



SGI[®] InfiniteStorage 6120 RAID User's Guide

007-5533-001

COPYRIGHT

© 2009 SGI. All rights reserved; provided portions may be copyright in third parties, as indicated elsewhere herein. No permission is granted to copy, distribute, or create derivative works from the contents of this electronic documentation in any manner, in whole or in part, without the prior written permission of SGI.

LIMITED RIGHTS LEGEND

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

The electronic (software) version of this document was developed at private expense; if acquired under an agreement with the USA government or any contractor thereto, it is acquired as “commercial computer software” subject to the provisions of its applicable license agreement, as specified in (a) 48 CFR 12.212 of the FAR; or, if acquired for Department of Defense units, (b) 48 CFR 227-7202 of the DoD FAR Supplement; or sections succeeding thereto. Contractor/manufacturer is SGI, 1140 E. Arques Avenue, Sunnyvale, CA 94085.

TRADEMARKS AND ATTRIBUTIONS

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

Windows is a registered trademark of Microsoft Corporation in the United States and/or other countries.

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

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 200VAC to 240VAC.
- 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
 - 1140 East Arques Avenue
 - Sunnyvale, CA 94085-4602

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 |
| 1.2.3 Power Cooling Module (PCM) | 8 |
| 1.2.4 Drive Modules | 9 |
| 1.2.5 Drive Expander Module (DEM) Card | 10 |
| 1.2.6 LED Indicators | 11 |
| 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 |
| 2.2.2 Rack-Mounting the InfiniteStorage 6120 Chassis | 17 |
| 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 |
| 2.3.5 Checking Status of Physical Disks | 26 |
| 2.3.6 Installing Licensed Features | 26 |
| 2.3.7 Creating Storage Pools | 27 |
| 2.3.8 Creating Virtual Disks | 29 |
| 2.3.9 Creating Spare Pools | 30 |
| 2.3.10 Initializing the Host System | 32 |
| 3.1 Managing the InfiniteStorage 6120 | 35 |
| 3.1.1 Management Interface | 35 |
| 3.1.2 User Logins | 36 |
| 3.1.3 Available Commands | 36 |
| 3.1.3.1 Basic Key Operations | 36 |
| 3.2 Configuration Management | 38 |
| 3.2.1 Monitor Status of Storage Assets | 38 |
| 3.2.1.1 Disk Information | 38 |
| 3.2.1.2 Storage Pool View | 39 |
| 3.2.1.3 Virtual Disk View | 40 |

| | | |
|------------|--|-----------|
| 3.2.1.4 | Spare Pool View | 41 |
| 3.2.1.5 | Visual Indication | 41 |
| 3.2.2 | Storage Pool and Virtual Disk Management | 42 |
| 3.2.2.1 | Changing the Disk's Status | 42 |
| 3.2.2.2 | Creating a Storage Pool | 42 |
| 3.2.2.3 | Storage Pool Initialization | 43 |
| 3.2.2.4 | Naming a Storage Pool | 43 |
| 3.2.2.5 | Deleting a Storage Pool | 43 |
| 3.2.2.6 | Creating a Virtual Disk | 44 |
| 3.2.2.7 | Naming a Virtual Disk | 44 |
| 3.2.2.8 | Deleting a Virtual Disk | 44 |
| 3.2.3 | Spare Pool Management | 45 |
| 3.2.3.1 | Creating a Spare Pool | 45 |
| 3.2.3.2 | Naming a Spare Pool | 45 |
| 3.2.3.3 | Deleting a Spare Pool | 45 |
| 3.2.4 | Drive Rebuild | 46 |
| 3.2.4.1 | Full and Partial Rebuilds | 46 |
| 3.2.4.2 | Sparing Policy | 46 |
| 3.2.4.3 | Manual Drive Replace/Rebuild | 47 |
| 3.2.5 | SATAssure | 47 |
| 3.2.6 | Cache Coherency | 48 |
| 3.2.7 | Setting the System's Date and Time | 48 |
| 3.2.8 | Restarting the InfiniteStorage 6120 | 49 |
| 3.3 | Performance Management | 50 |
| 3.3.1 | Background Job Priority | 50 |
| 3.3.2 | Rebuild Policy Priority | 51 |
| 3.4 | Firmware Update Management | 52 |
| 3.4.1 | Displaying Current Firmware Version | 52 |
| 3.4.2 | Firmware Update Procedure | 52 |
| 3.4.2.1 | Windows Environment | 53 |
| 3.4.2.2 | Linux Environment | 54 |
| 3.5 | The System Logs | 55 |
| 3.6 | Remote Management of InfiniteStorage 6120 | 57 |
| 3.6.1 | Network Connection | 57 |
| 3.6.2 | Network Interface Setup | 57 |
| 3.6.3 | Logins | 57 |
| 4.1 | Status LED Indicators | 61 |
| 4.1.1 | Front Panel Indicators | 61 |
| 4.1.2 | Power Cooling Module (PCM) | 62 |
| 4.1.3 | Controller | 63 |
| 4.1.4 | HDD (Hard Disk Drive) | 63 |
| 4.1.5 | DEM (Drive Expander Module) | 64 |
| 4.2 | Troubleshooting | 65 |
| 4.2.1 | Initial Start-up Problems | 65 |
| 4.2.2 | RAID\$ Prompt Not Shown Correctly | 65 |
| 4.2.3 | System Faults | 66 |
| 4.2.4 | Power Supply/Cooling Faults | 66 |

- 4.2.5 Thermal Alarm 67
- 4.3 Component Failure Recovery** 68
 - 4.3.1 Replacing a Power/Cooling Module 68
 - 4.3.2 Replacing a Drive Module 69
- 4.4 Recovering from a Drive Failure** 70
 - 4.4.1 Automatic Rebuild 70
 - 4.4.2 Drive Replacement 71
 - 4.4.3 Spare Not Available 71
 - 4.4.4 Manual Rebuild 72
- 4.5 Manual Intervention Required (MIR) Conditions** 73
- Appendix A. Technical Specifications** 77
- Appendix B. Safety Requirements for Rack Installation** 79
- Index** 81

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.

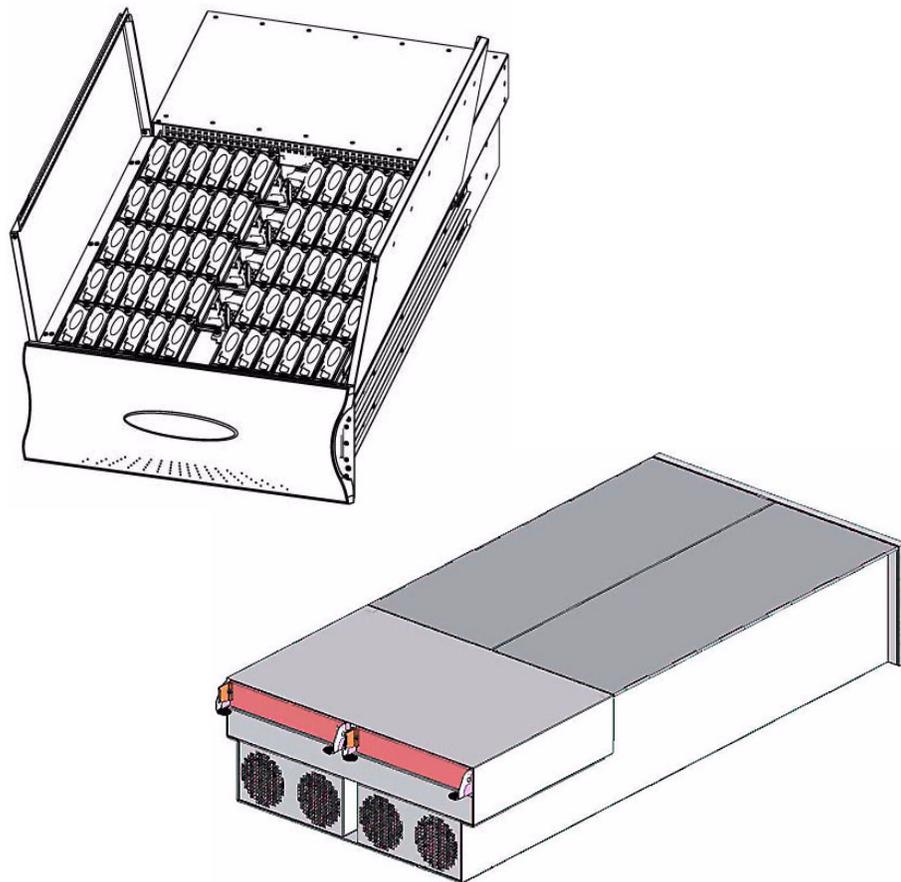
1.2 The InfiniteStorage 6120 System Hardware

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 (SSD). 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

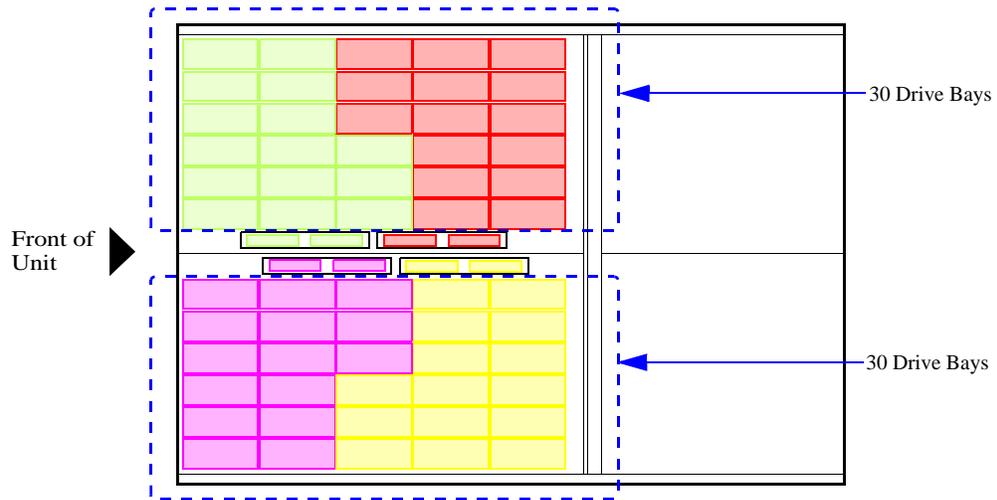
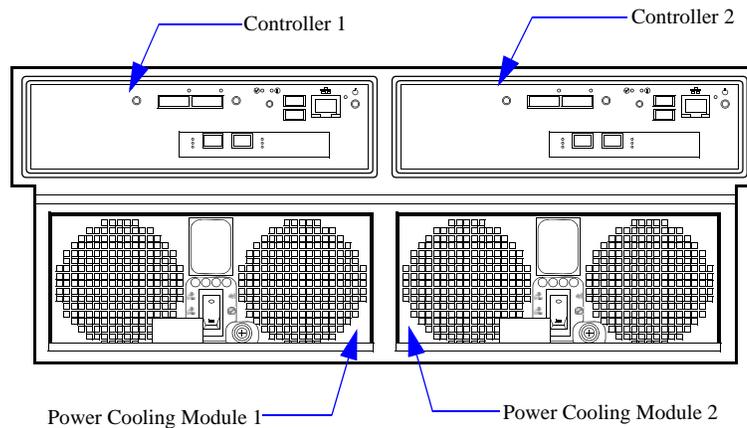


Figure 3. InfiniteStorage 6120 Rear View

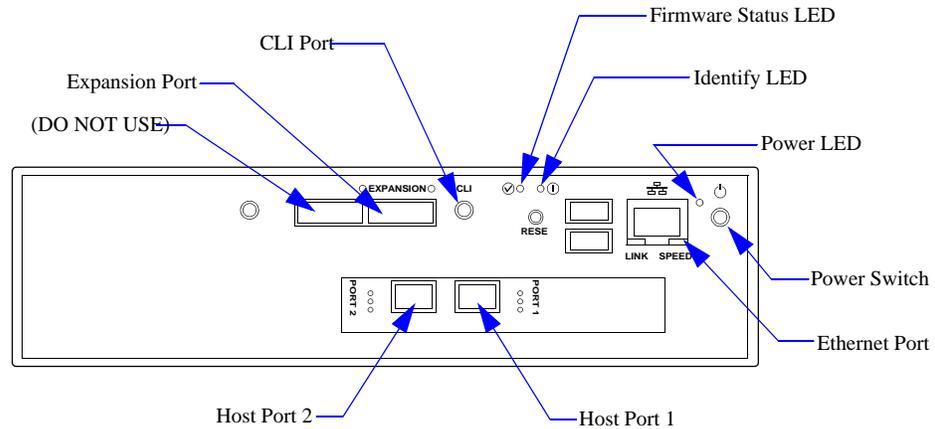


1.2.2 Controller

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.

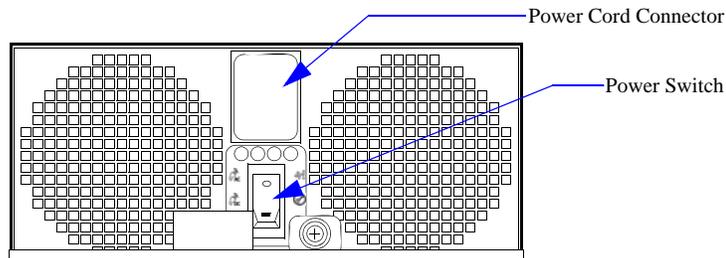
The Power LED turns green when there is power applied to the controller.

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

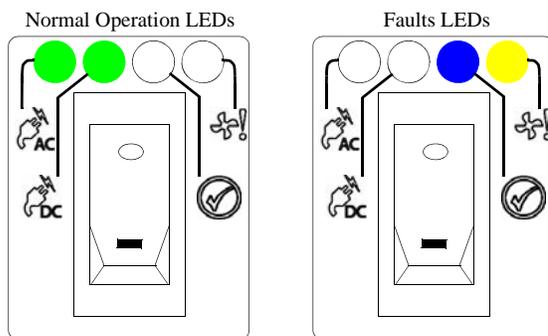


Figure 7. PCM LED Description

| LED | Description | Color | Status |
|---|-------------|-------|---|
|  | PCM AC | Green | ON - AC input to PCM within tolerances OFF - PCM failed |
|  | PCM DC | Green | ON - DC output of PCM within tolerances OFF - PCM failed |
|  | PCM Fault | Amber | ON - PCM fault detected OFF - no detected PCM faults |
|  | PCM ID | Blue | 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.

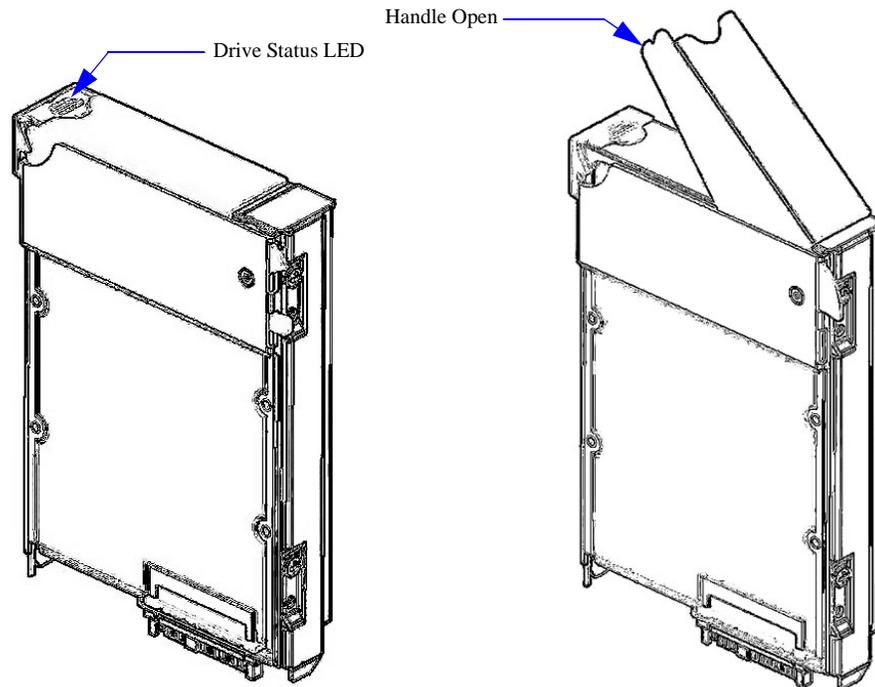
The module handle provides the following functions:

- 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

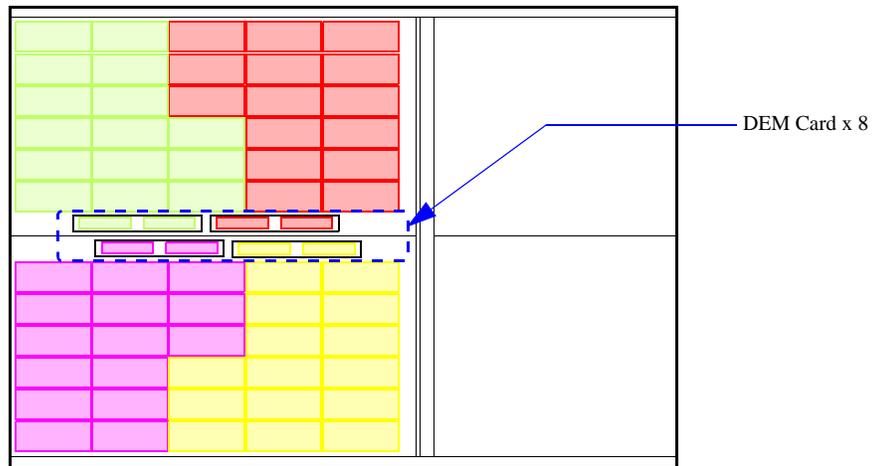


1.2.5 Drive Expander Module (DEM) Card

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.

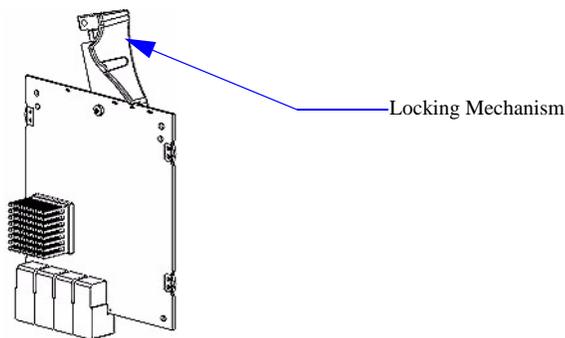
The DEMs are organized into four module pairs, with each module pair controlling a bank of 15 drives. If a DEM card fails, its sister card will provide continuous operation of the unit until a replacement is inserted bringing the unit back to full redundancy.

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

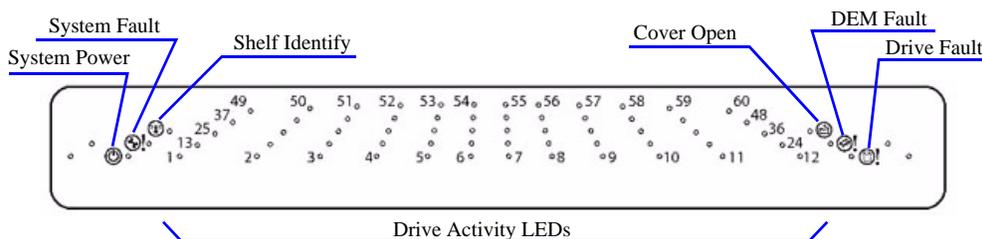


Figure 12. Front Panel LED Description

| 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:

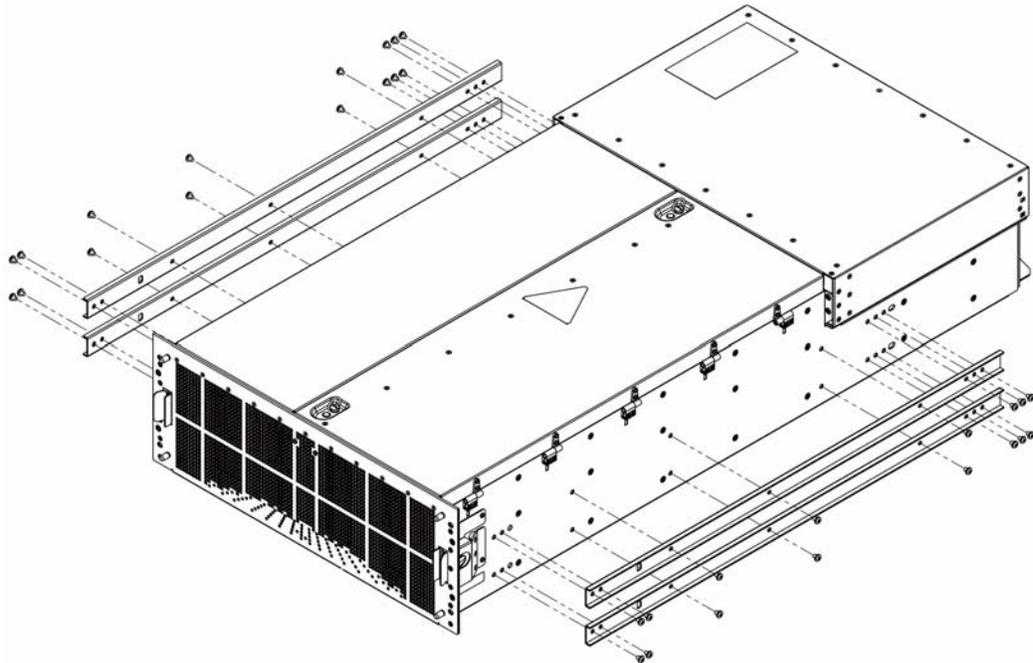
- InfiniteStorage 6120 controller chassis with integrated controllers
- Two (2) power cables
- One (1) mini-jack serial cable for initial configuration
- One (1) enclosure faceplate
- Rack-mounting hardware

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:

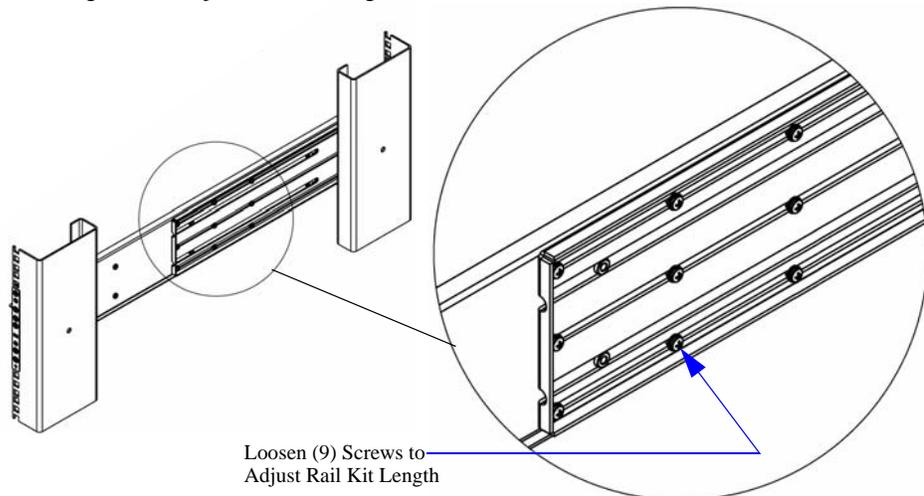
1. Using eight (8) 10-32×0.188" Phillips panhead screws for each slider, attach the four (4) inner sliders to the chassis as shown in [Figure 14](#) below.

Figure 14. Attach 4 Inner Sliders to Chassis



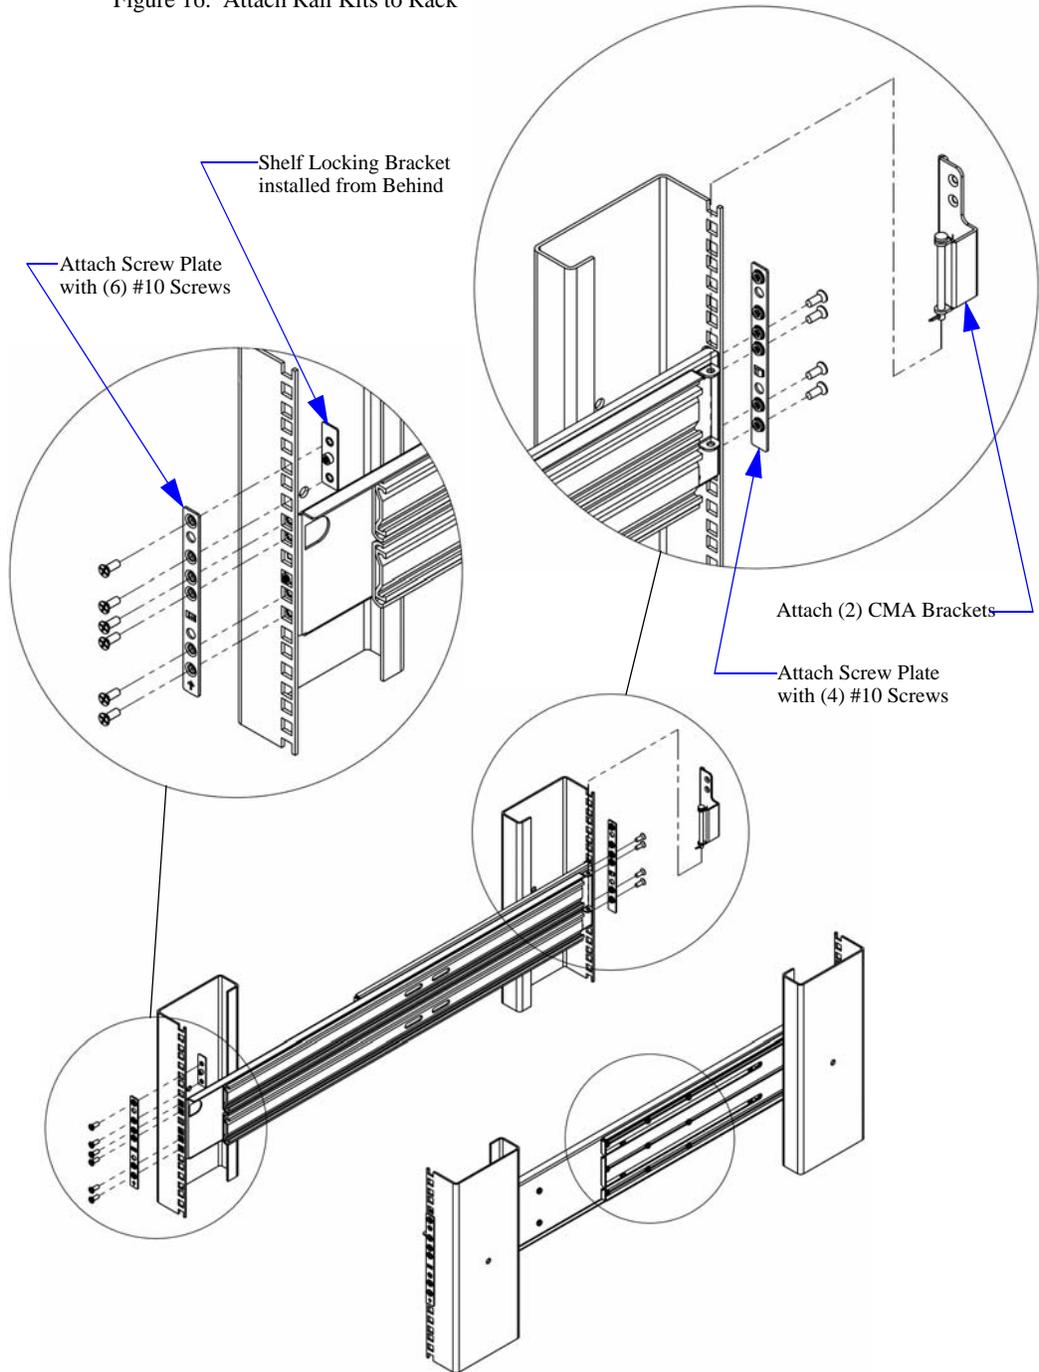
2. For expansion enclosure, skip this step and proceed to the next step. Loosen the nine (9) screws on the rail kit and set the length to fit the depth of your rack ([Figure 15](#)). Tighten all the screws once the rail kit is set to the desired length. Repeat for the other rail kit.

Figure 15. Adjust Rail Kit Length



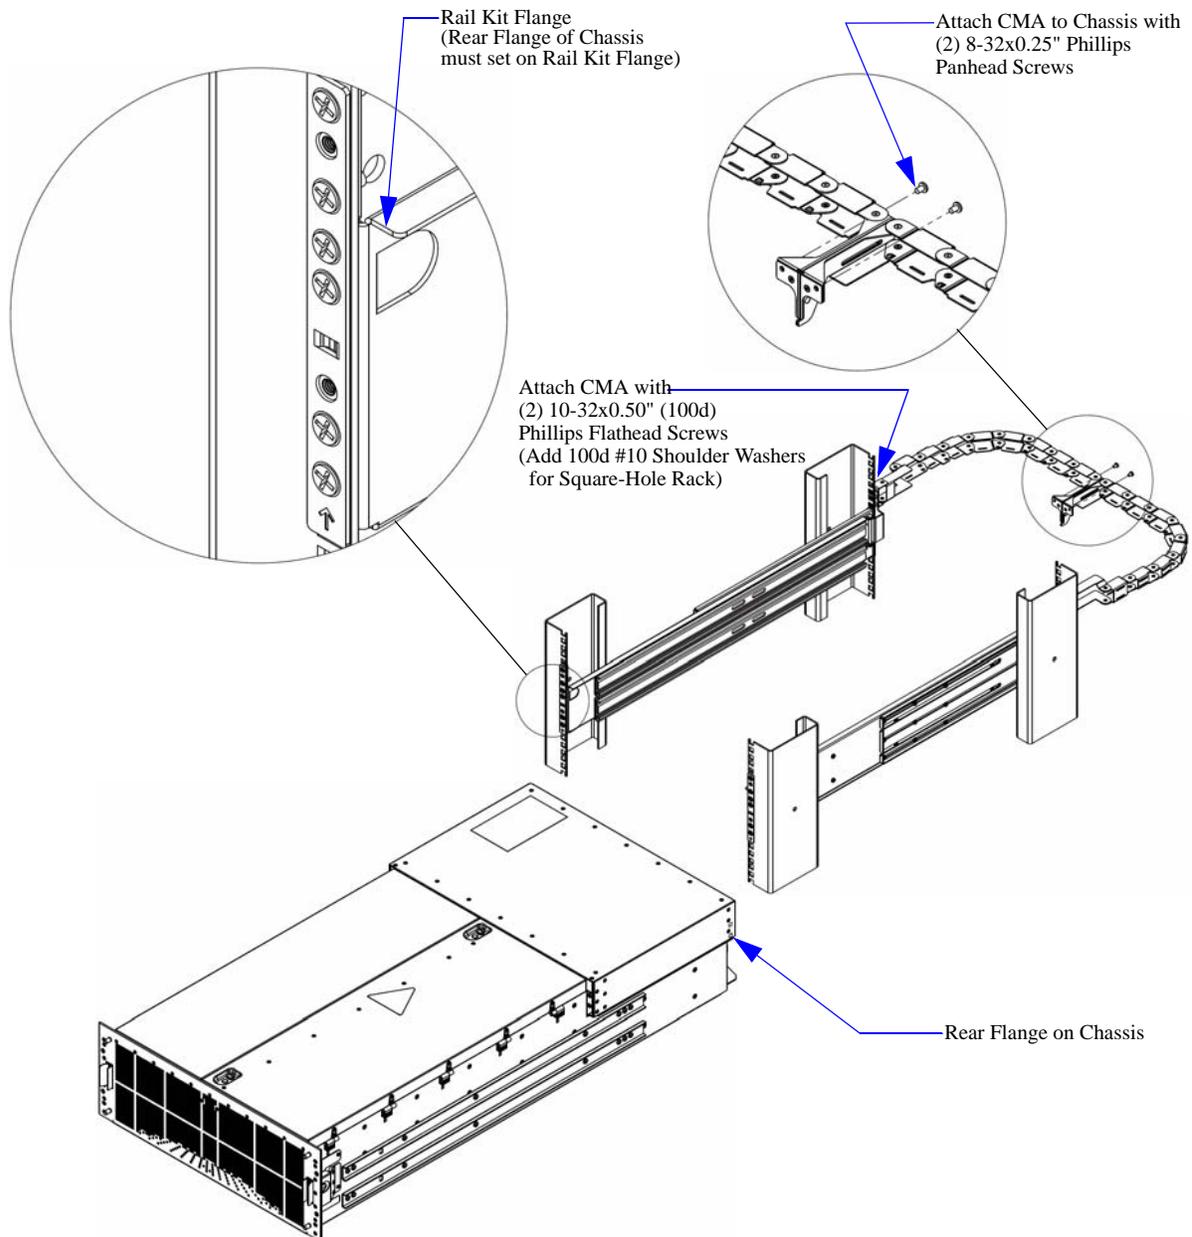
3. Attach the two (2) rail kits to your rack as shown in [Figure 16](#). Make sure you leave enough space for the unit—bottom of screw plate aligns to bottom of InfiniteStorage 6120 chassis. At the front, use one (1) screw plate, one (1) shelf locking bracket (installed from behind), and six (6) #10 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 (647mm), 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 25.5 inches (647 mm) 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

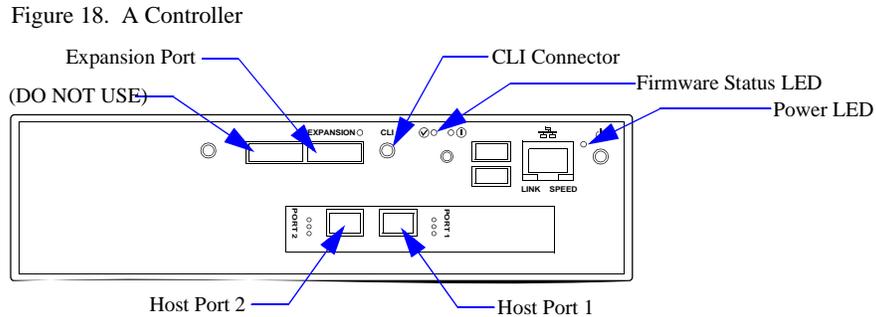


2.2.3 Connecting the InfiniteStorage 6120

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.

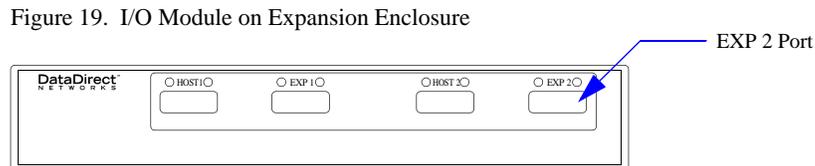


2.2.3.2 Expansion Enclosure Connections

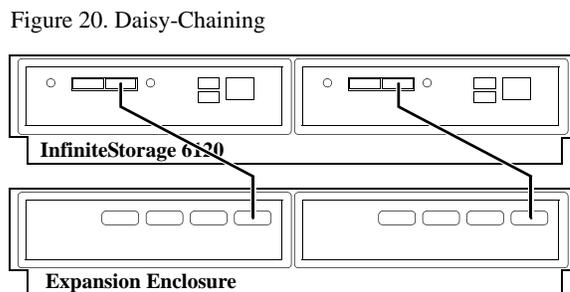
An expansion enclosure can be daisy-chained to the InfiniteStorage 6120 to increase storage capacity. Follow these steps to connect an expansion enclosure:

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).



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

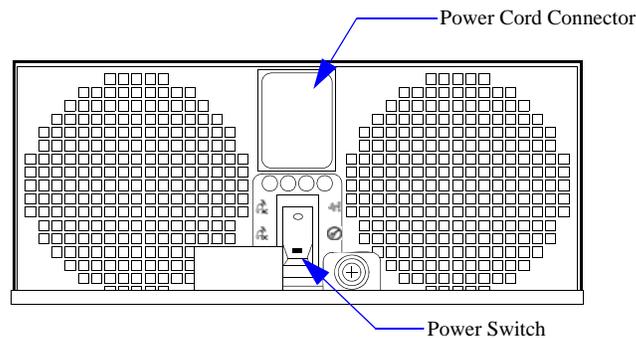


2.2.3.3 Power Connections

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.

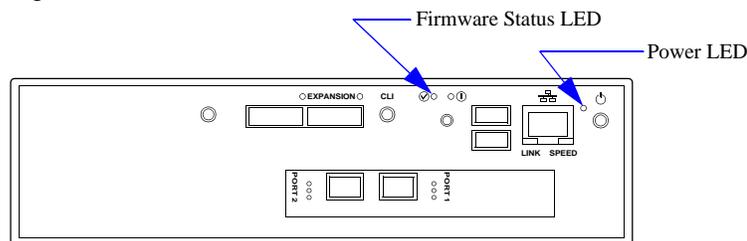
Figure 21. Power Cooling Module



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 8GiB increments. For example, 16GiB of storage space will be allocated when creating a virtual disk of 10GiB. 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



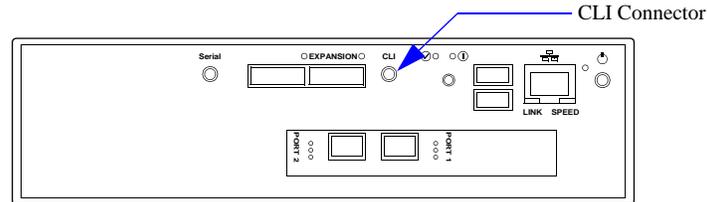
2.3.3 Initial Configuration via RS-232

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$ 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'
```

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

Figure 26. Show Network Interface Settings

```
RAID$ 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$ 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. Enter command **SHOW UNASSIGNED_POOL ALL_ATTRIBUTES** (Figure 28).

Figure 28. Show Unassigned Pool Example Screen

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

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 29).

Figure 29. Physical Disk List Example Screen

```
RAID$ 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:   GK80AB0A
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 SATAAssure) and activate the associated licenses.

License key generation requires that you have your feature license number, JBOD enclosure serial number, and controller ID information. 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.

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, SGI recommends the use of 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).

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

To create a storage pool, use the **CREATE POOL** command:

1. At the CLUI prompt, type:

```
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 (Figure 30).

Figure 30. CREATE POOL Example (1)

```
RAID$ 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 31), 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 31. CREATE POOL Example (2)

```
RAID$ 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 32).

Figure 32. SHOW POOL Example Screen

```
RAID$ show pool *
OID: 0x19b40000 Index: 0x0000 Name: 19b40000
OID: 0x19b60001 Index: 0x0001 Name: 19b60001
```

SHOW POOL <id> ALL_ATTRIBUTES will display detailed information of the storage pool (Figure 33).

Figure 33. SHOW POOL ALL_ATTRIBUTES Example Screen

```
RAID$ 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 34).

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

Figure 34. CREATE VIRTUAL DISKS Example Screen

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

RAID$ 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 35).

Figure 35. SHOW VIRTUAL_DISK List Example Screen

```
RAID$ show virtual_disk *
OID: 0x89ba0000 Index: 0x0000 Name: 89ba0000
OID: 0x89ba0001 Index: 0x0004 Name: 89ba0001
```

SHOW VIRTUAL_DISK <id> ALL_ATTRIBUTES will display detailed information of the virtual disk (Figure 36).

Figure 36. SHOW VIRTUAL_DISK ALL_ATTRIBUTES Example

```
RAID$ 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:                60001ff0722b200000000000089ba0000
```

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 37)

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

Figure 37. CREATE SPARE POOL Example Screen

```
RAID$ create spare_pool
SPARE_POOL 6 OID=0x19bc0006 create STATUS='Success'
```

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 38).

Figure 38. Assign Physical Disk to Spare Pool Example Screen

```
RAID$ assign physical_disk 0x004e to_pool 6
PHYSICAL_DISK 78 OID=0x20b6004e assigned to POOL 6 OID=0x19bc0006STATUS='Success'
```

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 39).

Figure 39. Show Spare Pool Example Screen (1)

```
RAID$ show spare_pool *
OID: 0x19bc0006 Index: 0x0006 Name: 19bc0006
```

SHOW SPARE_POOL <id> ALL_ATTRIBUTES will display detailed information of the spare pool (Figure 40). **SHOW SPARE_POOL <id> PHYSICAL_DISKS** will display the list of disks in the pool.

Figure 40. Show Spare Pool Example Screen (2)

```
RAID$ 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
```

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 41). Repeat this step for each storage pool in the system.

Figure 41. Set Storage Pool to Spare Pool Example Screen

```
RAID$ set pool 0 spare_pool 6
POOL 0 OID=0x19b40000 attributes set with STATUS='Success'
```

6. Use the **SHOW POOL** command to view the new storage pool information (Figure 42).

Figure 42. Show Pool Example Screen

```
RAID$ 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 43).

Figure 43. Show Job Example Screen

```
RAID$ 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 44).

Figure 44. Show Virtual Disk Example Screen

```
RAID$ 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
:
:
:
```

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 1.0.3.3.xxxx. To access the most up-to-date commands, use the CLUI's online help feature.

3.1.1 Management Interface

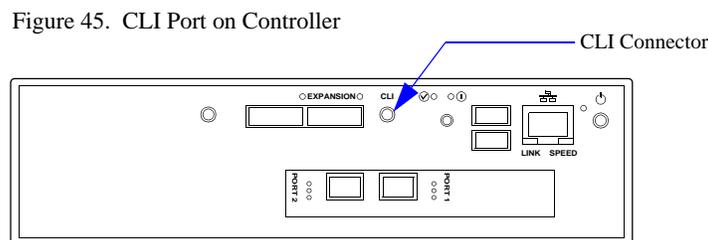
SAN management information for the InfiniteStorage 6120 can be accessed locally through a serial interface or remotely through SSH.

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 45). Connect the other end of the cable to a serial port on a standard PC.



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 46).

Figure 46. 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 47). 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 47. Help Screen

```
RAIDS ?
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 48.

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 48](#).

Figure 48. 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 49). **SHOW UNASSIGNED_POOL PHYSICAL_DISK** displays the list of unassigned disks.

Figure 49. Physical Disk List Example Screen

```
RAID$: show physical_disk *
OID: 0x20a6003e Index: 0x003e
OID: 0x20a7003f Index: 0x003f
OID: 0x20a80040 Index: 0x0040
:
:
OID: 0x220a0190 Index: 0x0190
OID: 0x220b0191 Index: 0x0191
OID: 0x220c0192 Index: 0x0192
```

SHOW PHYSICAL_DISK=<disk-id> ALL retrieves information about a specific disk (Figure 50).

Figure 50. Physical Disk Information Example Screen

```
RAID$ 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: GK80AB0A
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 51).

Figure 51. Failed Drive Example Screen

```
RAID$ 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: GK80AB0A
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 52).

Figure 52. Storage Pool List Example Screen

```
RAID$ show pool *
OID: 0x19b40000 Index: 0x0000 Name: 19b40000
OID: 0x19b60001 Index: 0x0001 Name: 19b60001
```

SHOW POOL <id> PHYSICAL_DISK displays the list of disks associated with the specified storage pool (Figure 53).

Figure 53. Storage Pool's Physical Disks Example Screens

```
RAID$ show pool 0 physical_disk
0x003e NORM 0x5000cca215c56e47
0x003f NORM 0x5000cca215c56456
0x0040 NORM 0x5000cca215c54c71
0x0041 NORM 0x5000cca215c5675c
0x0042 NORM 0x5000cca215c56e55
```

SHOW POOL <id> VIRTUAL_DISK displays the list of virtual disks associated with the specified storage pool (Figure 54).

Figure 54. Storage Pool's Virtual Disks Example Screens

```
RAID$ show pool 0 virtual_disk
OID: 0x89ba0000 Index: 0x0000 Name: 89ba0000
OID: 0x89bb0001 Index: 0x0001 Name: 89bb0001
```

SHOW POOL <id> ALL displays detailed information of the storage pool (Figure 55).

Figure 55. Storage Pool Details Example Screen

```
RAID$ 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 56).

Figure 56. Virtual Disk List Example Screen

```
RAID$ show virtual_disk *

OID: 0x89ba0000 Index: 0x0000 Name: 89ba0000
OID: 0x89ba0001 Index: 0x0004 Name: 89ba0001
```

SHOW VIRTUAL_DISK <id> ALL_ATTRIBUTES will display detailed information of the virtual disk (Figure 57).

Figure 57. Virtual Disk Details Example Screen

```
RAID$ 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:             60001ff0722b2000000000089ba0000
```

3.2.1.4 Spare Pool View

SHOW SPARE_POOL * will display the list of configured spare pool(s) (Figure 58).

Figure 58. List of Spare Pools Example Screen

```
RAID$ show spare_pool *
OID: 0x19bc0006 Index: 0x0006 Name: 19bc0006
```

SHOW SPARE_POOL <id> ALL_ATTRIBUTES will display detailed information of the spare pool (Figure 59). **SHOW SPARE_POOL <id> PHYSICAL_DISKS** will display the list of disks in the pool.

Figure 59. Spare Pool Example Screen

```
RAID$ 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 8GiB increments. For example, 16GiB of storage space will be allocated when creating a virtual disk of 10GiB. 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 60).

Figure 60. Storage Pool List Example Screen

```
RAID$ 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 61).

Figure 61. Set Pool Name Example Screen

```
RAID$ 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 62).

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

Figure 62. Create Virtual Disks Example Screen

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

RAID$ 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 63).

Figure 63. Set Virtual Disk Name Example Screen

```
RAID$ set virtual_disk 0 name="LUN0"
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 64).

Figure 64. Show Pool All Attributes

```
RAID$ show pool 3 all
OID:                0x19b003
Type:               STORAGE
Name:              RAID5set
Chunk Size:        64KB (0x80 blocks)
Block Size:        0x200
RAID Type:         RAID5
Free Raid5 Capacity: 2752512 MBs
Max VD Size:       2752512 MBs
Total Capacity:    3522560 MBs
UUID:              0x00
Global Spare Pool: 0x19bc0006
DiskTimeout(FRT):  10 minutes
:
:
```

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 65).

Figure 65. Show Pool All Attributes

```

RAID$ show pool 3 all
  OID:                0x19b003
  Type:               STORAGE
  :
  :
  :
  UUID:              0x00
  Global Spare Pool: 0x19bc0006
  DiskTimeout(FRT):  10 minutes
  Init Policy:       NO IO permitted
  Init Priority:     50%
  Full Rebuild Priority: 80%
  Fractional Rebuild Priority: 90%
  Sparing Policy:    AUTOMATIC
  :
  :

```

To change the sparing policy setting, use this command:

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

The default setting is automatic which is recommended.

A rebuild operation can take up to several hours to complete, depending on the size of the disk and rate of rebuild. Refer to Section 3.3, "Performance Management" on page 50 for information on how to adjust the rate of rebuild.

3.2.4.3 Manual Drive Replace/Rebuild

REPLACE PHYSICAL_DISK <id> NEW_DISK <new-disk-id> replaces the specified failed disk with the specified new disk.

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, enter the command

SHOW POOL <pool-id> ALL_ATTRIBUTES .

Figure 66. Show Storage Pool Details Example Screen

```

RAID$ show pool 0 all_attributes
  OID:                0x19b40000
  Type:               STORAGE
  Name:               0x19b40000
  :
  :
  Assign Policy:
  Device Type:       NA
  Rotation Speed:   NA
  Raw Capacity:     NA
  SATAssure:        FALSE
  IO ROUTING:       TRUE
  WBC:              TRUE
  MWBC:             FALSE
  Initializing:     TRUE
  Rebuilding:       FALSE
  :
  :

```

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.



Warning

Mirrored Write Back Cache performance is limited by the speed of the internal Ethernet as cache data is transmitted over this connection.

The write back cache settings are configurable for each storage pool. To display the current settings, enter command **SHOW POOL <id> ALL_ATTRIBUTES** (Figure 67).

Figure 67. Show Write Back Cache Setting

```

RAID$ show pool 3 all
OID:                0x19b003
Type:               STORAGE
Name:              RAID5set
Chunk Size:        64KB (0x80 blocks)
.
.
.
SATAssure:         FALSE
IO ROUTING:        TRUE
WBC:               TRUE
MWBC:              FALSE
Initializing:      FALSE
Rebuilding:        TRUE
Paused:            FALSE
.
.
.

```

WBC=Write Back Caching
MWBC=Mirrored Write Back Caching

To change the settings, use these commands:

```

SET POOL <id> MIRRORRED 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 68).

Figure 68. Show System Time

```

RAID$ 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 69), enter this command:

SHOW JOB * ALL_ATTRIBUTES

The type of job and percentage of completion are displayed.

Figure 69. Show Job Example Screen

```
RAID$ 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 70).

To resume the job, enter command **RESUME JOB=<id>**

Figure 70. Pause/Resume a Job

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

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

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

RAID$ 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 71).

Figure 71. Show Pool Information

```
RAID$ 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 72), 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 72. Set Full Rebuild Priority

```
RAID$ 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 73), 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 73. Set Partial Rebuild Priority

```
RAID$ 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 74).

Figure 74. Controller Information Screen

```
RAID$ 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

The firmware update procedure described here only applies to upgrades from firmware version 1.0.3.3.xxxx to a new version. Upgrades from an older version of firmware must be done by a trained technician.

NOTE : 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.



Warning

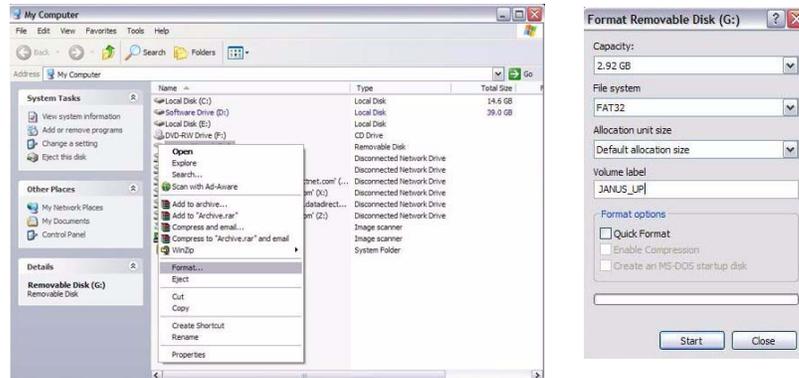
Firmware update should only be performed when both controllers are powered on and running.

3.4.2.1 Windows Environment

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

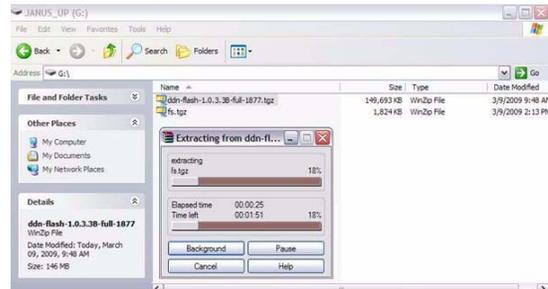
1. Create a USB stick for the code update:
 - Insert your 4GB USB stick into your computer.
 - In Windows Explorer, right-click on the device with the Volume Name “JANUS_UP” and select **Format** from the pop-up menu (Figure 75). Then click the **Start** button in the Format screen.

Figure 75. 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 76).

Figure 76. Firmware Update Example Screen (2)



2. Insert the USB stick into the back of the controller.
3. At the CLUI prompt, enter command **UPDATE CONTROLLER LOCAL** (Figure 77)

Figure 77. Update Firmware Screen

```
RAID$ 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
umount: /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!
```

4. Upon completion of reboot, login and enter command **SHOW CONTROLLER LOCAL ALL**. Verify that the firmware version is correct.
5. Repeat the above steps on the other controller. You do not need to back out either controller during this procedure.

3.4.2.2 Linux Environment

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 sgi-flash-1825-opt.tgz firmware@10.32.31.240:
```

2. At the CLUI prompt, enter this command :

```
UPDATE_FIRMWARE CONTROLLER=<controller-id> FILE="<file-specification>"
```

For example:

```
UPDATE_FIRMWARE CONTROLLER=LOCAL FILE="sgi-flash-1825-opt.tgz"
```

(Note that the file name must be enclosed within double quotation marks.)

3.5 The System Logs

To display the event log starting at the start sequence number in ascending sequence number order (Figure 78), enter this command:

```
SHOW CONTROLLER <id> LOG ASCEND
```

Figure 78. Controller Log Example Screen (1)

```
RAID$ 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 JDEBIAN 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 79), enter this command:

```
SHOW CONTROLLER <id> LOG DESCEND
```

Figure 79. Controller Log Example Screen (2)

```
RAID$ 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 80), enter this command:

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

Figure 80. Controller Log Example Screen (3)

```
RAID$ 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.
```

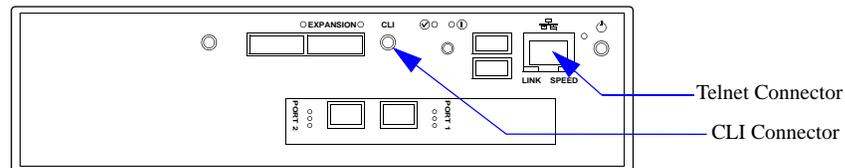

3.6 Remote Management of InfiniteStorage 6120

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 81). Then configure the network interface as described below.

Figure 81. 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 82). **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 82. Current Network Interface Settings Screen

```
RAID$ 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 you are connected to, enter this 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 83).

Figure 83. Set Network Interface Example

```
RAID$ 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)
```

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

3.6.3 Logins

By default, the login name is `user` and its password is `user`.

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.

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 84). The LEDs are defined in Figure 85.

Figure 84. Front Panel LED Indicators

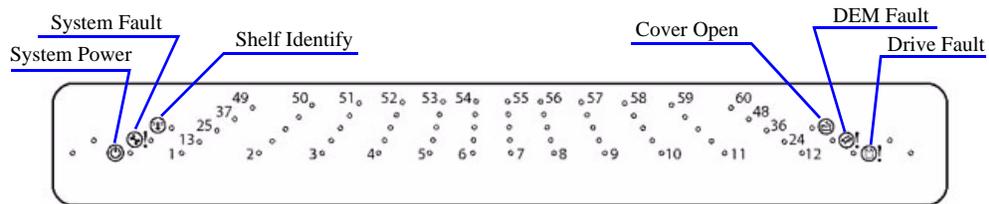


Figure 85. Front Panel LED Description

| 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 |

Figure 85. Front Panel LED Description

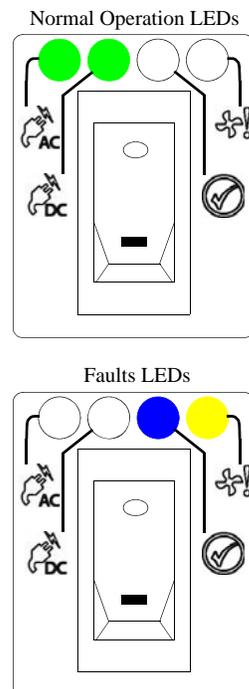
| LED | Definition | Color | Status |
|---|-------------|-------|--|
|  | DEM Fault | Amber | OFF - all DEMs operating correctly ON - at least one DEM has failed; service action required |
|  | Drive Fault | Amber | ON - one or more HDDs are failed; SES must determine exact HDD OFF - no detectable drive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

4.1.2 Power Cooling Module (PCM)

Figure 86 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 86. Power Cooling Module LEDs



| LED | Description | Color | Status |
|---|-----------------|-------|--|
|  | Enclosure front | 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 88 describes the controller LED (Figure 87) 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 87. LEDs on Controller

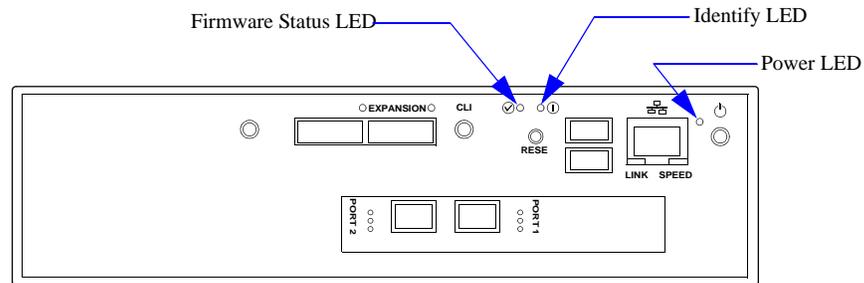


Figure 88. Controller LED Indicators

| 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 89 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 89. 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 90 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 90. 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\$ Prompt Not Shown Correctly*

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

- The firmware on this controller is in a *MIR* state and the *MIR* state needs to be resolved. Refer to Section 4.5, "[Manual Intervention Required \(MIR\) Conditions](#)" on page 73 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 |
|---|---|---|
| <p>1. The System Fault LED turns amber on front panel.</p> | <p>The ESI processor has detected an internal fault (for example, failure of an internal communications path)</p> | <p>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.</p> <p>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.</p> |

4.2.4 Power Supply/Cooling Faults

| Symptom | Cause | Action |
|--|--|---|
| <p>1. System Fault LED is amber on front panel.</p> <p>2. An amber LED on one or more power cooling modules.</p> | <p>1. Any power fault</p> <p>2. A fan failure</p> <p>3. A thermal condition which could cause PSU overheating</p> | <p>1. Ensure that the power On/Off switch on rear of power cooling module is turned ON.</p> <p>2. Ensure that the AC main connection to the power cooling module is live.</p> <p>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.</p> <p>4. Reduce the ambient temperature.</p> |

4.2.5 Thermal Alarm

| Symptom | Cause | Action |
|--|--|---|
| <ol style="list-style-type: none"> 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. | <p>If the air temperature through the enclosure exceeds a pre-set threshold, a thermal alarm will sound.</p> | <ol style="list-style-type: none"> 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 front. 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.



Warning

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.



Warning

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.

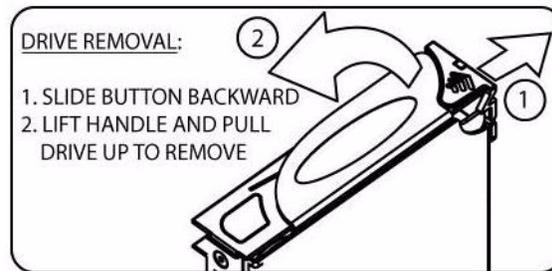
1. Identify the failed PCM with an amber Fault LED.
2. Turn off the module's power switch and disconnect the power cord.
3. Remove the module's thumbscrew, then slide the module out of the bay.
4. Ensure that the power switch on the new (replacement) module is OFF.
5. Slide the module into the bay making 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 91. 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 91).
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 92), use the **SHOW CONTROLLER LOCAL LOG** command.

Figure 92. Display Event Log

```

RAID$ show controller local log
000041 2009-02-04 15:03:17:7295774 G=4 S=0 T=1 RP=0 VP=1
LOG_ST_POOL_CHANGE STATE POOL CHANGE POOL:0003 STATE:0005
000042 2009-02-04 15:03:17:7295780 G=4 S=0 T=1 RP=0 VP=1
LOG_ST_MEMBER_CHANGE STATE MEMBER CHANGE ID:5000cca215c56e02 PDIDI:0145
POOL:0003 INDEX:0000 STATE:0003
000043 2009-02-04 15:03:17:7295810 G=4 S=0 T=1 RP=0 VP=1
LOG_ST_REBUILD_START STATE REBUILD START ID:5000cca215c56e02 PDIDI:0145
POOL:0003 INDEX:0000 FENCE:0000000000000000
  
```

To look at the failed drive (Figure 93), enter command **SHOW UNASSIGNED_POOL FAILED ALL**

Figure 93. Display Failed Drive

```

RAID$ 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: GR80AB0A
Serial Number: GTF200P8GBVPXF
Health State: FAILED
Rotation Speed: 7200 RPM
Device Type: SATA
Member State: UNASSIGNED
Spare: FALSE
Failed: TRUE
UUID: 0x5000cca215c564560
  
```

To monitor the rebuild progress (Figure 94), enter command **SHOW JOB * ALL_ATTRIBUTES**

Figure 94. Display Rebuild Progress

```
RAID$ show job * all_attributes
OID:          0x2b050003
Target:       0x19b40003
Type:         REBUILD
Status:       RUNNING
Priority:      80
Fraction Complete:12%
```

To show the storage pool information (Figure 95), enter command **SHOW POOL <id> ALL_ATTRIBUTES**. Once the rebuild is complete, the status of storage pool will return to “NORMAL”.

Figure 95. Display Storage Pool Information

```
RAID$ 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
.
.
Initializing: FALSE
Rebuilding:   TRUE
Paused:       FALSE
AutoWriteLock: FALSE
Data Lost:    FALSE
Current Home: 0x0015b2a122b20000 0x00000000
Future Home:  0xffffffffffffffff 0x00000000
Preferred Home: 0xffffffffffffffff 0x00000000
BkgdJob OID:  0x2b050003
BkgdJob Priority: 80%
Total Phy Disks 5
State:        NOREDUNDANCY
Member Size:  704512 MB
  PID   State  UUID
  0x0191 RBLD  0x5000cca215c56e02
  0x004e NORM  0x5000cca215c5709c
  0x0040 NORM  0x5000cca215c54c71
  0x0041 NORM  0x5000cca215c5675c
  0x0042 NORM  0x5000cca215c56e55
```

Spare Drive Rebuilding

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.

4.4.4 ***Manual Rebuild***

You may manually replace a failed disk using the **REPLACE** command: **REPLACE PHYSICAL_DISK <id> NEW_DISK <new-disk-id>**

A replace operation is used to replace a failed disk with a healthy spare disk. The operation can take several hours to complete depending on the size of the disk and speed of the replace operation. The rate of rebuild can be adjusted (see Section 3.3, "Performance Management" on page 50 for more information).

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 96).

Figure 96. Controller Information Screen

```
RAID$ 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

Appendices



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" (18cm) × 17.56" (45cm) × 37" (94cm) |
| InfiniteStorage 6120 chassis weight | 120 lbs (54.5kg) without drives, 240 lbs (109.1 kg) with drives |
| Electrical/AC | 200-240V @ 47-63Hz |
| 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, 5973BTU/h |
| Certification | UL, CE, CUL, FCC |

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 50

C

cable connections 20
cache
 mirrored write back, set 48
 write back, set 48
chunk size 42
CLUI, using 35
components, hardware 5
controller 7
 LED 65
cooling module 8
 fault 66, 67
 LED 62, 65
create
 spare pool 30, 45
 storage pool 27
 virtual disk 29
Customer service vi

D

date, set system 48
delete
 storage pool 43
 virtual disk 44
DEM 10
 LED 64
disk
 fail a 42
 failure recovery 70
 fail, clear 42
 information 38
 listing all 38
 location, physical 41
 replace, manual 47
 status, change 42
disk timeout 46
drive failure recovery 70
drive, see disk 38

E

enclosure
 connecting 20
 location, physical 41
ESI processor 66

Ethernet connection 57
expansion enclosure, connect 20

F

failure
 drive, fixing 69, 70
 fan, alerts 66
 fan, fixing 68
 PCM, fixing 68
 power supply, fixing 68
fan, see cooling module 8
features, S2A6620 3
firmware
 current version 52
 update 52
fractional rebuild 46
front panel LED 11
full rebuild 46

G

gateway, set 57
guidelines, rack installation safety 79

H

host
 connections 20
 port location 7

I

installation 15
IP address, set controller 57

J

job priority, set 50

L

LED
 controller 65
 cooling 62, 65
 cooling fault 66, 67
 DEM 64
 fault, system 67
 front panel 11
 power supply 62, 65

- power supply fault 66, 67
- system fault 66
- licensing instructions 26
- login 36
- logout 36

M

- Manual Intervention Required 73
- MIR 73
 - clear, state 73
 - conditions 73
- mirrored write back cache, set 48

N

- netmask, set 57
- network interface, set 57

P

- partial rebuild 46
- password 36
- PCM 8
 - failure recovery 68
- physical location
 - disk 41
 - enclosure 41
 - storage pool 41
 - virtual disk 41
- pool
 - spare, create 30, 45
 - storage, create 27
 - unassigned 38
- power supply module 8
 - fault 66, 67
 - LED 62, 65
- priority
 - job, set 50
 - rebuild, set 51
- Product support vi

R

- rack installation 17
 - safety guidelines 79
- RAID level 42
- rebuild
 - automatic, set 46
 - fractional 46

- full 46
- manual, initialize 47
- partial 46
 - set priority 51
- recovery
 - drive failure 70
 - PCM 68
- replace disk, manual 47
- restart
 - controller 49
 - subsystem 49
- RS-232 terminal set up 35

S

- safety guidelines, rack installation 79
- SATAAssure, set 47
- serial port, set up 35
- shutdown
 - controller 49
 - subsystem 49
- spare pool
 - create 30, 45
 - delete 45
 - global 45
 - name, change 45
 - view list of 41
- sparing policy 46
- specifications, technical 77
- SSH 57
- storage pool
 - configuration, checking 39
 - create 27, 42
 - delete 43
 - initialization 43
 - list of, view 39
 - location, physical 41
 - max number of 42
 - name, change 43
 - SATAAssure, set 47
 - sparing policy 46
 - status 39
- system faults LED 66

T

- technical specifications 77
- Technical support vi
- Telnet 57
- time, set system 48

troubleshooting 61

U

unassigned pool 38
unpack the system 16
user name 36

V

version information 52
virtual disk
 associated with storage pool,
 view 39
 create 29, 44
 delete 44
 location, physical 41
 max number of 42
 name, change 44
 status, checking 40
visual indication
 disk 41
 enclosure 41
 storage pool 41
 virtual disk 41

W

write back cache, set 48

