module is operating normally.

4.3.2 Replacing a Drive Module

NOTE : When a drive is failed by the InfiniteStorage 6120, replace it immediately so that the operation and performance of the system are not affected.

Figure 94. Drive Module Removal

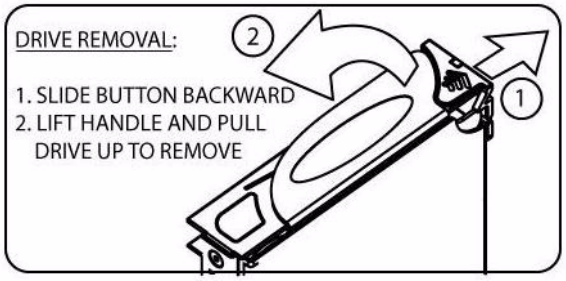

CAUTION : Drive spin down—Damage can occur to a drive if it is removed while still spinning. If possible, use the operating system to spin down the drives prior to removal. If this is not possible, we recommend that you perform ALL steps of the following procedure to ensure that the drive has stopped prior to removal.

To remove a module:

1. Slide the latch backward to release handle (Figure 94).
2. Lift handle and pull module up, just enough to disconnect the module from the backplane.
3. Wait for 30 seconds for the drive to completely spin down.
4. Pull the module gently out of the drive bay.

To insert a module:

1. Slide the latch backward to release handle.
2. Insert module into drive bay. Note orientation of module.
3. Cam the drive 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.

4.4 Recovering from a Drive Failure

When a drive failure occurs on the InfiniteStorage 6120, the storage pool containing that drive 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 drive failures on the same storage pool. If another disk drive fails on the same storage pool before the data on the first drive is rebuilt (to a replacement disk or hot spare), the storage pool will go offline.

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

4.4.1 Automatic Rebuild

A single drive failure in any storage pool does not result in the data loss. The virtual disk(s) on that storage pool will continue to operate in degraded mode. If a spare drive is available and automatic rebuild is enabled, the InfiniteStorage 6120 will automatically rebuild the data on the spare drive.

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

When a drive failure occurs, the error is written to the event log. The controller will rebuild the drive automatically once it finds a suitable spare drive. You may monitor the rebuild progress or adjust the rebuild rate to match the user load.

To display the event log (Figure 95), use the **SHOW CONTROLLER LOCAL LOG** command.

Figure 95. Display Event Log

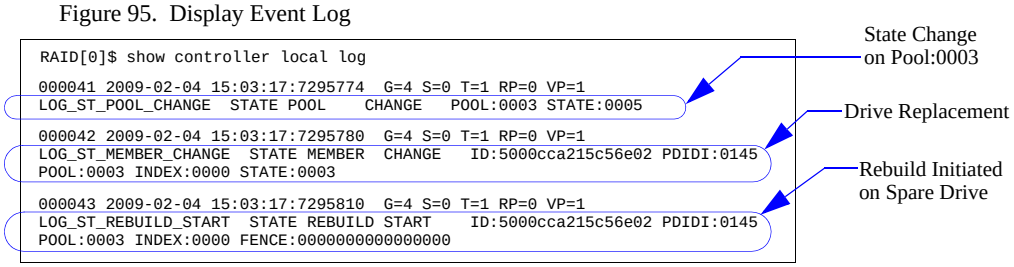

To look at the failed drive (Figure 96), enter command **SHOW UNASSIGNED_POOL FAILED ALL**

Figure 96. Display Failed Drive

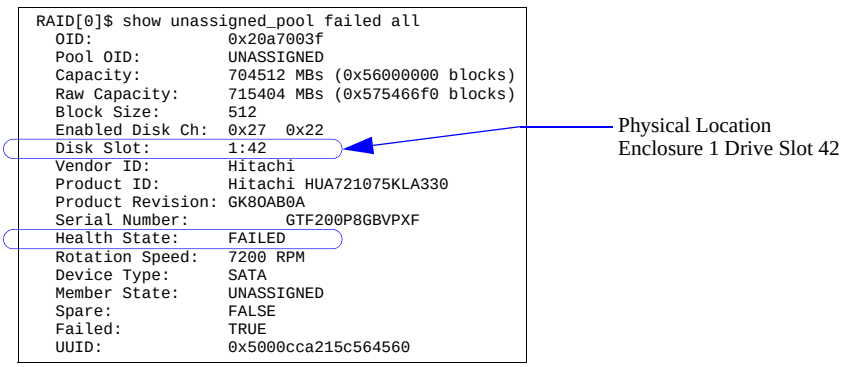

To monitor the rebuild progress (Figure 97), enter command **SHOW JOB * ALL_ATTRIBUTES**

Figure 97. 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 98), enter command **SHOW POOL <id> ALL_ATTRIBUTES**. Once the rebuild is complete, the status of storage pool will return to "NORMAL".

Figure 98. Display Storage Pool Information

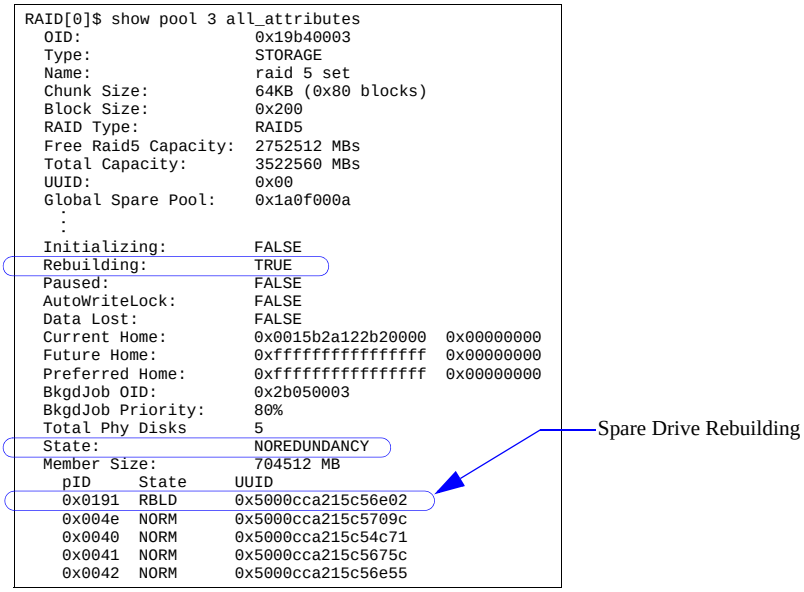

4.4.2 *Drive Replacement*

Upon completion of a rebuild, the spare drive becomes a member of the storage pool, replacing the failed drive. After you have replaced the failed drive with a new drive, the new drive is added to the unassigned pool. It is recommended that you assign this new drive to the spare pool to "replace" the spare drive that has been used.

4.4.3 *Spare Not Available*

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

1. Enter command **SHOW UNASSIGNED_POOL * ALL** to identify the new drive'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 drive and <pool-id> is the OID of the storage pool that had the failed drive.

4.5 Manual Intervention Required (MIR) Conditions

Manual Intervention Required (MIR) represents a condition with the controller that requires you 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. 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 you to choose whether or not to proceed.

SHOW CONTROLLER LOCAL ALL_ATTRIBUTES will display the MIR condition of the controller where you are logged in (Figure 99).

Figure 99. Controller Information Screen

```
RAID[0]$ show controller local all_attributes
OID:                    0x38000001
Firmware Version:
  Release:              1.0.3.3
  Source Version:       1877
  Fully Checked In:     Yes
  Private Build:        Yes
  Build Type:           Production
  Build Date and Time:  2009-03-02-09:49:EST
  Builder Username:     compiler
  Builder Hostname:     ereo-debian
  Build for CPU Type:   AMD-64-bit
Hardware Version:       0x0
State:                  RUNNING
Local AP OID:           0x00000000
Memory Size:            0x0
Max Q of S ID:          0x0
Up Time (seconds):      0x7fe3
Last Event Sequence #:  0x27b
Log Facility Mask:
  .
  .
  .
Crash Dump Enabled:     TRUE
Log Disk Enabled:       TRUE
RP Count:               0x1
Restart Pending:        FALSE
Name:                   B
Controller:             LOCAL        (SECONDARY)
Controller ID:          0x0015b2a1214a0000
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 the 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 drives on the backend. Install disk drives on the backend or fix the condition that prevents this controller from finding backend drives. Please refer to the installation guidelines for proper setup.

MIR_NO_CONFIG

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

MIR_NO_QUORUM

No quorum drives could be found within the disk drives on the backend. Fix the condition that prevents the quorum drives from being seen. This may be caused by the drives 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 one 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 drives from one system were installed in another system when the systems were powered down and then rebooted. If the drives 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 drives and the other controller saw a different subset of drives. 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 one existed.

MIR_NO_LOAD_CONFIG

A configuration could not be loaded from the backend drives. Fix the condition that prevents this controller from loading a configuration from the backend drives 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 one existed.

SECTION 5

GUI Management

5.1 The GUI Management Agent

This section provides information on configuring your InfiniteStorage 6120 using the 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 InfiniteStorage 6120.

5.2 Starting the GUI Management Agent

5.2.1 Login

Using a web browser, open a link to the IP address of the InfiniteStorage 6120. 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, use the user name `admin` with the password `password` to log into the system. The user name and password are case-sensitive.

Figure 100. Login Screen

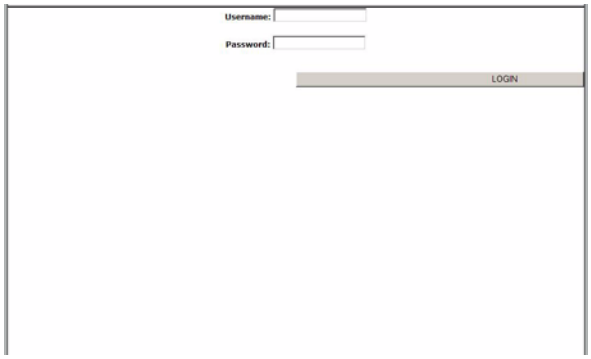

5.2.2 Initial Home Screen and Health Indicators

The home screen shows the Health Indicator. Since this is a new unconfigured system, at this point, the Health Indicator should be orange. You should see an alert message below the Health Indicator which show that manual intervention is required for both controllers.

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 Show Controller Screen. The home screen shows the Health Indicator.

2. Select the underlined **Manual Intervention Required**.
   Details of the MIR state display show `STATUS='MIR: no quorum disks found'`.

3. Press the CLEAR MIR button to clear the MIR state.
   Your subsystem is now ready to be configured.

Figure 101. Controllers-Show All Screen

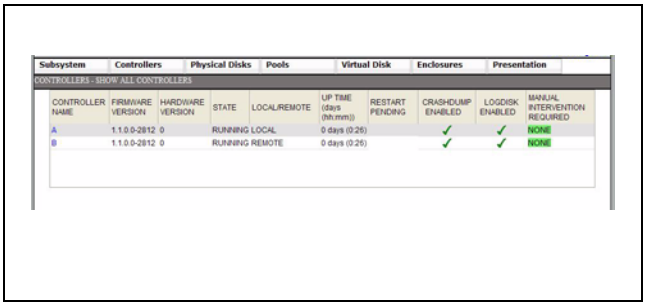

5.3 Physical Disks

5.3.1 Physical Disks Commands

Physical disk commands are located under the **Physical Disk** tab drop down menu:

- **Show Physical Disks**
- **Show Failed Disks**
- **Locate Disks**
- **Set Failed**
- **Assign to Pool**
- **Clear Failed**

5.3.2 Checking Status of Physical Disks

Before creating any storage pools, check the status of all the drives.

Under the **Physical Disks** tab, select **`SHOW ALL PHYSICAL DISKS`** to verify that all the drives are present and healthy.

Figure 102. Physical Disk-Show All Screen

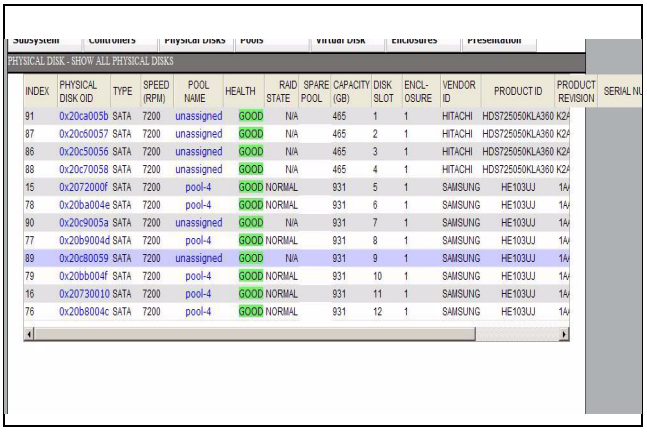

NOTE : If the drives are not visible, select **`SUBSYSTEM RESTART`** under the **Subsystem** tab to reboot the system.

5.4 Storage Pools

When creating a storage pool on an InfiniteStorage 6120, the following selectable attributes are available:

- RAID Level
  Storage pools can be configured to use either a RAID5 or RAID6 parity scheme. In RAID5, the capacity of one drive is reserved for parity, allowing data recovery in the event of a single drive loss in the storage pool. In RAID6, the capacity of two drives is reserved for parity, allowing data recovery if either one or two drives are down in a storage pool. ***For maximum data protection, use RAID6***.
- Chunk Size
  The chunk size defines the amount of data written to a single disk before proceeding to the next disk in the storage pool. For optimal performance, the chunk size should be proportional to the expected host request size. For example, a large chunk size should be used when transferring large blocks of data while a small chunk size should be used when transferring multiple small blocks of data.
- Drive Count
  A RAID5 storage pool may consist of 5 or 9 drives. A RAID6 storage pool may consist of 6 or 10 drives. For maximum performance, storage pools should be created using drives 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 drive.
- 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 drives and make sure that data integrity is always maintained for all IO operations. A license is required to run this feature.

  Note : SGI does not support the mixing of SAS and SATA drives in the same storage pool.

5.4.1 Storage Pool Commands

Storage pool commands are located under the **Pool** tab drop-down menu:

- **Show Storage Pools**
- **Create Storage Pool**
- **Set Storage Pools Attributes**
- **Delete Storage Pool**
- **Locate Storage Pool**

To create a storage pool :

1. Under **POOLS** tab, select **`CREATE STORAGE POOL.`**
   The Storage Pools-Create Storage Pool screen displays (Figure 103).

Figure 103. Create Storage Pool Screen 1

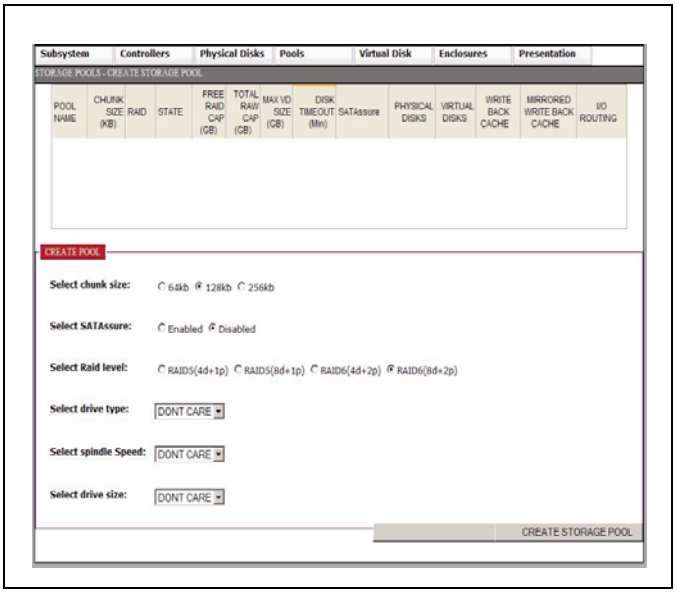

2. Select **`CHUNK_SIZE: 64kb, 128kb, 256kb.`**
3. Enable or disable **`SATAssure.`**
4. Select **RAID Level: `RAID5(4d+1p), RAID5(8d+1p), RAID6(4d+2p), RAID6(8d+2p)`**
5. Select **Drive Type** from scroll list.
6. Select **Spindle Speed** from scroll list.
7. Select **Drive Size** from scroll list.
8. Select **CREATE STORAGE POOL.**
   The system displays a success or error message and a list of storage pools created (Figure 104).

Figure 104. Storage Pool List

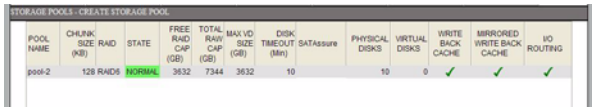

9. Repeat step 1 to add more storage pools as needed.

To show storage pools:

Under **POOLS** tab, select **`SHOW STORAGE POOLS`**.

To delete a storage pool:

1. Under **POOLS** tab, select **DELETE STORAGE POOL**.
   The Storage Pools-Delete Storage Pool screen displays.
2. Select the storage pool you want to delete.
3. Press the **DELETE** button or press **CANCEL** to cancel operation.
   A success or failure message displays.

To set storage pool attributes:

1. Under **POOLS** tab, select **SET STORAGE POOL ATTRIBUTES**.
   The Storage Pools-Set Attributes screen displays.
2. Select the storage pool whose attributes you want to change.
   The Set Attributes screen appears (Figure 105).

Figure 105. Storage Pool Set Attributes

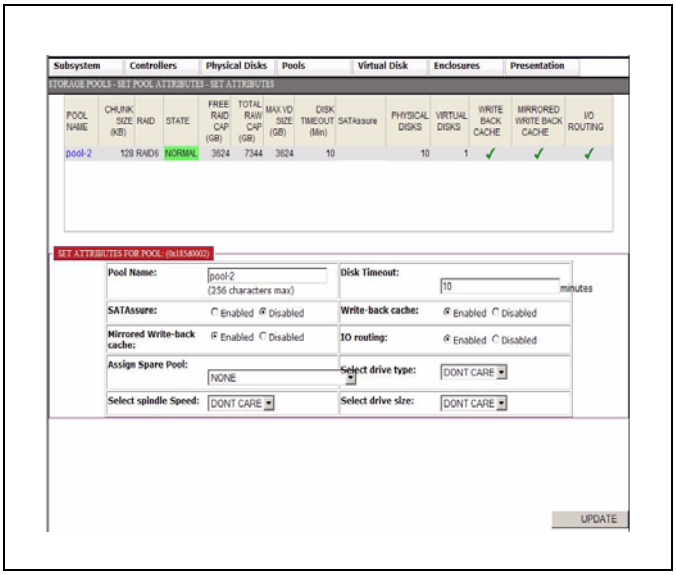

3. Make the changes desired.
4. Press the **UPDATE** button.
   A success or failure message displays.

To locate storage pools:

1. Under **POOLS** tab, select **LOCATE STORAGE POOL**.
   The Storage Pools-Locate screen displays (Figure 106).

Figure 106. Storage Pool Locate

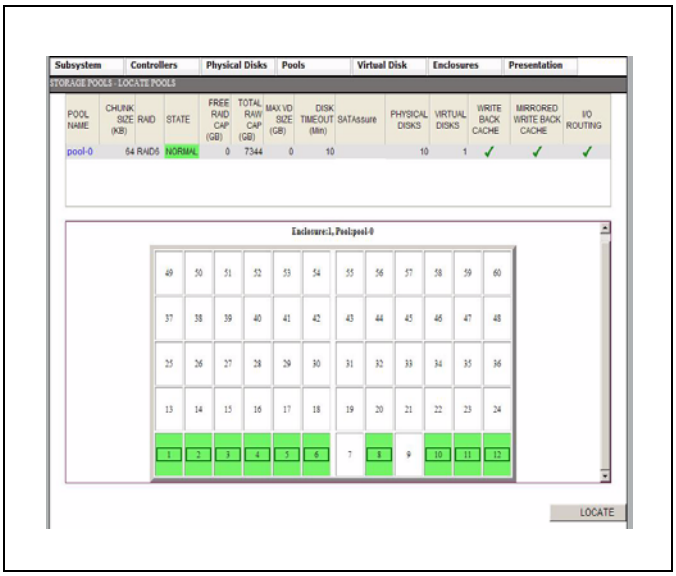

2. Select the storage pool you want to locate.

3. Press the **LOCATE** button.
   A success or failure message displays.

5.5 Virtual Disks

A virtual disk can span a portion of a storage pool or utilize the entire capacity of a single storage pool.

5.5.1 Virtual Disk Commands

Virtual disk commands are located under the **Virtual Disk** tab drop-down menu:

- **Show Virtual Disks**
- **Create Virtual Disk**
- **Set Virtual Disks Attributes**
- **Delete Virtual Disk**

To create a Virtual Disk:

1. Under the **Virtual Disk** tab, select **CREATE Virtual Disk .**
   The Virtual Disk-Create screen displays (Figure 107).

Figure 107. Create Virtual Disk Screen 1

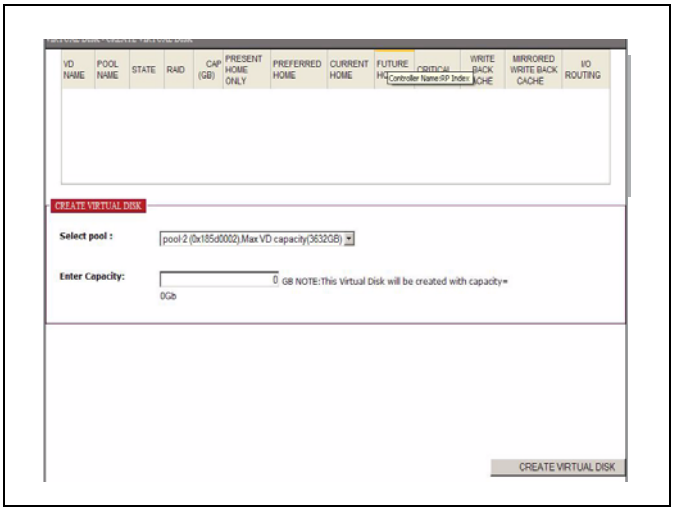

2. Select the pool from the drop-down **Select Pool** menu.
3. Enter **CAPACITY** in GiB.

4. Press the **CREATE VIRTUAL DISK** button.
   The newly created disk is listed (Figure 108).

Figure 108. Create Virtual Disk Screen 2

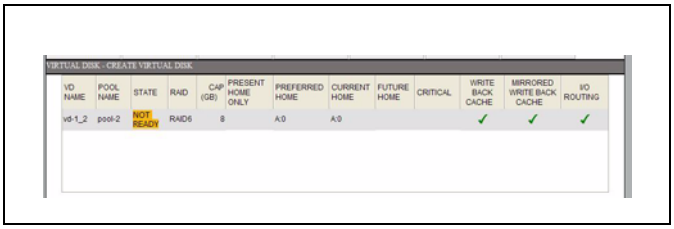

5. Repeat steps 1-4 to add more virtual disks as needed.

6. Under **Virtual Disk** tab, select **SHOW VIRTUAL DISKS** to view the list of configured virtual disks. Select the individual VD Name to display detailed information of the virtual disk selected.

Figure 109. Virtual Disk Detail Screen

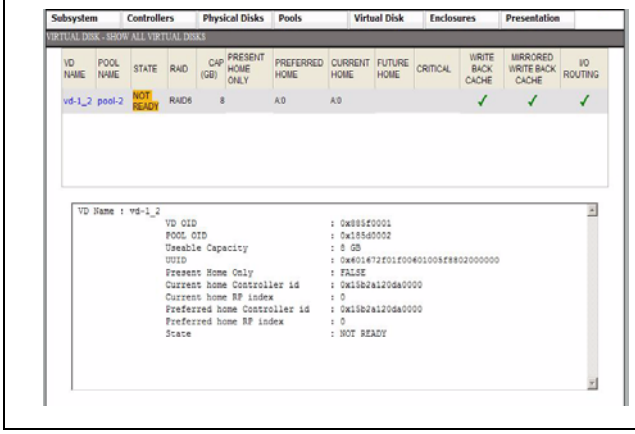

NOTE : A virtual disk on a storage pool is NOT accessible until its initialization is complete.

To delete a Virtual Disk:

1. Under **Virtual Disk** tab, select **DELETE Virtual Disk**. Delete Virtual Disk screen displays.
2. Select the virtual disk you want to delete.
3. Press the **DELETE** button or press **CANCEL** to cancel operation. Success or failure message displays.

To view a list of virtual disks:

Under **Virtual Disk** tab, select **Show Virtual Disk**.

To set virtual disk attributes:

1. Under **Virtual Disk** tab, select **SET VIRTUAL DISK ATTRIBUTES**.
   The Virtual Disks-Set Attributes screen displays (Figure 110).

Figure 110. Virtual Disk Set Attributes

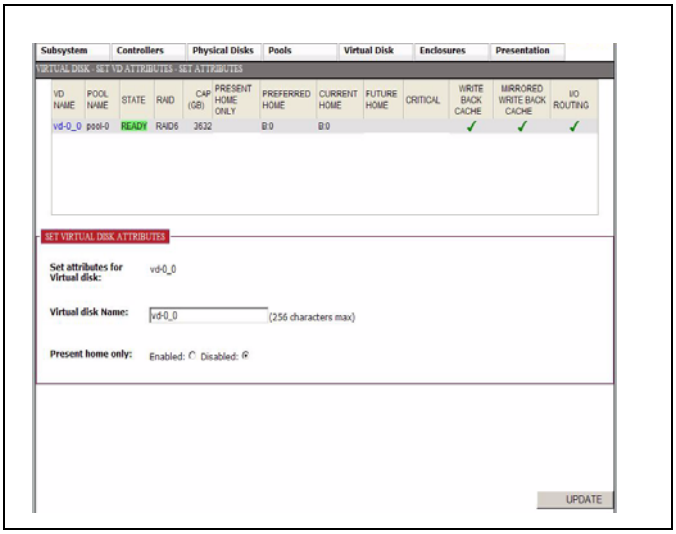

2. Select the virtual disk whose attributes you want to change.
   The Set Attributes screen appears.
3. Make the desired changes.
4. Press the **UPDATE** button.
   A success or failure message displays.

5.6 Spare Pools

The InfiniteStorage 6120 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 drive failure, a drive that is assigned to a spare pool is automatically swapped into the storage pool to replace the failed drive. A rebuild then occurs immediately, reducing the possibility of data loss due to additional drive 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.

5.6.1 Spare Pool Commands

Spare pool commands are located under the **Pool** tab drop-down menu:

- **Show Spare Pools**
- **Create Spare Pool**
- **Set Spare Pools Attributes**
- **Delete Spare Pool**
- **Locate Spare Pool**

To create a spare pool:

1. Under **POOLS** tab, select **CREATE SPARE POOL.**
   The Spare Pools-Create Spare Pool screen displays (Figure 111).

.

Figure 111. Create Storage Pool Screen 1

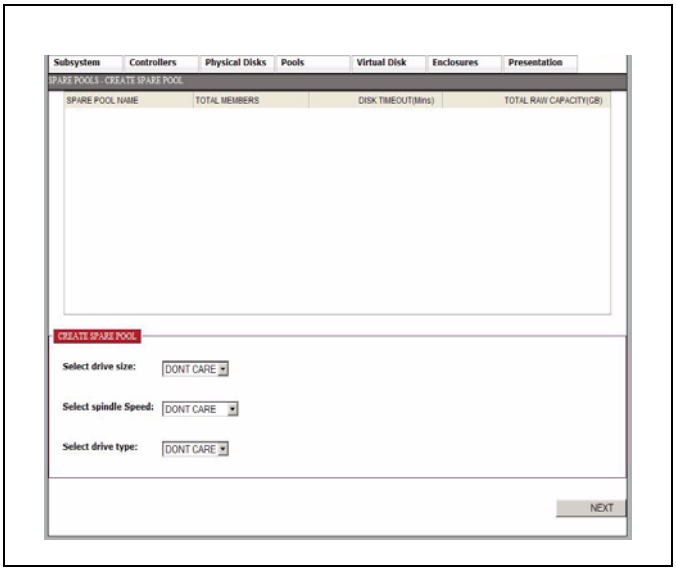

2. Select the drive size, spindle speed, and drive type from their respective menus.

3. Press **Next** button.
   The Storage Pools-Create Storage Pools-Advanced screen displays (Figure 112).

Figure 112. Create Storage Pools Advanced Screen

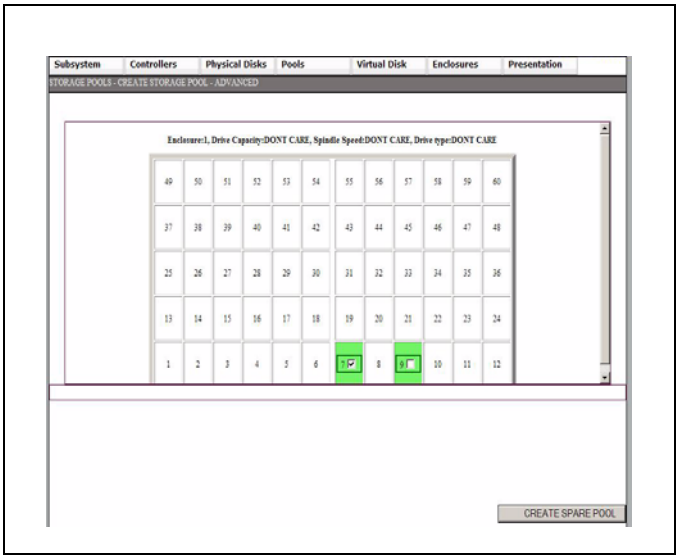

4. Select **PHYSICAL DISK** by checking in the appropriate box.

5. Select **CREATE SPARE POOL**.
   A message is displayed to indicate whether the new spare pool creation was successful.

6. Repeat steps 1 - 7 to add more spare pools as needed.

NOTE : Each storage pool must have a spare pool assigned to it.

To display a list of unassigned spare pools:

1. Under **POOLS** tab, select **SHOW UNASSIGNED POOL**.

2. Under **POOLS** tab, select command **SHOW SPARE POOL** to view the list of configured spare pool(s).

3. To assign the spare pool to a storage pool, select **SET POOL ATTRIBUTES** under **POOLS** tab.

4. Select the name of the pool that you will assign to the new spare pool (Figure 113).

Figure 113. Set Spare Pool Attributes

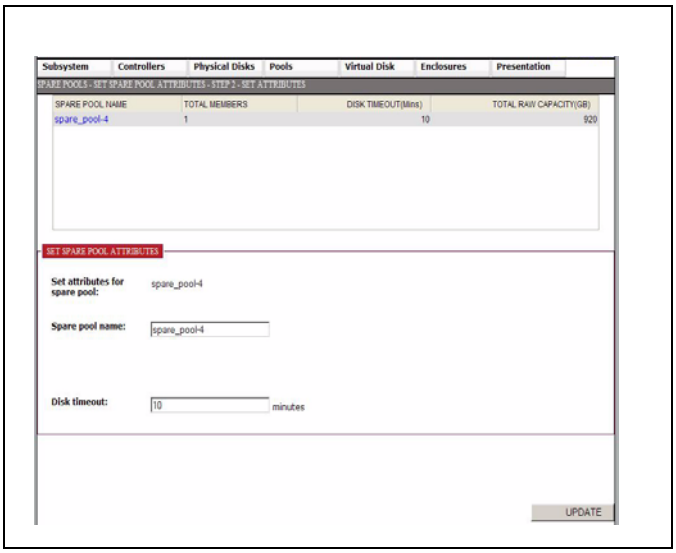

WARNING: When assigning disk drives to a spare pool that is to be assigned to a given storage pool, be sure to select a disk that is as large or larger than any disks in the assigned storage pool.

---

5. In the drop-down for **Assign Spare Pool**, select the name of the spare pool then press the **Update** button.

To show spare pools:

Under **POOLS** tab, select **SHOW SPARE POOLS**.

To delete spare pools:

1. Under **POOLS** tab, select **DELETE SPARE POOL**.
The Spare Pools-Delete Spare Pool screen displays.

2. Select the spare pool you want to delete.

3. Press the **DELETE** button or press **CANCEL** to cancel the operation.
A success or failure message displays.

To set spare pool attributes:

1. Under **POOLS** tab, select **SET SPARE POOL ATTRIBUTES**.
The Spare Pools- Set Attributes screen displays.

2. Select the storage pool whose attributes you want to change.
   The Set Attributes screen appears (Figure 114).

Figure 114. Storage Pool Set Attributes

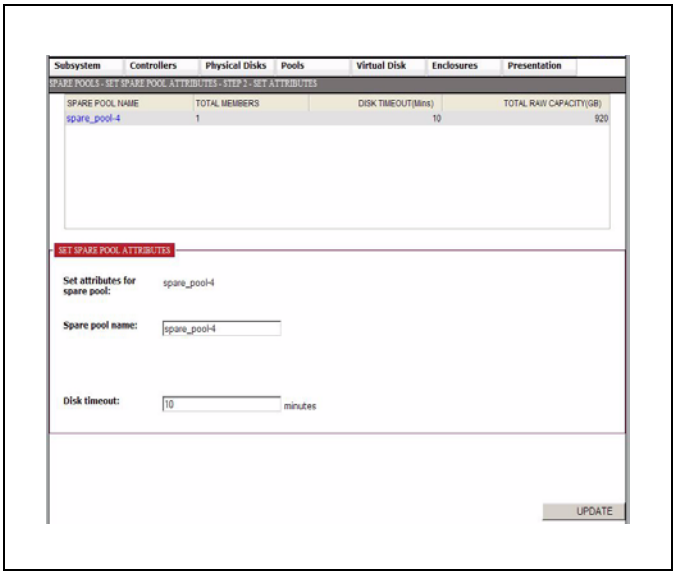

3. Make the desired changes.

4. Press the **UPDATE** button.
   A success or failure message displays.

To locate spare pools:

1. Under **POOLS** tab, select **LOCATE SPARE POOL**.
   The SPARE Pools- Locate screen displays.

2. Select the storage pool you want to locate.

3. Press the **LOCATE** button.
   A success or failure message displays. To display the list of available disks, under the **POOLS** tab, select **SHOW UNASSIGNED POOL**

4. Under the **POOLS** tab, select command **SHOW SPARE POOL** to view the list of configured spare pool(s).

NOTE : Each storage pool must 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.

---

5. To assign the spare pool to a storage pool, select **SET POOL ATTRIBUTES** under the **POOLS** tab.

6. Select the name of the pool that you will assign to the new spare pool.

7. In the drop-down menu for **Assign Spare Pool**, select the name of the spare pool then press the **Update** button.

5.7 Initializing the Host System

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 disks can be presented to host systems. However, you may immediately create the presentations.

To monitor the progress of a job, under **SUBSYSTEM** tab, select the **SHOW BACKGROUND JOBS**.

Figure 115. Subsystem Menu

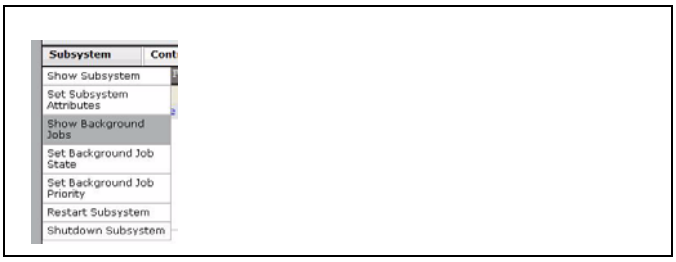

It will display the type of job and the percentage of completion. Select a background job to display its attributes (Figure 116).

Figure 116. Background Job Attributes

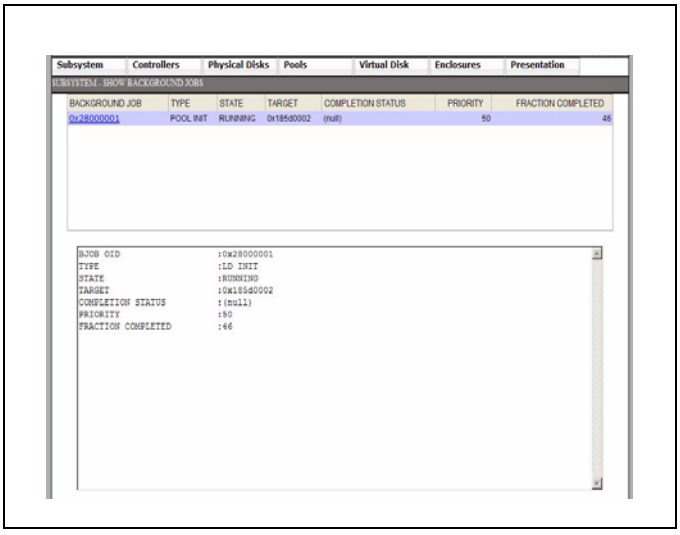

You may also check if a virtual disk is ready for access by selecting **SHOW VIRTUAL DISKS** under the **Virtual Disk** tab (Figure 117).

Figure 117. Show Virtual Disk Readiness

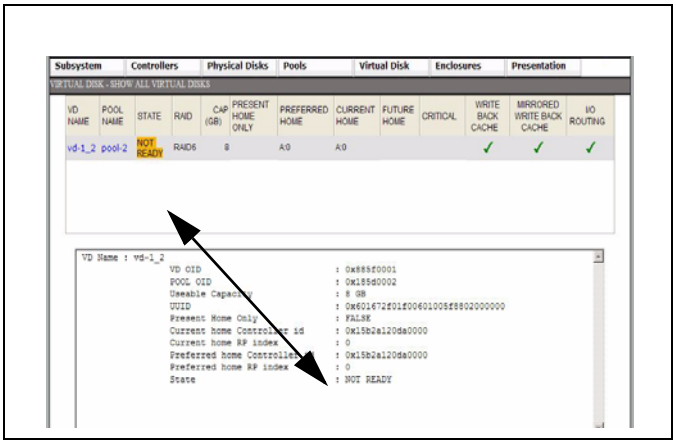

5.8 Presentations

A *presentation* on an InfiniteStorage 6120 has the following components:

- Discovered Initiators
  A *discovered initiator* is defined as a Fibre Channel initiator that has logged into the InfiniteStorage 6120. This is normally a single port on a Fibre Channel 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 initiators. The host attribute OS (operating system) allows you to 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 ports on the controller. The channel attribute MODE allows you to 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 InfiniteStorage 6120, 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. The following are attributes of a presentation:
  - PORT – 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 – LUN number (user-specified) that the virtual disk will show to the host.

5.8.1 Presentation Commands

Presentation commands are located under the **Presentation** tab drop-down menu:

- **Show Presentation**
- **Create Presentation**
- **Delete Presentation**
- **Update Presentation**

To present the virtual disk to an external host:

1. Under the **Presentation** tab, select **Create Host** (Figure 118)**..**

Figure 118. Presentation Menu

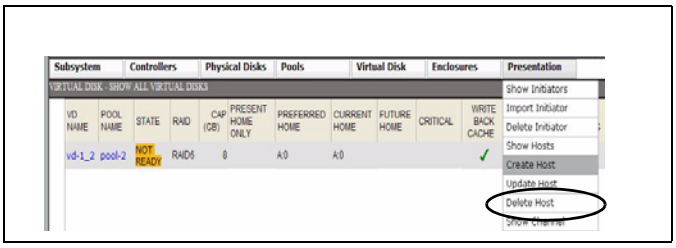

The Presentation-Create Host screen displays (Figure 119).

Figure 119. Presentation-Create Host Screen 1

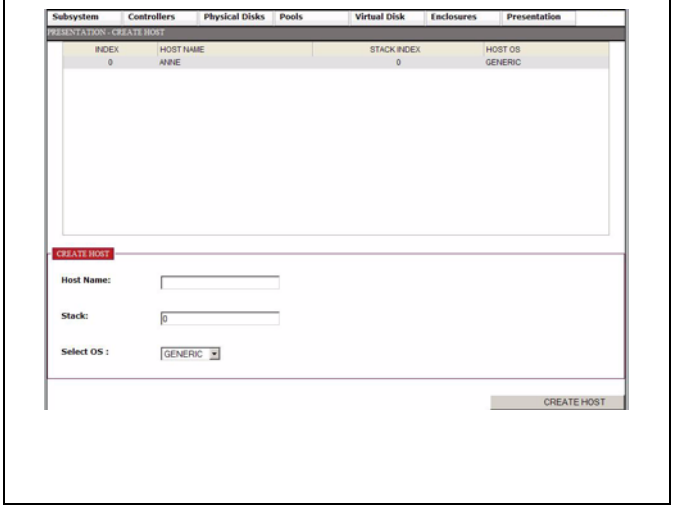

2. Specify the OS attribute using the pull-down menu if necessary.

3. Press the **Create Host** button.
   The system displays a success or error message and a list of hosts created (Figure 120).

Figure 120. Create Host Screen 2

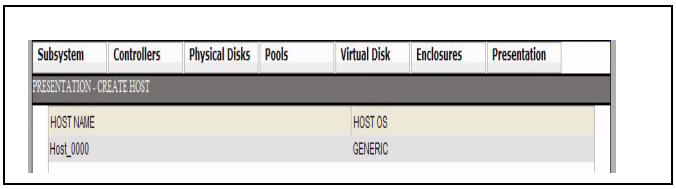

4. Repeat steps 1-3 to create more hosts.

To import one or more discovered initiators to your newly created host:

1. Under the **Presentation** tab, select **`Import Initiators.`**
   The Presentation-Create Host screen displays (Figure 121).

Figure 121. Presentation-Import Initiator Screen 1

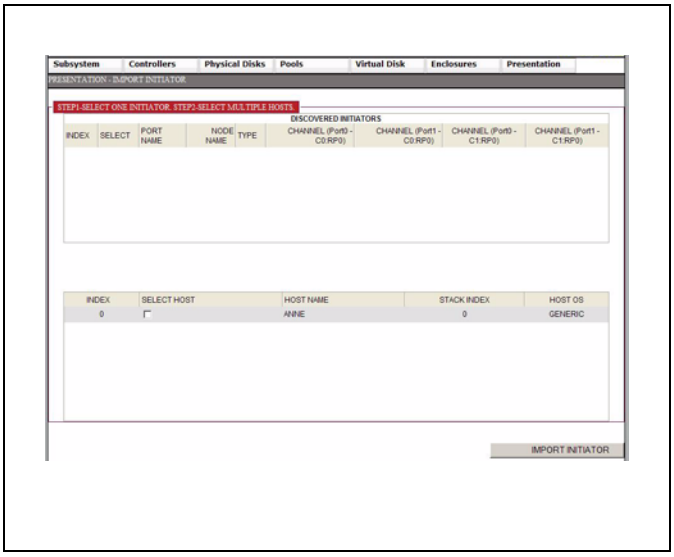

2. Select the initiator.
3. Select the corresponding host.
4. Press the **`Import Initiator`** button.
   A success or error message displays.

To create the presentation:

1. Under the Presentation tab, select **`Create Presentation.`**
   The Presentation-Create Presentation screen displays (Figure 122).

Figure 122. Create Presentation Screen 1

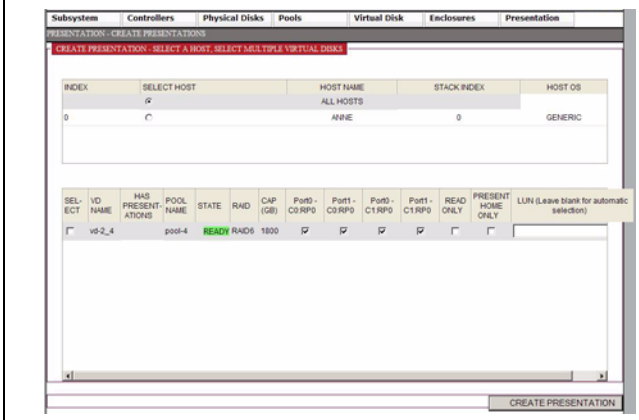

2. Select a host and a virtual disk.

3. Press the **`Create Presentation`** button.
   A success or error message displays.

NOTE : You can view the new presentations by selecting the **Show Presentation** menu selection under the **Presentation** tab (Figure 123).

Figure 123. Show Presentation Screen

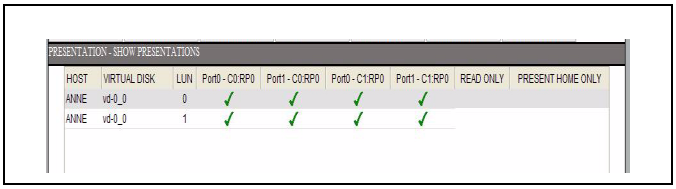

The external computer can access the pool once it has completed its initialization. You may also check if a virtual disk is ready for access by selecting **`SHOW VIRTUAL DISKS`** under the **Virtual Disk** tab.

To present the virtual disks to all hosts on all ports (optional):

1. Under the **Presentation** tab, select **`Create Presentation`**.
   The Presentation-Create Presentation screen displays (Figure 124).

Figure 124. Presentation-Create Presentation Display Screen

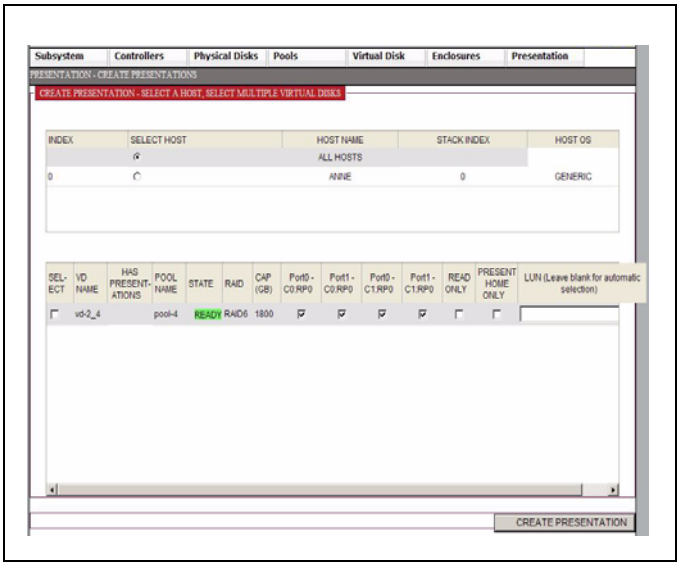

2. Select the **ALL HOSTS** option.

3. Select **Virtual Disks**.

4. Press the **Create Presentation** button.

5.9 Email and Critical Event Notification Setup

Critical event notices will be sent via email. You must provide the email address.

To set up or change an email address:

1. Under the **SUBSYSTEM** tab, select **UPDATE EMAIL SETTINGS**.
   The upper portion of the screen will display the current email agent attributes, if any. Use the lower portion of the screen to enter email information (Figure 125).

Figure 125. Update Email Settings Screen.

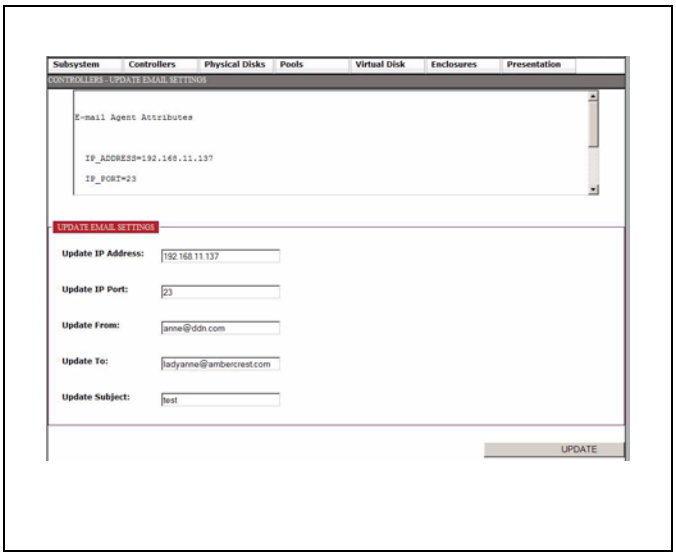

2. In the **Update IP Address** field, enter the IP address of the SMTP server to which the subsystem will send email notifications.
3. In the **IP Port** field, enter the IP port of the SMTP server to which the subsystem will send email notifications.
4. In the **Update From** field, enter an identification 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, enter an Internet Mail Format (RFC2822) email address to which the subsystem will send email notifications.
6. In the **Update Subject** field, enter a string to be included on the `Subject` line of the email notification sent by the email agent.
7. Press the **UPDATE** button.
   A success or failure message displays.

SECTION 6

Appendices

A

Technical Specifications

Here are the technical specifications for the InfiniteStorage 6120.

NOTE : Specifications subject to change without notice.

Configuration, Performance, & Capacity

Host interface	Fibre channel 8Gb/s
Drive interface	SAS / SATA
Management interface	RS-232 and Ethernet (SSH)
Cache memory	6GB
Number of host ports	4
Number of host connections	Up to 512
Number of drives supported	Up to 60 with a single unit, 120 with expansion enclosure
Number of storage pools supported	Up to 24
Number of virtual disks supported	Up to 64 per storage pool, Up to 256 per system
Hot spare capability	Yes
Full duplex 10/100/1000 BaseT onboard	Yes

Reliability

SES (SCSI Enclosure Services) protocol support	Yes
Temperature monitoring	Yes
Redundant hot-swappable power & cooling modules	2

Physical, Power & Environmental

InfiniteStorage 6120 chassis dimensions (H×W×D)	7.0" (18 cm) × 17.56" (45 cm) × 37" (94 cm)
InfiniteStorage 6120 chassis weight	120 lbs (54.5 kg) without drives, 240 lbs (109.1 kg) with drives
Electrical/AC	200-240 V @ 47-63 Hz
Power consumption (maximum current)	10.3A @ 230VAC
Power consumption (average current)	7.5A @ 230VAC
Operating environment (temperature / relative humidity)	5 °C to 35 °C / 20% - 80%, non-condensing
Non-operating environment (temperature/ relative humidity)	-10 °C to 35 °C / 20% - 80%, non-condensing
Cooling (maximum, peak, average)	13324 BTU/h, 6560 BTU/h, 5973 BTU/h
Certification	UL, CE, CUL, FCC

B

Safety Requirements for Rack Installation

The following safety requirements must be considered when the InfiniteStorage 6120 is mounted in a rack.

The rack when configured with the units must meet the safety requirements of UL 60950 and IEC 60950.

Mechanical Loading

- The mounting of the system must conform to even and safe mechanical loading. Uneven mechanical loading could cause a rack to tip and fall.
- The rack construction must be capable of supporting the total weight of the installed enclosure(s) and the design should incorporate stabilizing features suitable to prevent the rack from tipping or being pushed over during installation or in normal use.
- When loading a rack with the units, fill the rack from the bottom up and empty from the top down.
- Always remove all modules and drives, to minimize weight, before loading the chassis into a rack.

WARNING : It is recommended that you do not slide more than one enclosure out of the rack at a time, to avoid danger of the rack tipping over.

- When mounting in a rack, ensure that the enclosure is pushed fully back into the rack.
- The system must be operated with low pressure rear exhaust installation. (The back pressure created by rack doors and obstacles should not exceed 5 pascals [0.5mm water gauge].)

Circuit Overloading

The rack should have a safe electrical distribution system. It must provide over current protection for the unit and must not be overloaded by the total number of units installed in the rack. Consideration of the units nameplate rating should be used when addressing these concerns.

Reliable Earthing

- The electrical distribution system must provide a reliable earth ground for each unit and the rack.
- Each power supply in each unit has an earth leakage current of 1.5mA. The design of the electrical distribution system must take into consideration the total earth leakage current from all the power supplies in all the units. The rack will require labelling with `HIGH LEAKAGE CURRENT. Earth connection essential before connecting supply`.

Elevated Operating Ambient Temperature

If the InfiniteStorage 6120 system is installed in a closed or multi-unit rack assembly, the operating ambient temperature of the rack environment may be greater than room ambient. The rack design should take into consideration the maximum operating ambient temperature for the unit is 35°C.

Air Flow

When installing the system, do not compromise on the amount of air flow required for safe operation of the system.

Index

B

background job, see job 56

C

cable connections 20
cache
  mirrored write back, set 54
  write back, set 54
channel 35, 67, 104
chunk size 49
CLUI, using 39
components, hardware 5
controller 7
  LED 75
cooling module 8
  fault 76, 77
  LED 72, 75
create
  spare pool 32, 51
  storage pool 28
  virtual disk 31
Customer service vi

D

date, set system 54
delete
  storage pool 50
  virtual disk 50
DEM 10
  LED 74
discovered initiator 35, 67, 104
disk
  fail a 48
  failure recovery 80
  fail, clear 48
  information 42
  listing all 42
  location, physical 47
  replace, manual 53
  status, change 48
disk commands, physical 90
disk timeout 52
drive failure recovery 80
drive, see disk 42

E

enclosure
  connecting 20
  location, physical 48
ESI processor 76
Ethernet connection 65
expansion enclosure, connect 20

F

failure
  drive, fixing 79, 80
  fan, alerts 76
  fan, fixing 78
  PCM, fixing 78
  power supply, fixing 78
fan, see cooling module 8
features, S2A6620 3
firmware
  current version 58
  update 58
fractional rebuild 52
front panel LED 11
full rebuild 52

G

gateway, set 65
GUI management system 88
guidelines, rack installation safety 113

H

host 35, 67, 104
  connections 20
  port location 7

I

installation 15
IP address, set controller 65

J

job priority, set 56

L

LED
- controller 75
- cooling 72, 75
- cooling fault 76, 77
- DEM 74
- fault, system 77
- front panel 11
- power supply 72, 75
- power supply fault 76, 77
- system fault 76

licensing instructions 27
login 40
logout 40

M

management system, GUI 88
Manual Intervention Required 82
MIR 82
- clear, state 82
- conditions 82

mirrored write back cache, set 54

N

netmask, set 65
network interface, set 65

P

partial rebuild 52
password 40
PCM 8
- failure recovery 78

physical location
- disk 47
- enclosure 48
- storage pool 48
- virtual disk 48

pool
- spare, create 32, 51
- storage, create 28
- unassigned 42

power supply module 8
- fault 76, 77
- LED 72, 75

presentation
- GUI management 104
- virtual disk to external host 35, 67, 104

priority
- job, set 56
- rebuild, set 57

Product support vi

R

rack installation 17
- safety guidelines 113

RAID level 49
rebuild
- automatic, set 52
- fractional 52
- full 52
- manual, initialize 53
- partial 52
- set priority 57

recovery
- drive failure 80
- PCM 78

replace disk, manual 53
restart
- controller 55
- subsystem 55

RS-232 terminal set up 39

S

safety guidelines, rack installation 113
SATAssure, set 53
serial port, set up 39
shutdown
- controller 55
- subsystem 55

spare pool
- concept 98
- create 32, 51
- delete 51
- global 51
- GUI management of 98
- name, change 51
- view list of 47

sparing policy 52
specifications, technical 111
SSH 65
stack 35, 67, 104
storage pool
- configuration, checking 44

create 28, 49
delete 50
GUI management of 91
initialization 49
list of, view 44
location, physical 48
max number of 48
name, change 49
SATAssure, set 53
sparing policy 52
specifications 91
status 44
system faults LED 76

T

technical specifications 111
Technical support vi
Telnet 65
time, set system 54
troubleshooting 71

U

unassigned pool 42
unpack the system 16
user name 40

V

version information 58
virtual disk
associated with storage pool, view 45
create 31, 50
delete 50
GUI management of 95
location, physical 48
max number of 48
name, change 50
status, checking 46
visual indication
disk 47
enclosure 48
storage pool 48
virtual disk 48

W

write back cache, set 54