

SGI® InfiniteStorage 11000 RAID User's Guide

007-5539-001

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 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 Silicon Graphics, Inc., 1140 East Arques Avenue, Sunnyvale, CA 94085-4602.

TRADEMARKS AND ATTRIBUTIONS

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

Microsoft and Windows are registered trademarks of the Microsoft Corporation in the United States and/or other countries.

HyperLink is a registered trademark of Hilgraeve, Inc.

OpenView is a registered U.S. trademark of Hewlett-Packard Company.

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

Record of Revision

Version	Description
-001	March 2008 First publication.

List of Figures

Figure 1-1	Front and Rear Views of the InfiniteStorage 11000 Controller	5
Figure 1-2	Controller Power Supply/Cooling Module	6
Figure 1-3	Controller Fan Module	6
Figure 1-4	I/O Ports On the Rear Panel of the Controller	7
Figure 1-5	LED Status Indicators - Rear Panel of the InfiniteStorage 11000 Controller	8
Figure 2-1	Striping Across Tiers - directRAID	19
Figure 2-2	Tier Map CLI Command Screen	21
Figure 2-3	Tier Changemap CLI Command Screen	22
Figure 2-4	Tier CLI Command Screen	23
Figure 2-5	Tier Configuration CLI Command Screen	24
Figure 2-6	Dual CLI Command Screen	25
Figure 2-7	Dual-Label CLI Command Screen	26
Figure 2-8	Cache CLI Command Screen	27
Figure 2-9	Logical Unit Status - Formatting.	29
Figure 2-10	LUN Status - Ready	29
Figure 2-11	LUN Zoning CLI Screen.	31
Figure 2-12	User Add CLI Command Screen.	31
Figure 3-1	CLI Port on the Controller's Rear Panel	36
Figure 3-2	Password on the CLI Command Screen Example	39
Figure 3-3	The whoami Screen	40
Figure 3-4	Host CLI Command Screen Example	41
Figure 3-5	Host Status CLI Command Screen Example.	43
Figure 3-6	Host Speed CLI Command Screen Example.	44
Figure 3-7	Disk CLI Command Screen Example	45
Figure 3-8	Disk List CLI Command Example Screen	46
Figure 3-9	Disk Status CLI Command Screen Example.	47
Figure 3-10	Disk Defectlist CLI Command Screen	48

Figure 3-11	Disk PLS Tier 1 Status Screen Example	50
Figure 3-12	Tier CLI Command Screen Example	51
Figure 3-13	Tier Configuration CLI Command Screen Example	52
Figure 3-14	Tier Config=ALL CLI Command Screen Example	53
Figure 3-15	LUN CLI Command Screen	54
Figure 3-16	LUN List Command Screen Example	54
Figure 3-17	LUN Config CLI Command Screen Example	56
Figure 3-18	LUN Reservations CLI Command Screen Example	57
Figure 3-19	Displaying the Current Disk Enclosure Failures Using SES	58
Figure 3-20	Route CLI Command Screen Example	63
Figure 3-21	Saving the System Parameters Example Screen	66
Figure 3-22	Save Status CLI Command Screen Example	66
Figure 3-23	LUN CLI Command Screen Example	68
Figure 3-24	LUN Label CLI Command Screen Example	70
Figure 3-25	Moving a LUN's Ownership	71
Figure 3-26	Tier Status and Automatic Disk Rebuilding Parameters Screen	72
Figure 3-27	DUAL CLI Command Screen	78
Figure 3-28	Dual Label CLI Command Screen Example	80
Figure 3-29	Cache CLI Command Screen Example	82
Figure 3-30	AV CLI Command Screen Example	85
Figure 3-31	Initial Cache Example.	87
Figure 3-32	Cache Allocation Example	87
Figure 3-33	Initial Cache Example.	88
Figure 3-34	Cache Allocated After I/O Example	88
Figure 3-35	LUN Lock Status Example Screen	89
Figure 3-36	STATS CLI Command Screen	90
Figure 3-37	Command Delay Statistics Screen Example	91
Figure 3-38	Host Delay Statistics Screen Example	92
Figure 3-39	Tier Delay Statistics Screen Example.	92
Figure 3-40	Host Command Offsets Screen Example	93
Figure 3-41	Dual Message Statistics Screen Example	94
Figure 3-42	Command Length Statistics Screen Example.	94
Figure 3-43	Host Command Offsets Screen Example	95

Figure 3-44	Tier Status Screen Example	96
Figure 3-45	LUN Verify CLI Command Screen Example	97
Figure 3-46	Tier Verify CLI Command Screen Example	99
Figure 3-47	Tier Verify Disable CLI Command Screen Example	100
Figure 3-48	Tier Journal CLI Command Screen	102
Figure 3-49	SES Device Monitoring Rate Example	103
Figure 3-50	Mapping Internal LUNs to External LUNs	104
Figure 3-51	User Login Messages.	105
Figure 3-52	Zoning CLI Command Screen	106
Figure 3-53	Zoning Edit CLI Command Screen	107
Figure 3-54	User CLI Command Screen	108
Figure 3-55	Version CLI Command Screen Example	109
Figure 3-56	Faults CLI Command Screen	112
Figure 3-57	Uptime CLI Command Screen	114
Figure 3-58	ACP UPS SNMP Trap Monitor Status	114
Figure 3-59	API CLI Command Screen	115
Figure 3-60	API Stats CLI Command Screen.	115
Figure 3-61	Console CLI Command Screen	116
Figure 3-62	Changing the Console Baud Rate	116
Figure 3-63	Spare CLI Command Screen	118
Figure 3-64	Telnet Port on the Controller.	120
Figure 3-65	Network CLI Command Screen	121
Figure 3-66	Route CLI Command Screen Example	123
Figure 3-67	Password and Password Defaults CLI Commands Screen	124
Figure 4-1	Power Supply and Cooling Modules on the Controller Front	128
Figure 4-2	Controller Fan Module Removal.	129
Figure 4-3	Example SES Message	132
Figure 4-4	Front Panel Drive Activity Indicators Panel	137
Figure 4-5	Drive Carrier - Handle Released	145
Figure 5-1	The Drive Enclosure System - front open view	147
Figure 5-2	The Drive Enclosure System - rear isometric and rear views	148
Figure 5-3	Drive Enclosure Module Locations	148
Figure 5-4	Power Cooling Module (1)	150

Figure 5-5	I/O Module	152
Figure 5-6	Drive Carrier Module: Closed and Opened	154
Figure 5-7	DEM	155
Figure 5-8	Drive Carrier and DEM Locations	155
Figure 5-9	Front Panel Drive Activity Indicators Panel	156
Figure 6-1	DEM Pair Locations	165
Figure 6-2	1 X 60 Drive Numbering Table.	166
Figure 6-3	2 x 30 Drive Numbering Table	167
Figure 6-4	Full Rack without Controllers.	168
Figure 7-1	I/O Panel LEDs	172
Figure B-1	1 x 60 Drive Layout Configuration	177
Figure B-2	2 x 30 Drive Layout Configuration	178
Figure B-3	Controller-Drive Enclosure Couplet (1x Configuration)	179
Figure B-4	Controller-Drive Enclosure Couplet (2x Configuration)	180
Figure C-1	Controller Cabling (2 Controllers and 5 Drive Enclosures)	182
Figure C-2	Drive Enclosure Cabling (2 Controllers and 5 Drive Enclosures)	183
Figure C-3	Controller Cabling (2 Controllers and 10 Drive Enclosures).	184
Figure C-4	Drive Enclosure Cabling (2 Controllers and 10 Drive Enclosures)	185

Contents

List of Figures	v
1. SGI InfiniteStorage 11000 Introduction	1
Overview	1
Introducing the SGI InfiniteStorage 11000 User's Guide	1
SGI InfiniteStorage 11000 High-Bandwidth RAID Solution	1
The InfiniteStorage 11000 Controller	2
Dual "Coupled" Controllers	2
Features of the InfiniteStorage 11000	2
The InfiniteStorage 11000 System Hardware	4
Power Supply and Fan Modules	5
I/O Connectors and Status LED Indicators	7
Uninterruptible Power Supply (UPS)	10
Host and Disk Connectivity	11
Contacting the SGI Customer Service Center	11
Storage Administration Training from SGI	12
2. InfiniteStorage 11000 Controller Installation	13
Setting Up the InfiniteStorage 11000	13
Unpacking the InfiniteStorage 11000	14
Rack-Mounting the InfiniteStorage 11000 Chassis	14
Connecting the InfiniteStorage 11000 Controllers in Dual Mode	15
Connecting the InfiniteStorage 11000	15
Selecting SAS- ID for Your Drives	15
Laying Out your Storage Drives	15
Connecting the RS-232 Terminal	16
Basic Key Operations	16
Powering On the InfiniteStorage 11000 Controller	17

Configuring the InfiniteStorage 11000	18
Planning Your Setup and Configuration	18
Configuration Interfaces	20
Login as Administrator	20
Setting System Time & Date	21
Setting Tier Mapping Mode	21
Checking Tier Status and Configuration	22
Tier Configuration Heading Definitions	24
Cache Coherency and Labeling in Dual Mode	25
Configuring the Storage Arrays	26
Setting Security Levels	30
User Authentication (Recommended for SAN Environment)	30
To add a user:	31
Host Port Zoning (Anonymous Access)	33
3. Storage System Management	35
Managing the InfiniteStorage 11000	35
Management Interface	35
Locally - Serial Interface	35
Remotely - Telnet	37
Available Commands	37
Administrator and User Logins	37
Login	37
Logout	38
Password	38
Who Am I	40
Configuration Management	40
Configure and Monitor Status of Host Ports	41

Configure and Monitor Status of Storage Assets	45
Disk and Channel Information	45
Tier View	50
Tier Configuration	52
LUN View	53
LUN Configuration	56
LUN Reservations	57
Adding/Removing Storage Assets	57
Status of Drive Enclosures	58
Display SES Devices Information	58
Visual Indication of Drive	58
Visual Indication of Tier	59
Visual Indication of Channel	59
Tier Mapping for Enclosures	59
System Network Configuration.	59
Telnet	60
SNMP & Syslog	61
API Server Connections	62
Displaying and Editing the Routing Table	62
Restart and Shutdown the InfiniteStorage 11000	64
System Restart	64
System Shutdown	64
Setting the System's Date and Time	65
System Date	65
System Time	65
Saving the InfiniteStorage 11000's Configuration	66
Restoring the System's Default Configuration	67
LUN Management	67
Configuring the Storage Array.	67
Creating a LUN.	68
Formatting a LUN	69
Interrupting a LUN Format Operation.	69
Changing a LUN Label	70

Moving a LUN (Dual Mode Only) 71
Deleting a LUN 71
SCSI Reservations 71
Automatic Drive Rebuild 71
Manual Drive Rebuild 72
Drive Rebuild Verify 72
Manual Drive Replace 72
Interrupting a Rebuild Operation 73
SMART Disk Monitoring and Test 73
SMART Support 73
Monitoring SMART Support 73
SMART Enable and Disable 74
Status and Show 74
SMART Self-test Support 74
Start Self-test 75
Stop Self-Test 75
Self-test Status 75
SMART and Proactive Sparing. 75
SMART and Drive Sparing 76
8+2 Channel Sparing 76
8+1+1 Sparing 77
Couplet Controller Configuration (Cache/Non-Cache Coherent) 77
Cache Coherent 78
Non-Cache Coherent 78
Fail / Restore the Other Controller Unit in the Couplet Pair 79
Labeling the InfiniteStorage 11000 Controller Unit(s) 80
Singlet 81

Performance Management	82
Optimizing I/O Request Patterns	82
Display Current Cache Settings	82
Cache Segment Size	83
Writeback Cache Settings	83
Prefetch Settings	83
Cache Settings Reset	84
Disk Configuration Settings	84
Audio/Visual Settings of the System	85
Locking LUN in Cache	87
Locking / Unlocking a LUN	89
System Performance Statistics	89
Resources Allocation	95
Background Format/Rebuild Operations	95
Background LUN Verify Operations	97
Background TIER Verify Operations	98
Rebuild Journaling	100
SES Device Monitoring Rate	103
Host Command Timeout	103
Security Administration	104
Monitoring User Logins	105
Zoning (Anonymous Access)	105
User Authentication.	108
Firmware Update Management	109
Displaying Current Firmware Version	109
Firmware Update Procedure	109
Remote Login Management	110
When a Telnet Session is Active	111
System Logs	111
Message Log	111
System and Drive Enclosure Faults	112
Displaying the System's Uptime	114
Saving a Comment to the Log	114

Other Utilities	114
APC UPS SNMP Trap Monitor	114
API Server Connections	115
Changing Baud Rate for the CLI Interface	116
CLI/Telnet Session Control Settings	116
Disk Diagnostics.	117
Disk Reassignment and Miscellaneous Disk Commands	117
SPARE Commands	117
Remote Management of the InfiniteStorage 11000 controller	120
Network Connection.	120
Network Interface Set Up	120
Login Names and Passwords	123
SNMP Set Up on Host Computer	125
InfiniteStorage 11000 Controller Implementation of SNMP	126
InfiniteStorage 11000 Controller MIB (The MIB is read only.)	126
Traps	126
4. Troubleshooting the SGI InfiniteStorage 11000 Controller and Drive Enclosures	127
Component Failure Recovery	127
Controller Power Supply Failure	127
Controller Fan Failure	129
Recovering from Drive Failures.	129
Single Drive Failures	130
Returning the system to a fault-tolerant state	130
Manually Replacing a Failed Disk with a Spare Disk	131
Changing the Rate of Rebuild	131
Interrupting the Rebuild Operation	131
Multiple Drive and Channel Failures	131
Component Failure on Enclosures	132

Drive Enclosure Troubleshooting Introduction133
Initial Start-up Problems133
Faulty Cords133
Alarm Sounds On Power Up133
Green “Signal Good” LED on I/O Module Not Lit133
Computer Doesn’t Recognize the Drive Enclosure Subsystem133
Status LEDs134
HDD (Hard Disk Drive)134
PCM (Power Cooling Module) Status LEDs134
DEM (Drive Expander Module)135
I/O Module136
Front Panel Drive Activity Indicators137
Audible Alarm138
Top Cover Open138
SES Command138
General Drive Enclosure Troubleshooting139
Thermal Control139
Thermal Alarm140
Thermal Shutdown140
Dealing with Drive Enclosure Hardware Faults.141
Continuous Operation During Replacement141
Replacing a Module141
Power Cooling Modules.142
To remove an AC PCM/Cooling Module- Handle Model:.142
To remove an AC PCM - Thumb Screw Model:143
To install an AC PCM - Handle Model:143
To install an AC PCM - Thumbscrew Model:144
I/O Module144
Replacing the Drive Carrier Module145
Replacing the DEM146
5. SGI InfiniteStorage 11000 Drive Enclosure Overview147
The Drive Enclosure.147
Enclosure Core Product149

Enclosure Chassis	149
The Plug-in Modules	149
Power Cooling Module (PCM)	150
Input/Output (I/O) Module	152
Drive Carrier Module and Status Indicator	154
DEM Card	154
Indicators	156
Front Panel Drive Activity Indicators	156
Internal Indicators	158
Rear of Enclosure Activity Indicators	158
Visible and Audible Alarms	159
Drive Enclosure Technical Specifications	160
Dimensions	160
Weight	160
AC INPUT PCM	161
DC INPUT PCM	161
DC OUTPUT PCM	162
PCM Safety and EMC Compliance	162
Power Cord	162
Environmental Compliance	162
6. Drive Enclosure Installation	163
Introduction	163
Planning the Installation	163
Enclosure Bay Numbering Conventions	165
Enclosure Installation Procedures	167
I/O Module Configurations	169
Controller Options	169
SAS DEM	169
SATA Interposer Features	169
Grounding Checks	169
7. Drive Enclosure Operation Overview	171
Before You Begin	171

	Power On / Power Down171
	PCM LEDs172
	I/O Panel LEDs172
A.	Technical Specifications175
B.	Drive Addressing Information177
C.	Drive Cabling Configuration Examples181
D.	Safety Guidelines for Rack Installation187
	Elevated Operating Ambient Temperature187
	Air Flow187
	Mechanical Loading187
	Circuit Overloading187
	Reliable Earthing188

SGI InfiniteStorage 11000 Introduction

Overview

Your SGI® InfiniteStorage 11000 RAID controller storage system is an intelligent storage infrastructure device designed and optimized for the high bandwidth and capacity requirements of IT departments, rich media, and high performance workgroup applications.

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

The controller's design is based on an advanced pipelined, parallel processing architecture, caching, RAID, and system and file management technologies. These technologies have been integrated into a single plug-and-play device—the SGI InfiniteStorage 11000 RAID controller—providing simple, centralized, and secure data and SMNP management.

The SGI InfiniteStorage 11000 is designed specifically to support medium-high to very-high bandwidth, rich content, and shared access to and backup of large banks of data. 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.

Introducing the SGI InfiniteStorage 11000 User's Guide

Welcome to the SGI InfiniteStorage 11000 User's Guide; this document covers features relevant to installation, operation, service and troubleshooting of this high-performance storage system.

SGI InfiniteStorage 11000 High-Bandwidth RAID Solution

The SGI (InfiniteStorage 11000 RAID system) is an intelligent network infrastructure device designed and optimized for the high bandwidth and capacity requirements of departments, rich media and high performance workgroup applications.

The InfiniteStorage 11000 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 11000 is designed based on advanced pipelined, parallel processing architecture, caching, directRAID, ASIC (Application Specific Integrated Circuits), system and file management technologies. These technologies are integrated into a single plug and play RAID control device—the InfiniteStorage 11000—providing simple, centralized and secure data and network management.

The InfiniteStorage 11000 Controller

The InfiniteStorage 11000 is designed specifically to support high bandwidth, shared access to and backup of large banks of data, and rich content. 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.

Using zero-latency High Parallel Port Technology (HPPT™), each InfiniteStorage 11000 orchestrates a coherent flow of data throughout the SAN, from users to storage, managing data at up to 1600MB/second (400MB/s per FC-AL port). This task is accomplished through virtualized host and storage connections, a DMA-speed shared data access space, advanced network-optimized directRAID™ data protection and security—all acting in harmony with sophisticated Fibre Channel storage management intelligence embedded within the InfiniteStorage 11000.

Dual “Coupled” Controllers

With its modular design, the InfiniteStorage 11000 controllers can be “coupled” together forming data access redundancy while maintaining fully pipelined, parallel bandwidth to the same disk storage (see “Moving a LUN (Dual Mode Only)” on page 71 and “Couplet Controller Configuration (Cache/Non-Cache Coherent)” on page 77). This modular architecture ensures high data availability and uptime along with application performance. The system provides full bandwidth to all host ports simultaneously, without host striping.

Features of the InfiniteStorage 11000

The InfiniteStorage 11000 incorporates the following features:

Simplifies Deployment of Complex SANs

The InfiniteStorage 11000 provides SAN administration with the management tools required for large number of clients and complex Fibre Channel configuration through its topology-independent FC-user management features.

Fibre Channel Connectivity with up to 1600 MB/s Throughput

The InfiniteStorage 11000 provides up to 4 individual 400MB/s Fibre Channel host port connections, including simultaneous access to the same data through multiple ports. Each FC host port supports Class 3, point-to-point, FC-AL, and switched fabric operation.

High-Parallel Port Technology Implementation

HPPT architecture provides system performance and fault-tolerance.

- 10-independent dual-ported FC drive channels; including separate (asynchronous, “self-optimized”) I/O queues for each channel
- 10-independent directRAID engines
- 10-independent RAID-protected high speed data caches

Highly Scalable Performance and Capacity

The controller provides both fault-tolerance and capacity scalability. Performance will remain the same even in degraded mode. Internal data striping provides generic load balancing across drives. Formattable capacity can be scaled from hundreds of GBs to hundreds of TBs depending on drive type and optional enclosure type chosen. See your SGI sales or service representative for more information.

Comprehensive, Centralized Management Capability

The InfiniteStorage 11000 RAID controller provides a wide range of management capabilities: Configuration Management, Performance Management, LUN Management, Security Administration, and Firmware Update Management.

Management Options via RS-232 and Ethernet (Telnet)

A RS-232 port and Ethernet port are included to provide local and remote management capabilities. SNMP and GUI are also supported.

Data Security with Dual-Level Protection

Non-host based data security is maintained with scalable security features including restricted management access, dual-level protection and authentication against authorized listing (up to 512 direct host logins are supported). No security software is required on the host computers.

Storage Virtualization and Pooling

Storage pooling enables different types of storage to be aggregated into a single logical storage resource from which virtual volumes can be served up to multi-vendor host computers. Up to 128 LUNs are supported. Each LUN can be subdivided into up to 64 smaller equally-sized LUN segments, giving a total of 8192 LUN segments. The LUN segments are managed together and share the same characteristics.

SES (SCSI Enclosure Services) Support for Enclosure Monitoring

Status information on the condition of enclosure, disk drives, power supplies, and cooling systems are obtained via the SES interface.

Absolute Data Integrity and Availability

Automatic drive failure recovery procedures are transparent to users. Up to **125 hot spares** are supported.

Hot-Swappable and Redundant Components

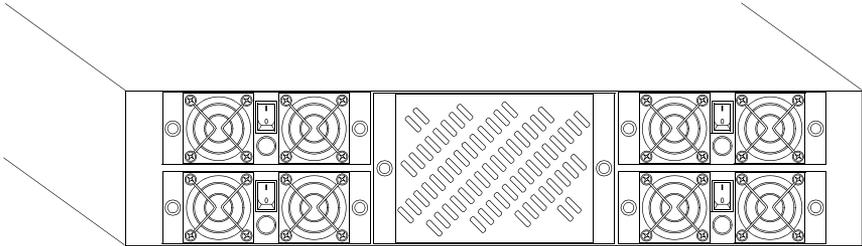
The InfiniteStorage 11000 utilizes redundant, hot-swappable power supplies and cooling modules that can be replaced while the system is running.

The InfiniteStorage 11000 System Hardware

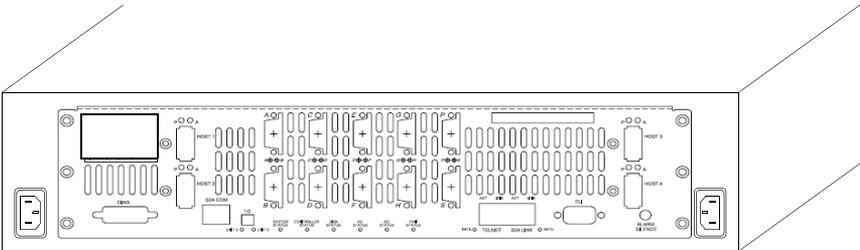
The InfiniteStorage 11000 is a high-performance controller designed to be rack-mounted in standard 19 inch racks. Each controller is 3.5 inches in height, requiring 2U of rack space. The system uses 10 independent SAS drive channels to manage data distribution and storage for up to 120 disk drives per channel (note that this can be limited by drive enclosure type).

The basic InfiniteStorage 11000 (see Figure 1-1 on page 5) controller unit consists of:

- A single unit (with a minimum of 2.56GB cache memory)
- Ten SAS connectors that connect the InfiniteStorage 11000 to drive enclosures
- Connector(s) for host Fibre Channel (FCS) connection(s)
- Serial connectors for maintenance/diagnostics
- Ethernet RJ-45 connector.



▲ Front of InfiniteStorage 11000 Behind Cover Panel



▲ Rear of InfiniteStorage 11000 Controller

Figure 1-1 Front and Rear Views of the InfiniteStorage 11000 Controller

Power Supply and Fan Modules

Each InfiniteStorage 11000 is equipped with four (4) Power Supply/Cooling modules and one (1) fan module (Figure 1-3). The PSU (power supply unit) voltage operating ranges are nominally 110V or 230V AC, selected automatically.

The four Power Supply/Cooling modules provide redundant power supply and cooling system for the unit. If one module fails, the other will maintain the power supply and cooling while you replace the faulty module. The faulty module will still be providing proper air flow for the system so do not remove it until a new module is available for replacement.

The four modules are installed in the left and right slots at the front of the unit, behind the cover panel (see Figure 1-1). Each PSU module is held in place by two thumbscrews.

The LED mounted on the front of the Power Supply/Cooling module indicates the status of the PSU and the fans (Figure 1-2). It is green when the module is operating normally and turns red when a fault occurs. (Note: See Status LEDs on rear panel for clarification.)

The fan module (Figure 1-3) is installed in the front center slot, behind the cover panel, and is held in place by two thumbscrews.

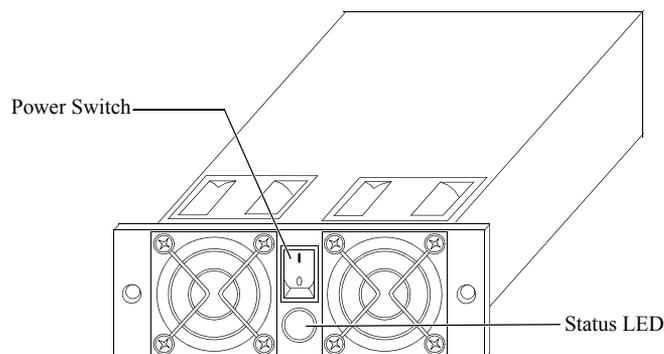


Figure 1-2 Controller Power Supply/Cooling Module

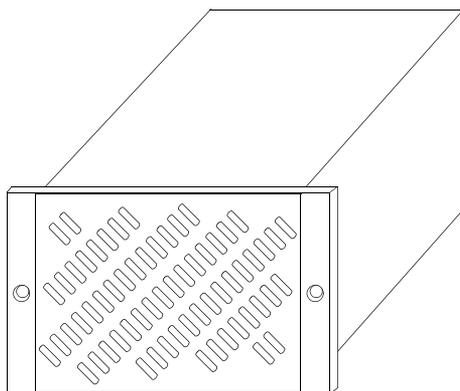


Figure 1-3 Controller Fan Module

I/O Connectors and Status LED Indicators

Figure 1-4 shows the ports at the back of the InfiniteStorage 11000 controller unit.

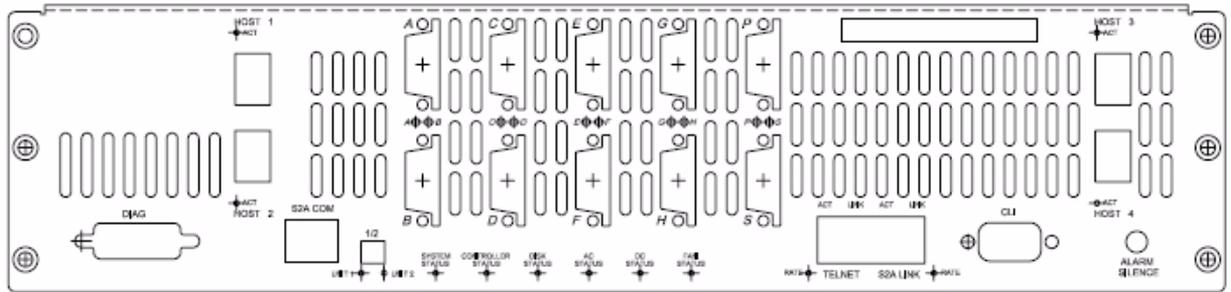


Figure 1-4 I/O Ports On the Rear Panel of the Controller

The four **HOST** ports (latch-style connectors) are used for FC host connections. You can connect your host systems directly to these ports or connect the ports to optional switches and hubs.

The **FC LEDs** are located next to each FC host port. There are 3 LEDs per host port. These LEDs indicate whether the connection is running at 8 GB (left LED), 4 GB (middle LED), or 2 GB (right LED). The respective LED will be a solid **green** to indicate that there is a physical connection. If the LED is flashing, this indicates data transfer. If the connector is taken from the host port, all 3 LEDs for that port will flash.

The **DISK CHANNEL** ports (jackscrew style connectors) are for disk connections. There are 10 ports, labeled by data channels (**ABCDEFGHIPS**). **Flashing** LEDs indicate activity.

The **RS-232** connector provides local system monitoring and configuration capabilities and uses a standard DB-9 null modem female-to-male cable.

The **TELNET** port provides remote monitoring and configuration capabilities. The **ACT** (Activity) LED flashes **green** when there is Ethernet activity. The **LINK** LED turns **green** when the link speed is **1000MB/s**, **amber** when the link speed is **100MB/s**, and is **unlit** when the link speed is **10 MB/s**. It is **unlit** when there is no Ethernet link.

The **LINK** port is used to connect individual InfiniteStorage 11000 controller units to form a couplet via a cross-over Ethernet cable. The **ACT** (Activity) LED flashes **green** when there is Ethernet activity. The **LINK** LED turns **green** when the link speed is **1000MB/s**, **amber** when the

link speed is **100MB/s**, and is **unlit** when the link speed is **10 MB/s**. The LED is **unlit** when there is no Ethernet link.

The **COM** port is an RS-232 Interface that uses an RJ-45 cable and connects controller units. The **COM** port has two(2) LEDs associated with it: **HDD ACT** (Activity) and **LINK ACT**.

The **Controller ID Selection Switch (labeled as 1/2)** allows the user to configure the units in a couplet as **Unit 1** or **Unit 2**. Each unit has an **activity** LED, which is **green** for the selected unit and is comprised of two DIP switches. The first DIP switch (indicated by the 1/2 label) is used to select the unit configuration. Flip the switch up for Unit 1--down for Unit 2.

When two InfiniteStorage 11000 controllers are paired together to form a couplet, one controller must be configured as **Unit 1** and its partner must be configured as **Unit 2**. If you are only configuring a single controller, this unit is usually labelled **Unit 1**.

There are two **AC Fail** LEDs. Each LED is connected to its power supply independent of the other supply. The LEDs are **green** to indicate that the AC input to the supply is present. The LEDs turn **red** if the AC input to the supply is not present. If this occurs, check the LEDs on the front side of the unit. If you lose AC power from one supply outlet, the LED for that outlet will turn **red**. Figure 1-5 shows the following status LEDs: System, Controller, Disk, AC, DC, and Fan.



Figure 1-5 LED Status Indicators - Rear Panel of the InfiniteStorage 11000 Controller

The **SYSTEM STATUS** LED is solid **green** when the storage system is operating normally.

The **CTRL (CONTROLLER) STATUS** LED is **green** when the InfiniteStorage 11000 is operating normally and turns **red** when the InfiniteStorage11000 controller unit has failed.

The **DISK STATUS** LED is **green** when a disk enclosure is operating normally and turns **amber** when there is a problem.

The **AC STATUS** LED is green indicating normal operating status (if all power supplies indicate Power Good).

If two (2) supplies from the same AC cord both indicate power bad, the AC has failed and LED glows red. (Note: Check LEDs on front of unit.)

The **DC STATUS** LED is **green** when indicating normal operating status. It turns **amber** if there is a non-critical power Supply DC fault (that is, a power supply is not installed or is not indicating “**Power Good**”). It turns **red** if an on-board supply fails or if there is a critical supply fault. If this occurs, check the LEDs on the front side of the unit.

The **FAN STATUS** LED is **solid green** when fans are operating normally. A **flashing green** LED indicates system monitoring activity is being updated. The LED flashes **amber** if one of the fans in the module fails. If 2 or more fans fail, the LED flashes a solid **red** and the system will begin the shut down process at 5 seconds, for a total of 30 seconds to complete shutdown.

Table 1-1 Controller LED Indicators and Status Activity

Status Indicator	Led Activity	Explanation
DISK ports	Flashing Green	Activity. There is an LED for each of the ten ports/channels (ABCDEFGPS)
	Unlit	
Telnet ACT	Flashing Green	Activity
	Unlit	No activity
Telnet LINK (Speed)	Solid Green	Link Speed=1000 mb/s
	Solid Amber	Link Speed=100 mb/s
	Unlit	Link Speed= 0 mb/s
Link ACT	Flashing Green	Activity
	Unlit	No activity
Link LINK (Speed)	Solid Green	Link Speed=1000 mb/s
	Solid Amber	Link Speed=100 mb/s
	Unlit	Link Speed=10 mb/s
System	Solid Green	System is operating normally
Ctrl	Solid Green	System is operating normally
	Solid Amber	System is shutting down

Table 1-1 (continued) Controller LED Indicators and Status Activity

Status Indicator	Led Activity	Explanation
Disk	Solid Green	All related disk enclosures are operating normally
	Solid Amber	There is a problem with 1 or more of the disk enclosures
Temp Status	Solid Green	All temp sensors operating normally
	Solid Amber	At least 1 temp sensor has reported over-temperature conditions
	Solid Red	2 or more temp sensors has reported over-temperature condition
DC	Solid Green	Operating normally
	Solid Amber	Non-critical power supply fault
	Solid Red	Critical power supply fault
Fan Status	Solid Green	Operating normally
	Flashing Green	System monitoring activity
	Flashing Amber	1 fan has failed and needs to be replaced
	Solid Red	2 or more fans have failed or are undetected and the system will shutdown in 30 seconds
AC Fail	Solid Green	Operating normally
	Solid Red	Power input to supply not present. AC failure
FC	Solid	Physical connection has been made
	Flashing	Data is being transferred

Uninterruptible Power Supply (UPS)

Using an Uninterruptible Power Supply (UPS) with the InfiniteStorage 11000 is highly recommended. The UPS can guarantee power to the system in the event of a power failure for a short time, which will allow the system to power down properly.

There are two types of UPS: *basic* and *redundant*. The basic UPS is rack-mountable, and can maintain power to a five (5) enclosure system for seven (7) minutes while the system safely shuts down during a power failure. The redundant UPS contains power cells that provide a redundant UPS solution.

Note: The UPS should be installed by a licensed electrician. Contact your service representative to obtain circuit and power requirements.

Host and Disk Connectivity

The InfiniteStorage 11000 includes four (4) 4Gb/s full duplex Fibre Channel host ports. The host ports are hard-wired for non-OFC optical connections utilizing SFP (Small Form-factor Pluggable) connectors.

Contacting the SGI Customer Service Center

To contact the SGI Customer Service Center, call 1-800-800-4SGI, or visit:
<http://www.sgi.com/support/customerservice.html>
From outside the United States contact your local SGI sales office.

To reach SGI for other purposes, use the following contact information:

SGI Corporate Office

1140 E. Arques Avenue

Sunnyvale, CA 94085

<http://www.sgi.com>

North America +1 800.800.7441

Latin America +55 11.5185.2860

Europe +44 118.912.7500

Japan +81 3.5488.1811

Asia Pacific +1 650.933.3000

Storage Administration Training from SGI

SGI offers customer training classes covering all current compute and storage systems. If you have a maintenance agreement in place with SGI, contact SGI Customer Education at 1-800-361-2621 for information on the time, location and cost of the applicable training course you are interested in. Or, go to the following URL site for more education information:
<http://www.sgi.com/support/custeducation/>

Customers with support contracts can also obtain information from:
<https://support.sgi.com/login>

InfiniteStorage 11000 Controller Installation

This chapter provides an overview of the controller installation process. The steps are explained in detail in the following sections of this chapter.

1. Unpack the controller system.
2. If it is necessary to install the controller in the 19-inch cabinet(s), contact your service provider as necessary.

Note: Most controller configurations arrive at sites pre-mounted in a rack supplied by SGI. The units are heavy and should only be installed by trained personnel with proper tools.

3. Set up and connect the drive enclosures to the controller.
4. Connect the controller to your Fibre Channel (FC) switch and host computer(s).
5. Connect your RS-232 terminal to the controller.
6. Power up the system.
7. Configure the storage array (create and format LUNs - Logical Units) via RS-232 interface, Telnet, or GUI.
8. Define and provide access rights for the clients in your SAN environment. Shared LUNs need to be managed by SAN management software. Individual dedicated LUNs appear to the client as local storage and do not require management software.
9. Initialize the system LUNs for use with your server/client systems. Partition disk space and create file systems as needed.

Setting Up the InfiniteStorage 11000

This section provides information on installing the hardware components of your InfiniteStorage 11000 controller system. Refer to the Rackmounting Installation Guide.



Warning: If your configuration is 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 (ramp).

Unpacking the InfiniteStorage 11000

Before you unpack your InfiniteStorage 11000, inspect the shipping container for damage. If you detect damage, report it to your carrier immediately. Retain all cartons and packing materials in case you need to store or ship the system in the future.

Visually inspect the InfiniteStorage 11000 chassis and all components for signs of external damage. If you detect any problems, contact your shipper and SGI to resolve.

Your InfiniteStorage 11000 shipping carton contains the following items:

- InfiniteStorage 11000 chassis with:
- four (4) power supply/cooling modules
- one (1) fan module
- two (2) power cords
- RS-232 and Ethernet cables for monitoring and configuration
- cover panel and rack-mounting hardware
- Optional CDs and documentation

Caution: Electrostatic discharge can damage the circuit boards. Be certain to wear an ESD wrist strap or otherwise ground yourself when handling the modules and components.

Rack-Mounting the InfiniteStorage 11000 Chassis

Refer to the Rack-Mounting Installation Guide.

Note: The final step is to attach the cover panel to the front of chassis using two thumb screws. However, since you need access to power switch to turn it on, do not attach the cover panel until your unit is up and running.

Connecting the InfiniteStorage 11000 Controllers in Dual Mode

For dual mode configuration only:

1. Connect the “LINK” ports on the two InfiniteStorage 11000 units using an Ethernet cable.
2. Connect the “COM” ports on the two units using an RJ45 cable.

Connecting the InfiniteStorage 11000

1. There are ten disk channels on the controller. The disk ports are labeled as follows:
DISK A =Channel A DISK B=Channel B
DISK C =Channel C DISK D=Channel D
DISK E=Channel E DISK F=Channel F
DISK G=Channel G DISK H=Channel H
DISK P=Channel P (parity)
DISK S=Channel S (spare)
2. Using the ten SAS cables provided, connect these disk ports to your ten disk channels. Make sure the latches on the transceivers are engaged.
3. Each InfiniteStorage 11000 controller supports up to four (4) Fibre Channel (FC) host connections. You may connect more than four client systems to the InfiniteStorage 11000 with the use of hubs/switches and you can restrict user access to the LUNs.
4. The Host ports (SFP) are numbered 1 through 4. Connect your host system(s), hubs or switches to these ports. Make sure the latches on the transceivers are engaged.

Selecting SAS- ID for Your Drives

The InfiniteStorage 11000 uses a select ID of 1.

Laying Out your Storage Drives

Tiers, or RAID groups, are the basic building blocks of the InfiniteStorage 11000. A tier can be catalogued as 8+1 or 8+2. In 8+1 mode, a tier contains 10 drives—eight (8) data drives

(Channels A through H), one parity drive (Channel P), and one *optional* spare drive (Channel S). In 8+2 mode, a tier contains 10 drives—eight (8) data drives (Channels A through H) and two parity drives (Channel P and S).

The InfiniteStorage 11000 can manage up to 128 tiers.

Configuration of disks in the enclosures must be in sets of complete tiers (Channels A through P). Allocating one spare drive per tier gives you the best data protection but this is not required. The spare drives on the InfiniteStorage 11000 are global hot spares.

Connecting the RS-232 Terminal

For first time set-up, you will need access to an RS-232 terminal or terminal emulator (such as Windows hyperterminal). Then you may set up the remote management functions and configure/monitor the InfiniteStorage 11000 remotely via Telnet.

1. Connect your terminal to the CLI port at the back of the InfiniteStorage 11000 using a standard DB-9 female-to-male null modem cable.
2. Open the terminal window.
3. Use the settings for your serial port as listed in Table 2-1:

Table 2-1 Serial Port Settings

Setting	Value
Bits per second:	115,200
Data bits:	8
Parity:	None
Stop bits:	1
Flow Control:	None

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 telnet sessions when entering new commands. Basic Key Assignments are listed in Table 2-2 on page 17.

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

- CLI prompts the user for more information.

- Alternate CLI prompt is active. (The alternate CLI is used on the RS-232 connection during an active telnet session.)

Note: Not all telnet programs support all the keys listed in the table. The Backspace key in the terminal program should be setup to send 'Ctrl-H'.

Table 2-2 Basic Key Assignments

Key	Escape Sequence	Description
Backspace	Ctrl-H, 0x08	deletes preceding character
Del	Del, 0x7F 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~	moves cursor to the start of the line.
End	Esc [K Esc [4~	moves cursor to the end of the line
Ins	Esc [2~	toggles character insert mode, on and off
NOTE: Insert mode is ON by default and resets to ON for each new command.		
PgUp	Esc [5~	retrieves oldest command in the history buffer
PgDn	Esc [6~	retrieves latest command in the history buffer

Powering On the InfiniteStorage 11000 Controller

Note: Systems that have dual InfiniteStorage 11000 controllers (couplets) should have the controllers powered on simultaneously to insure correct system configuration.

1. Verify that the power switches on the two (2) power supply module at the back of each controller are off.
2. Connect the two AC connectors, using the power cords provided at the back to the AC power source for each controller unit. For maximum redundancy, connect the two power connectors to two different AC power circuits for each unit.

3. Check that all your drive enclosures are powered up.
4. Check that the drives are spun up and ready.
5. Turn on the power supplies on the InfiniteStorage 11000 unit(s). The controller will undergo a series of system diagnostics and the bootup sequence is displayed on your terminal.
6. Wait until the bootup sequence is complete and the InfiniteStorage 11000 system prompt is displayed.

Note: Do not interrupt the boot sequence without guidance from SGI Technical Support.

The system is now configurable, reference the following sections as needed for example configuration information.

Configuring the InfiniteStorage 11000

This section provides information on configuring your InfiniteStorage 11000.

Note: The configuration examples provided here represent only a general guideline. These examples should not be used directly to configure your particular InfiniteStorage 11000.

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

Planning Your Setup and Configuration

Before proceeding with your InfiniteStorage 11000 configuration, determine the requirements for your SAN environment, including the types of I/O access (random or sequential), the number of storage arrays (LUNs) and their sizes, and user access rights.

The InfiniteStorage 11000 uses either an 8+2 or an 8+1+1 parity scheme called **directRAID**. It is a unique implementation that combines the virtues of RAID 3, RAID 0, and RAID 6 (see Figure 2-1). Like RAID 3, a dedicated parity drive is used per 8+1 parity group; two parity drives are dedicated in the case of an 8+2 parity group or RAID 6. A parity group is also known as a **Tier**.

The directRAID parity scheme exhibits RAID 3 characteristics such as tremendous large block-transfer—READ and WRITE—capability with NO performance degradation in crippled mode. This capability also extends to RAID 6, delivering data protection against a double disk drive failure in the same tier with no loss of performance.

		Tier Configuration			
		Capacity	Space Available	Disk Status	Lun List
		(Mbytes)	(Mbytes)		
Striping across tiers when a LUN is created across multiple tiers	Tier				
	1	280012	271820	ABCDEFGHPS	0
	2	280012	271820	ABCDEFGHPS	0
	3	280012	271820	ABCDEFGHPS	0

----->
Parity Protection within same tier

Figure 2-1 Striping Across Tiers - directRAID

However, Like RAID 5, directRAID does not lock drive spindles and does allow the disks to re-order commands to minimize seek latency, and the RAID 0-like functionality allows multiple tiers to be striped, providing “PowerLUNs” that can span 100’s of disk drives. These PowerLUNs support very high throughput and have a greatly enhanced ability to handle small I/O (particularly as disk spindles are added) and many streams of real-time content.

LUNs can be created on just a part of a tier, a full tier, across a fraction of multiple tiers, or across multiple full tiers. A minimum configuration for tiers of drives require either 9 drives in an 8+1 configuration or 10 drives in an 8+2 configuration. When configured in 8+1+1 mode, the tenth data segment is reserved for global hot spare drives. When configured in 8+2 mode, spares may reside on each data segment and are global only to that data segment.

The InfiniteStorage 11000 supports various disk drive enclosures, such as the StorageScaler 6000, that can be used to populate the 10 <ABCDEFGHPS> disk channels in both SAS 1x and SAS 2x modes. Each chassis has a limit to the tiers that can be created and supported. Refer to the specific disk enclosure user guides for further information.

You can create up to 1024 LUNs in an InfiniteStorage 11000. LUNs can be shared or dedicated to individual users, according to your security level setup, with Read or Read/Write privileges granted per user. Users only have access to their own and “allowed-to-share” LUNs. Shared LUNs need to be managed by SAN management software. Individual dedicated LUNs appear to users as local storage and do not require external management software.

Note: In dual mode, LUNs are “owned” by the unit on which they are created. Hosts only see the LUNs on the InfiniteStorage 11000 to which they are connected, unless cache coherency is enabled.

For random I/O applications, use as many tiers as possible and create one or more LUNs. For applications that employ sequential I/O, use individual or small grouping of tiers. If you need guidance in determining your requirements, contact **your SGI or other responsible Technical Support** organization.

Configuration Interfaces

There are two interfaces you can use to configure the InfiniteStorage 11000 system:

- Command Line Interface (CLI)
- Graphical user interface (GUI) Management Tool, a JAVA-based GUI RAID client-server application manager.

Note: This user guide provides information for set up using the CLI. If you want to configure the InfiniteStorage 11000 using the Management Tool, refer to the *Management Tool User Guide*. Configuration examples are given for both options.

Login as Administrator

The default Administrator account name is **admin** and its default password is **password**. (See Section 3.1.3 "Administrator and User Logins" for information on how to change the user and administrator passwords.) Only users with administrator rights are allowed to change the controller's configuration.

To login:

1. At the login prompt, type: **login admin**<Enter>
2. At the password prompt, type: **password**<Enter>

Setting System Time & Date

The system time and date for the InfiniteStorage 11000 controllers 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 system date and time so that the time stamps for all events are correct. In dual mode, changes should always be made on Unit 1. New settings are automatically applied to both units.

- To set the system date, at the prompt, type: **date mm dd yyyy** <Enter> where **mm** represents the two digit value for month, **dd** represents the two digit value for day, and **yyyy** represents the four digit value for year. For example, to change the system date to **March 1, 2009**, enter: **date 3 1 2009** <Enter>
- To set the system time, at the prompt, type: **time hh:mm:ss** <Enter> where **hh** represents the two digit value for hour (00 to 24), **mm** is the two digit value for minutes, and **ss** represents the two digit value for seconds. For example, to change the system time to **2:15:32 p.m.**, enter: **time 14:15:32** <Enter>

Note: 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.

Setting Tier Mapping Mode

When the InfiniteStorage 11000 system is first configured, it is necessary to select a tier mapping mode for the attached enclosures.

The InfiniteStorage 11000 currently supports SAS drive enclosures.

- To display the current mapping mode (see Figure 2-2 on page 21), type: **tier map** <Enter>
- To change the mapping mode (see Figure 2-3 on page 22):

```
11000[1]: tier map

Current mapping mode: 15: 'SAS_2X      Enclosure'

11000[1]:
```

Figure 2-2 Tier Map CLI Command Screen

1. Type **tier changemap**_{<Enter>}

```
11000[2]: tier changemap
Supported mapping modes:
-----

15) SAS_2X   Enclosure mapping.
16) SAS_1X   Enclosure mapping.

Current mapping mode: 15, SAS_2X   Enclosure.

Enter the new mapping mode (15..16), 'e' to escape:
```

Figure 2-3 Tier Changemap CLI Command Screen

2. Enter the appropriate mapping mode.
 - For 10 and 20 box solutions, choose SAS_2X.
 - For 5 box solutions, choose SAS_1X.
3. For the changes to take effect, enter: **restart**_{<Enter>}

Checking Tier Status and Configuration

Use the **tier** command to display your current tier status. Figure 2-4 on page 23 illustrates the status of a system containing 80 drives on 8 tiers with both parity modes of 8+1 and 8+2 tiers. The plus sign (+) adjacent to the tier number indicates that the tier is in 8+2 mode.

```

11000[2]: tier

                                Tier Status

                                Space
                                Capacity Available
                                (Mbytes) (Mbytes)
Tier Owner                      Disk
                                Status
-----
  1 + 1    1121384    1111384  ABCDEFGHPS  0
  2 + 1    1121384    1101384  ABCDEFGHPS  1 6
  3 + 1    1121384    1110384  ABCDEFGHP.  2 4
  4 + 1    1121384    1103384  ABCDEFGHPS  3 5 7 8 9 10 11 12 13

Disk 3S Failed.

Automatic disk rebuilding:  Enabled
Automatic SMART replacement: Disabled
Maximum active rebuilds:    4
Maximum active verifies:   2
System rebuild extent:     32 Mbytes
System rebuild delay:      30
Journaling:                 Disabled

System Capacity 4485536 Mbytes, 4426536 Mbytes available.

```

Figure 2-4 Tier CLI Command Screen

Each letter under the “Disk Status” column represents a healthy drive at that channel. Verify that all drives can be seen by the InfiniteStorage 11000.

“Unhealthy” drives appear as follows:

- A blank space indicates that the drive is not present (or detected) at that location.
- A period (.) denotes that the disk was failed by the system.
- A question mark (?) indicates that the disk has failed the diagnostics tests or is not configured correctly.
- The character “r” indicates that the disk at that location is being replaced by a spare drive.

After entering the tier command, perform the following steps if necessary:

1. If a drive is not displayed at all (that is, it is “missing”), check to ensure that the drive is properly seated and in good condition. To search for the drive, type **disk scan** <Enter>.
2. If the same channel is missing on all tiers, check the cable connections for that channel.
3. If “automatic disk rebuilding” is not enabled, enable it by typing **tier autorebuild=on** <Enter>.

- To display the detailed disk configuration information for all of the tiers (see Figure 2-5), type **tier config**<Enter>.

```
11000[2]: tier config

                                Tier Configuration

Tier Owner LUNs   Disk Status   Installed Disks   Healthy Disks   F R   Sp Sp   Spare   Spare Repl
-----
1 + 1   1   ABCDEFGHPS ABCDEFGHPS ABCDEFGHPS
2 + 1   2   ABCDEFGHPS ABCDEFGHPS ABCDEFGHPS
3 + 1   2   ABCDEFGHP. ABCDEFGHPS ABCDEFGHP S
4 + 1   9   ABCDEFGHPS ABCDEFGHPS ABCDEFGHPS
121
*****

Disk 3S Failed.

Automatic disk rebuilding: Enabled
Automatic SMART replacement: Disabled
```

Figure 2-5 Tier Configuration CLI Command Screen

Tier Configuration Heading Definitions

- **Total LUNs.** LUNs that currently reside on the tier.
- **Healthy Disk.** The “health” of the spare disk currently being used (if any is being used) to replace a disk on the listed tier. The health indication for the spare channel that is physically on the listed tier is found under **SP H**.
- **F** indicates the failed disk (if any) on the tier.
- **R** indicates the replaced disk (if any) on the tier.
- **Sp H** indicates if the spare disk that is physically on the tier is healthy.
- **Sp A** indicates if the spare disk that is physically on the tier is available for use as a replacement.
- **Spare Owner** indicates the current owner of the physical spare, where ownership is assigned when the spare is used as a replacement. “**RES-#**” will appear under the Spare

Owner heading while a replacement operation is underway to indicate that unit “#” currently has the spare reserved.

- **Spare Used on** indicates the tier (if any) where this physical spare is being used as a replacement.
- **Repl Spare from** indicates the tier (if any) whose spare disk is being used as a replacement for this tier.

Note: Tiers are 8+1 mode by default.

Cache Coherency and Labeling in Dual Mode

Use the **DUAL** command to check the status of the units that are healthy and verify that the “Dual” (COM2) and “Ethernet” (LINK) communication paths between the two InfiniteStorage 11000 controller units are established (Figure 2-6).

```

11000[2]: dual

                                Dual Unit Configuration
                                -----
                                Unit 1                               Unit 2
                                -----
Label                            11000[1]                         11000[2]
Status                            Healthy                          Healthy

Dual communication:              established.
Ethernet communication:          healthy!
Cache coherency:                 established.
Cache coherency timeout:         0

```

Figure 2-6 Dual CLI Command Screen

If you require multi-pathing to the LUNs, enable cache coherency. If you do not require multi-pathing, disable cache coherency. To enable/disable the cache coherency function, enter the following (**ON** enables, **OFF** disables): **dual coherency=on|off**<Enter>

You may change the label assigned to each controller unit, (see Figure 2-7 on page 26). This allows you to uniquely identify each unit in the system. Each controller can have a label of up to 31 characters long.

1. To change the label (Figure 2-7), type **dual label** <Enter>.

```

11000[2]: dual label
Enter the number of the unit you wish to rename.
  LABEL=1 for unit 1, 11000[1]
  LABEL=2 for unit 2, 11000[2]
Unit: 2

Enter a new label for unit 2, or DEFAULT to return to the default label.
Up to 31 characters are permitted.
Current unit name: 11000[2]
New unit name: SYSTEM[2]

                                Dual Unit Configuration
                                -----
                                Unit 1                               Unit 2
-----
Label 11000[1]                   SYSTEM[2]
Status      Not connected                Healthy

Dual communication:      not established.
Ethernet communication:  failed!
Cache coherency:        not established.
Cache coherency timeout: 0

SYSTEM[2]:
    
```

Figure 2-7 Dual-Label CLI Command Screen

2. Select which unit you want to re-label.
3. When prompted, type in the new label for the selected unit. The new name is displayed.

Configuring the Storage Arrays

When you have determined your array configuration, you need to create and format the LUNs. You have the option of creating a 32-bit or a 64-bit address LUN. In the example below, two LUNs (32-bit addressing) are created:

- LUN 0 on Tiers 1 to 8 with capacity of 8192MB each.
- LUN 1 on Tiers 1 and 2 with capacity of 8192MB each.

Press **e** at any time to exit and cancel the command completely.

Note: In dual mode, LUNs are “owned” by the controller unit where they are created. Hosts only see the LUNs on the controller they are connected to, unless cache coherency is enabled.

1. To display the current cache settings (see Figure 2-8), type **cache**<Enter>

```
11000[2]: cache

                          Current Cache settings

LUN   Write   Maximum  MF  Prefetch
      Caching Prefetch Bit  Ceiling
-----
  0   Enabled   x    1  On   65535
  1   Enabled   x    1  On   65535
  2   Enabled   x    1  On   65535
  3   Enabled   x    1  On   65535
  4   Enabled   x    1  On   65535
  5   Enabled   x    1  On   65535
  6   Enabled   x    1  On   65535
  7   Enabled   x    1  On   65535
  8   Enabled   x    1  On   65535
  9   Enabled   x    1  On   65535
 10   Enabled   x    1  On   65535
 11   Enabled   x    1  On   65535
 12   Enabled   x    1  On   65535
 13   Enabled   x    1  On   65535

Writeback Limit: 75%

                2560.0 Mbytes of Cache Installed
                (2048 Segments of 1024 Kbytes)
```

Figure 2-8 Cache CLI Command Screen

2. Select a cache segment size for your array. For example, to set the segment size to 128KBytes, type **cache size=128**<Enter>
3. This setting can also be adjusted on-the-fly for specific application tuning: see Chapter 3 for additional information). The default setting is 1024.
4. Type **lun**<Enter>.
5. The Logical Unit Status chart should be empty, as no LUN is present on the array.
6. To create a new LUN, type **lun add=x**<Enter>

7. where **X** is the LUN number. Valid LUN numbers are 0..1023.
If only **lun add** is entered, you are prompted to enter a LUN number.
8. You will be prompted to enter the parameter values for the LUN. In this example:
 - Enter a label for the LUN (you can include up to 12 characters). The label may be changed later using the LUN LABEL command.
 - Enter the capacity (in Mbytes) for a single LUN in the LUN group: **8192**<Enter>
 - Enter the number of tiers to use: **8**<Enter>
 - Select the tier(s) by entering the Tier number. Enter each one on a new line. Tiers are numbered from 1 through 125.
 - **1** <Enter>
 - 2** <Enter>
 - 3** <Enter>
 - 4** <Enter>
 - 5** <Enter>
 - 6** <Enter>
 - 7** <Enter>
 - 8** <Enter>
 - Enter the block size in Bytes: **512**<Enter>

Note: **512** is the recommended block size. A larger block size may give better performance. However, verify that your OS and file system can support a larger block size before changing the block size from its default value.

This message will display: `Operation successful: LUN 0 added to the system`

9. When you are asked to format the LUN, type: **y**<Enter>

After you have initiated LUN format, the message `Starting Format of LUN` is displayed. You can monitor the format progress by entering the command **LUN** (see Figure 2-9 on page 29).

Upon completion, this message: `Finished Format of LUN 0` displays.

```

11000 [1]: lun
                Logical Unit Status
LUN  Label  Owner   Status   Capacity  Block   Tiers  Tier List
      (Mbytes) Size
-----
0    1      1      Format 14%  8192     512     8     1 2 3 4 5 6 7 8
                System Capacity 2240096 Mbytes, 2207328 Mbytes available.

```

Figure 2-9 Logical Unit Status - Formatting

10. Enter the command **LUN** to check the status of the LUN, which should be “Ready” (reference Figure 2-10).

```

11000 [1]: lun
                Logical Unit Status
LUN  Label  Owner   Status   Capacity  Block   Tiers  Tier List
      (Mbytes) Size
-----
0    1      1      Ready    8192     512     8     1 2 3 4 5 6 7 8
                System Capacity 2240096 Mbytes, 2207328 Mbytes available.

```

Figure 2-10 LUN Status - Ready

11. To create the LUN 1, type: **lun add=1**<Enter>
12. Enter these parameters:
- Enter a label for the LUN 1
 - For capacity, enter the value in MBytes: **8192**<Enter>
 - Enter the number of tiers to use: **2**<Enter>
 - Select the tier(s) by entering the Tier number. Enter each one on a new line and press the <Enter> key. The tiers are numbered from 1 through 125.
 - **1**<Enter>
 - **2**<Enter>
 - Enter the block size in Bytes: **512**<Enter>
 - -When asked to format the LUN, type: **y**<Enter>

Note: LUN format is a background process and you can start adding the next LUN as soon as the format for the previous LUN has started.

Setting Security Levels

After you have formatted all the LUNs, you can define users' access rights. Configurations come in two types:

- authorized user
- host port zoning

The **Authorized User** configuration is highly recommended for use in a SAN environment-- your data is completely secured and no accidental plug-in is allowed to do damage such as data change or deletion. Authorized users have access only to their own and "allowed to share" data. Administrators can also restrict users' access to the host ports and their read/write privileges to the LUNs. Another advantage of this configuration is that the users see the same LUN identification scheme regardless of the host port connection.

The **Host Port Zoning** configuration provides the minimum level of security. The LUN mappings change according to the host port connection. The read-only and read/write privileges can be specified for each LUN.

The **place holder** LUN feature allows the InfiniteStorage 11000 administrator to map a zero-capacity LUN to a host or group of hosts (via zoning or user authentication). The administrator can then create a real LUN and map it to the host(s) to replace the **place holder** LUN in the future. In most cases, the host does not have to reboot since it already mapped to the **place holder** LUN.

Note: Support of place holder LUNs is dependent upon the OS (operating system), the driver, and the Host Bus Adapter.

User Authentication (Recommended for SAN Environment)

Each user connected to the InfiniteStorage 11000 is identified by a **World Wide Name** (WWN) or GUID, and is given a unique user ID number. The controller can store configurations for up to 512 users and the security settings apply to all host ports.

Below is an example for adding two users to a system containing two LUNs (numbered 0 and 1). Each user has an internal LUN 1 is shared and "read-only." Both users see the shared LUN as LUN 0 and they see their own LUN as LUN 1. User 1 has access to host ports 1 and 4 while User 2 only has access to host port 2.

Prior to adding any users, verify that no “anonymous” access is allowed to the system:

1. Type **zoning**_{<Enter>}
2. Check to ensure that the LUN Zoning chart is empty (see Figure 2-11).

```
11000[2]: zoning

Port Zoning Summary:

          LUN Zoning
Port  World Wide Name  (External LUN, Internal LUN)
-----
  1   25000001FF060002  000,000    004,004    013,013
  2   26000001FF060002  000,001
  3   27000001FF060002
  4   28000001FF060002

11000[2]:
```

Figure 2-11 LUN Zoning CLI Screen

To add a user:

1. Type: **user audit=on**_{<Enter>} The 11000 reports which users are connected.
2. Type **user add**_{<Enter>} (see Figure 2-12).

```
11000[2]: USER ADD

Currently logged-in Anonymous Users:

ID  User Name      World Wide Name  S_ID  Port Time Logged in
-----
-- There are No Anonymous Users currently logged in to the system. --

Enter: 's' to specify a new Host User's world wide name, or
      'e' to escape: s

Enter: A 64-bit world wide name, or
      'e' to escape:
```

Figure 2-12 User Add CLI Command Screen

3. Specify a new Host User’s world wide name, enter **s**.

4. Specify a 64-bit world wide name or GUID, taken from the list of available anonymous users.
5. Enter an alias name for the user. The name may contain up to 12 characters. Type in a name and press `<Enter>`.
6. Host users can have their port access zoned. Enter **y** to specify host port zoning.
7. For Unit 1, enter each active port on a new line and then exit. For this example, type:


```

1<Enter>
4<Enter>
e<Enter>
            
```

8. For Unit 2, enter each active port on a new line and then exit. For this example, type:


```

1<Enter>
4<Enter>
e<Enter>
            
```

Host users are limited to accessing specific LUNs, as follows:

- a host user may have its own unique LUN mapping, or
- a host user may use the anonymous LUN mapping.

The anonymous user LUN mapping is handled by the port ZONING command. In either case, the LUN mapping applies on all the ports for which the user has been zoned.

9. Enter **y** to specify the unique LUN mapping.
10. Enter a new unique LUN mapping for this user. Options are shown in Table 2-3.

Table 2-3 LUN Mapping Options

Option	Description
G.1	GROUP.LUN number
P	place-holder
R	Read-Only. Place before the GROUP.LUN
N	Clear current assignment
<cr>	No Change
E	EXIT
?	Display detailed help text.

11. Connect user 2 and repeat steps 2--10 to specify the host port zoning and LUN mappings with the following changes:
12. For active host port (step 6), enter port 2 only.

13. For LUN mapping:
14. External LUN 0 is mapping to internal LUN: R1
External LUN 1 is mapping to internal LUN: 1
External LUN 2 is mapping to internal LUN: q

Note: In this scheme, users 1 and 2 have their own custom LUN identification scheme. The internal LUN 1 that is shared by the users needs to be managed by SAN management software. The individual dedicated LUN appears to the user as local storage and does not require external management software.

15. To display the new security settings, type: **user**_{<Enter>}

Host Port Zoning (Anonymous Access)

Host Port Zoning (Anonymous Access) should only be used for a non-SAN environment. Users are given “general admission” to the data.

Caution: Anonymous Access (host port zoning) provides only the minimum level of security.

One zoning configuration is supported for each of the host ports. Any unauthorized user accessing the controller storage is considered “anonymous” and granted zoning access for the host port to which they are connected. Given below is an example for adding LUN zoning to host port 1. External LUN 1 is mapped to internal LUN 0 and it is read-only for the users.

1. To edit the default zoning on a host port, type: **zoning edit**_{<Enter>}. The current settings are displayed.
2. Select a host port (1..4): **1**_{<Enter>}
3. Specify the internal LUN (0..1023) to be mapped to the external LUN. The new settings will display.
4. Repeat steps 1-3 to configure the other host ports.

Storage System Management

Managing the InfiniteStorage 11000

The InfiniteStorage 11000 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, Remote Login Management, Security Administration, and Firmware Update Management. Bundled together, this is called the InfiniteStorage 11000's *Administrator Utility*.

Management Interface

SAN management information for the InfiniteStorage 11000 can be accessed locally through a serial interface, or remotely through Telnet.

Important: An InfiniteStorage 11000 may have only one active login (serial or Telnet) at any given time.

Locally - Serial Interface

Any RS-232 terminal or terminal emulator (such as Microsoft Windows HyperTerminal) can be used to configure and monitor the InfiniteStorage 11000.

1. Connect your terminal to the CLI port at the back of the InfiniteStorage 11000 using a standard DB-9 female-to-male null modem cable (see Figure 3-1).

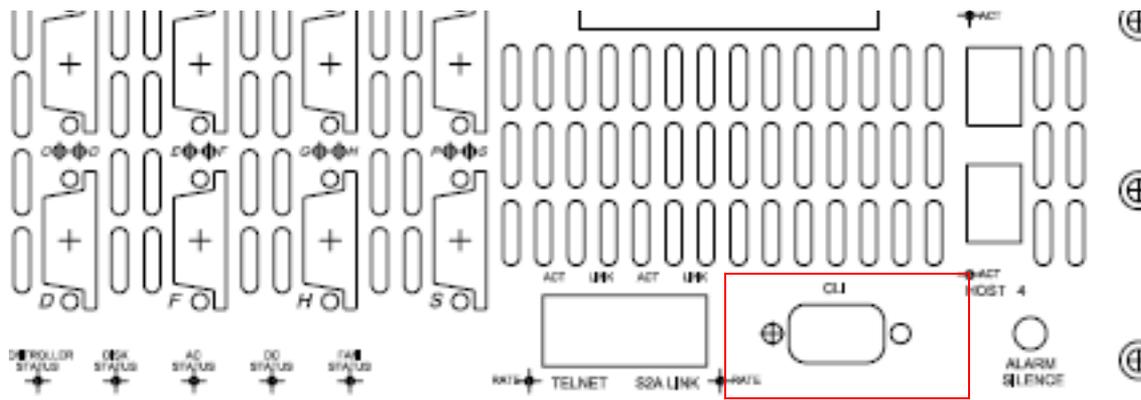


Figure 3-1 CLI Port on the Controller’s Rear Panel

2. Open your terminal window and set the settings for your serial port (see Table 3-1).

Table 3-1 Serial Port Settings and Values

Setting	Value
Bits per second:	115,200
Data bits:	8
Parity:	None
Stop bits:	1
Flow Control:	None

3. With the InfiniteStorage 11000 ready, press **<Enter>** to get the controller prompt.

Note: To change the baud rate on InfiniteStorage 11000, see section "Changing Baud Rate for the CLI Interface" in this guide.

Remotely - Telnet

To configure and monitor the InfiniteStorage 11000 remotely, connect the controller to your Ethernet network. Refer to the Remote Management section in this chapter for information on how to set up the controller's network interface.

Available Commands

Use the **Help** command to display the available commands within the InfiniteStorage 11000 Utility. To get help information on a command, type the command followed by a question mark. For example, to display help on cache options on the system, type: **cache?**_{<Enter>}

Administrator and User Logins

The **login** command allows the user to log into a (new) terminal or Telnet session at a specific security user level—*administrative* or *general purpose*. You will need Administrator access on the InfiniteStorage 11000 controller in order to change the system configurations.

For RS-232 terminal session, the general purpose user does not require login. For a Telnet session, you are required to login as either an administrator or a general purpose user. If you login as an administrator, you will have access to all the management and administrative functions. You can obtain status information and make changes to the system configuration.

At the general purpose user access level, you are only allowed to view status and configuration information. If the InfiniteStorage 11000 determines that the individual does not have the proper privileges, it will return a message (where the “user entered command” represents a command keyed in by the user):

```
<user entered command>: Permission denied
```

Login

To login to the system, do the following:

1. To login, enter: **login**_{<Enter>}
The prompt will display `Enter a login name:`
2. Enter a login name.

This prompt will display `Enter the password:`

3. Enter a password. If correct user name and password is entered, the appropriate prompt and label will display.

Note: The default administrator account name is “admin” and its password is “password.” Similarly, the default user account name is “user” and its password is “password.”

Logout

To logout of the system, enter: **logout**_{<Enter>} For a terminal session, you are returned to the general purpose user level. For Telnet, the current session is disconnected.

Password

Entering the **PASSWORD** command allows the administrator to change the login names and passwords for administrative and general purpose users (Figure 3-2 on page 39). The associated privileges remain the same regardless of the name or password changes.

```
11000[2]: password

Enter current password:
*****
Enter a new name to replace 'admin', or return to leave unchanged:

Administrative user name 'admin' unchanged.

Enter new password:
*****
Re-enter the new password:
*****

Enter a new name to replace 'user', or return to leave unchanged:

General user name 'user' unchanged.

Enter new password:

Password for general user 'user' unchanged.

Committing changes.
```

Figure 3-2 Password on the CLI Command Screen Example

Login names and passwords can be changed using the **PASSWORD** command, via RS-232 or Telnet. By default, the administrator name is “admin” and its password is “password”. Similarly, the default user name is “user” and its password is “password.” If a user forgets the password, entering **PASSWORD DEFAULTS** while logged in as “admin”, will restore all passwords and user names to the default values.

Who Am I

To display the owner and security level of the current terminal or Telnet session, enter:
whoami<Enter> (see Figure 3-3).

```
11000 [1]: whoami
CLI session:
Current owner      : admin.
Current security level: Administrative.
```

Figure 3-3 The whoami Screen

Configuration Management

The InfiniteStorage 11000 provides uniform configuration management across heterogeneous SANs. Status of host ports (see Table 3-2) and storage assets are continuously being monitored.

Table 3-2 InfiniteStorage 11000 Controller Configuration Limits

Item	Limit
Number of LUNs	1024 (0-1023)
Total Number of Users	4096
Number of LUNs Per User	255
Number of LUNs per port (zoning)	255
Number of FC logins per port	128
Max number of tiers per LUN	8
Max number of tiers	128
Max size of 32-bit LUN	0xFFFF0000 (blocks)
Max size of 64-bit LUN	128 TB
Granularity of LUN size	2 MB x number of tiers
Support LUN block sizes	512 bytes, 1K, 2K, 4K
Active host commands	16 per port
Internal Log Size	192 KB

Table 3-2 (continued) InfiniteStorage 11000 Controller Configuration Limits

Item	Limit
Active disk commands	128
Max disk pieces per disk command	16
Max queued commands per host port (does not include active host commands)	256
Max commands per disk	32 (can be lower for SATA)
Max supported capacity per drive	2 TB

Configure and Monitor Status of Host Ports

The status information of the host ports can be obtained at any time. The **HOST** command displays the current settings and status for each host port (see Figure 3-4). It also displays a list of the users currently logged into the system. An unauthorized user is given the user name **Anonymous**.

```

11000[2]: host

                               Host Fibre-Channel Port Configuration

Host   Hard   Current   Port Speed   Timeout
Port  Loop ID  Loop ID   Desired/Actual seconds   World Wide Name   Loop Status
-----
  1    EF     EF        Auto/-NA-    75    25000001FF060002 Not connected
  2    EF     EF        Auto/-NA-    75    26000001FF060002 Not connected
  3    EF     EF        Auto/-NA-    75    27000001FF060002 Not connected
  4    EF     EF        Auto/-NA-    75    28000001FF060002 Not connected

                               Current Logins

                               Frame/ S_ID/
User      Port  MTU   LID   World Wide Name   Login
-----

```

Figure 3-4 Host CLI Command Screen Example

The **PORT=X|ALL** parameter specifies the specific host port(s) (1 to 4) to be affected when used in combination with any of the other parameters: **ID**, **TIMEOUT**, **SPEED** (for FC only), **WWN**. The default is to apply changes to ALL host ports.

HOST ID=<new ID> changes the hard loop ID of a host port. The system selects a soft ID if the hard loop ID is already taken by another device. This parameter is entered as an 8-bit hex value. The default value is EF.

HOST WWN=X|DEFAULT overrides the system ID and specifies a different World Wide Name (WWN) for a host port. This parameter is entered as an 64-bit hex value. Default WWN is based on the serial number of the unit.

HOST STATUSCLEAR resets the error counts.

HOST STATUS displays the loop status of each host port (see Figure 3-5 on page 43).

```

11000[2]: host status

```

Host Port FC Status				
	Port 1	Port 2	Port 3	Port 4
Link Failures	0	0	0	0
Loss of Sync	0	0	0	0
Loss of Signal	0	0	0	0
Primitive_seq_errs	0	0	0	0
Word Alignment	0	0	0	0
CRC Errors	0	0	0	0
CTIO timeouts	0	0	0	0
CTIO PCI Errors	0	0	0	0
CTIO Xmit Errors	0	0	0	0
CTIO Other Errors	0	0	0	0
Receive Errors	0	0	0	0
Queue Full Errors	0	0	0	0
No-op Timeouts	0	0	0	0
No-op Resets	0	0	0	0

Host Port Data Integrity Status				
	Port 1	Port 2	Port 3	Port 4
Parity Errors	0	0	0	0
Soft Bus PE	0	0	0	0
Recovered Bus PE	0	0	0	0
UnRecovered Bus PE	0	0	0	0
BusPE on Unused ch	0	0	0	0
Soft Array PE	0	0	0	0
Recovered Array PE	0	0	0	0
UnRecovered Arr PE	0	0	0	0
WFIFO Cor ECC Errs	0	0	0	0
WFIFO Unc ECC Errs	0	0	0	0
WFIFO code Errors	0	0	0	0
WFIFO Length Errs	0	0	0	0


```

Link Status      Not connected  Not connected  Not connected  Not connected

```


Host Slot Error Counts		
	Slot12	Slot34
Stage Buff ECC Correctable Errors:	0	0
Stage Buff ECC Uncorrectable Errors:	0	0
Stage Buff ECC PCIX PERR:	0	0
Stage Buff ECC PCIX SERR:	0	0

Figure 3-5 Host Status CLI Command Screen Example

HOST SPEED lets you display and change the port speed on the host port(s). You are prompted for the desired speed as well as for the choice of host port(s) (see Figure 3-6 on page 44).

```

11000[1]: host speed

                                Host Fibre-Channel Port Configuration
Host   Hard   Current   Port Speed   Timeout
Port  Loop ID Loop ID   Desired/Actual seconds   World Wide Name   Loop Status
-----
  1    EF     EF       Auto/ 4.0    75    21000001FF040002 Good
  2    EF     EF       Auto/ 4.0    75    22000001FF040002 Good
  3    EF     EF       Auto/ 4.0    75    23000001FF040002 Good
  4    EF     EF       Auto/ 4.0    75    24000001FF040002 Good
Enter the host port to change (1..4, 'A' for all or 'Q' to quit): 1

Host port 1 selected
Please select the desired port speed from the following menu:
  Auto negotiate           (Enter 0)
  1 Gigabit per second    (Enter 1)
  2 Gigabit per second    (Enter 2)
  4 Gigabit per second    (Enter 4)
Please enter selection (0, 2, 4 or 'Q' to quit): 0

Auto negotiate selected

                                Host Fibre-Channel Port Configuration
Host   Hard   Current   Port Speed   Timeout
Port  Loop ID Loop ID   Desired/Actual seconds   World Wide Name   Loop Status
-----
  1    EF     EF       Auto/ 4.0    75    21000001FF040002 Good
  2    EF     EF       Auto/ 4.0    75    22000001FF040002 Good
  3    EF     EF       Auto/ 4.0    75    23000001FF040002 Good
  4    EF     EF       Auto/ 4.0    75    24000001FF040002 Good

                                Current Logins
User      Port  MTU   Frame/ S_ID/ LID   World Wide Name   Login
-----
Anonymous  1  2048   E8 210000E08B86CE9C THU SEP 11 17:28:15 2008
Anonymous  2  2048   E8 210000E08B9DF1E5 THU SEP 11 17:28:41 2008
Anonymous  3  2048   E8 210000E08B88882B THU SEP 11 17:27:30 2008
Anonymous  4  2048   E8 210000E08B9D4FE6 THU SEP 11 17:27:56 2008

```

Figure 3-6 Host Speed CLI Command Screen Example

Configure and Monitor Status of Storage Assets

Use the information and examples in the following sections to help configure and monitor the status of your InfiniteStorage 11000 system components.

Disk and Channel Information

The **DISK** command displays the current disk configuration and the status of the ten disk channels (ABCDEFHGHS) on the InfiniteStorage 11000InfiniteStorage 11000 controller (see Figure 3-7).

```
11000[2]: disk
  Disk Channel Status
Channel A SAS 3.0 Gb - healthy.
Channel B SAS -NA- - not ready.
Channel C SAS 3.0 Gb - healthy.
Channel D SAS -NA- - not ready.
Channel E SAS 3.0 Gb - healthy.
Channel F SAS -NA- - not ready.
Channel G SAS 3.0 Gb - healthy.
Channel H SAS -NA- - not ready.
Channel P SAS 3.0 Gb - healthy.
Channel S SAS -NA- - not ready.
Disk 3S Failed.
Disk rebuild verify:           Disabled
Disk total command timeout:    68 seconds
Disk command abort/retry timeout: 31 seconds
Disk maximum commands:        16
Disk maximum read length:      128 Kbytes
Disk maximum write length:     2048 Kbytes
Auto Reassign:                 Enabled
Write Same command:           Disabled
Disk commands outstanding: 0
```

Figure 3-7 Disk CLI Command Screen Example

If the channel status is “acquiring loop synchronization,” this may indicate a channel problem. Refer to Chapter 4 in this guide for recovery information.

DISK INFO=<tier><channel> displays information about a specific disk (tier, channel).

DISK LIST displays a list of the disks installed in the system and indicates how many were found (see Figure 3-8 on page 46).

```

11000[1]: disk list
                Disks Installed
            Vendor  Product ID      Mbytes RPM Revision Serial Number
-----
Disk  1A HITACHI  HUS153014VLS300  140205 15K   A410 J4VV94VA
Disk  1B HITACHI  HUS153014VLS300  140205 15K   A410 J4WG78HA
Disk  1C HITACHI  HUS153014VLS300  140205 15K   A410 J4WG7J5A
Disk  1D HITACHI  HUS153014VLS300  140205 15K   A410 J4WG7HGA
Disk  1E HITACHI  HUS153014VLS300  140205 15K   A410 J4WG7M1A
.
.
Disk  8F HITACHI  HUS153014VLS300  140205 15K   A410 J4WG7H5A
Disk  8G HITACHI  HUS153014VLS300  140205 15K   A410 J4WG7PNA
Disk  8H HITACHI  HUS153014VLS300  140205 15K   A410 J4WG7MMA
Disk  8P HITACHI  HUS153014VLS300  140205 15K   A410 J4WG7A8A
Disk  8S HITACHI  HUS153014VLS300  140205 15K   A410 J4W5DV5A

Found: 79 HITACHI  HUS153014VLS300  140205 15K   A410
Found:  1 HITACHI  HUS153014VLS300  140205 15K   A140

```

Figure 3-8 Disk List CLI Command Example Screen

DISK SCAN checks each disk channel in the system for any new disks and verifies that the existing disks are in the correct location. **DISK SCAN** also starts a rebuild operation on any failed disks which pass the disk diagnostics.

DISK STATUS displays the loop status of each disk channel and a count of the SAS errors encountered on each channel (see Figure 3-9 on page 47).

```

11000[2]: disk status
                Disk Channel Status
                -----
                A      B      C      D      E      F      G      H      P      S
-----
InValDw:0      0      0      0      0      0      0      0      0      0
InValDw:1      0      0      0      0      0      0      0      0      0
InValDw:2      0      0      0      0      0      0      0      0      0
InValDw:3      0      0      0      0      0      0      0      0      0
LsDwSyn:0      0      0      0      0      0      0      0      0      0
LsDwSyn:1      0      0      0      0      0      0      0      0      0
LsDwSyn:2      0      0      0      0      0      0      0      0      0
LsDwSyn:3      0      0      0      0      0      0      0      0      0
PhyRst:0       0      dfe6  0      dfe7  0      dfe6  0      dfe6  0      dfe6
PhyRst:1       0      dfe6  0      dfe6  0      dfe7  0      dfe7  0      dfe7
PhyRst:2       0      dfe6  0      df35  0      dfe7  0      dfe7  0      dfe7
PhyRst:3       0      dfa8  0      dfe5  0      dfe7  0      dfe7  0      dfe7
RunDisp:0      0      0      0      0      0      0      0      0      0
RunDisp:1      0      0      0      0      0      0      0      0      0
RunDisp:2      0      0      0      0      0      0      0      0      0
RunDisp:3      0      0      0      0      0      0      0      0      0
Recovery       0      0      0      0      0      0      0      0      0

```

Figure 3-9 Disk Status CLI Command Screen Example

DISK STATUSCLEAR=<tier><channel> resets the channel error counts.

DISK DEFECTLIST=<tier><channel> displays the defect list information for a specified disk. The tier is in the range of <1..128>. The channel is one of the following: **A**, **B**, **C**, **D**, **E**, **F**, **G**, **P**, or **S**. The list is classified into two types: **G** (Grown) and **P** (Permanent) (see Figure 3-10).

- The **G** list consists of the sectors that have become bad after the disk has left the factory and which can be added to at any time.
- The **P** list consists of the bad sectors that are found by the disk manufacturer.

```

11000[2]: disk defectlist

Enter the disk by specifying its tier and channel locations 'tc', where:
  't' indicates the tier in the range <1..128>, and
  'c' indicates the channel in the range <ABCDEFGHPS>.

Disk: 1a
Number of defects in P List = 722
Number of defects in G List = 13
11000[2]:

```

Figure 3-10 Disk Defectlist CLI Command Screen

DISK FAIL=<tier><channel> instructs the system to fail the specified disk at the physical tier in the range of <1 ...128> and channel in the range of <ABCDEFGHPS>. 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 and it is currently in use as a replacement for a failed disk, the disk that the spare is replacing is put back to a failed status and the spare is released, but marked as unhealthy and unavailable.

DISK PLS=<tier><channel> requests/displays the PHY Link Error Status Block (PLS) information for the specified drive (see Figure 3-11 on page 50). The disk is specified by its physical tier and channel locations, 'tc', where:

- 't' indicates the tier in the range <1..128>, and
- 'c' indicates the channel in the range <ABCDEFGHPS>.

If neither the tier nor the channel are specified, the PLS information is requested from all drives. If only the tier is specified, the PLS information is requested from all the drives on the specified tier. Table 3-3 shows the types of PHY errors. Note that SATA and SAS drives report PHY errors differently.

Table 3-3 PHY Link Error Status Block Information

SATA AAMUX PHY ERRORS	
H-RX	The number of SATA FIS CRC errors received on the host port of the AAMUX
H-TX	The number of SATA R_ERR primitives received on the host port indicating a problem with the transmitter of the AAMUX
H-Link	The number of times the PHY has lost link on the host port.
H-Disp	The number of frame errors for the host port of the AAMUX. These include: code error, disparity error, or realignment
O-RX	The number of SATA FIS CRC errors received on the other host port of the AAMUX
O-TX	The number of SATA R_ERR primitives received on the other host port indicating a problem with the transmitter of the AAMUX
O-Link	The number of times the PHY has lost link on the other host port.
O-Disp	The number of frame errors for the other host port of the AAMUX. These include: code error, disparity error, or realignment.
D-RX	The number of SATA FIS CRC errors received on the device port of the AAMUX

Table 3-3 (continued) PHY Link Error Status Block Information

D-TX	The number of SATA R_ERR primitives received on the device port indicating a problem with the transmitter of the AAMUX.
D-Link	The number of times the PHY has lost link on the device port.
D-Disp	The number of frame errors for the device port of the AAMUX. These include: code error, disparity error, or realignment.
SAS PHY ERRORS	
InvDW	Invalid DWORD Count - The number of invalid dwords received outside of the PHY reset sequence.
RunDis	Running disparity Count - The number of dwords containing running disparity errors received outside of the PHY reset sequence
LDWSYN	Loss of DWORD synchronization count - The number of times the PHY has lost synchronization and restarted the link reset sequence
PHYRES	PHY Reset Problem count - The number of times the PHY reset sequence has failed.

Note: SATA drives have an Active/Active MUX (AAMUX) installed. Error counts are read directly from the AAMUX.

```

11000 [1]: disk pls

                Tier 1
            PHY Error Status Blocks

Channel  A  B  C  D  E  F  G  H  P  S
-----
H-RX:    .  .  .  .  .  .  .  .  .  .
H-TX:    .  .  .  .  .  .  .  .  .  .
H-Link:   .  .  .  .  .  .  .  .  .  .
H-Disp:   .  .  .  .  .  .  .  .  .  .
O-RX:    .  .  .  .  .  .  .  .  .  .
O-TX:    .  .  .  .  .  .  .  .  .  .
O-Link:   .  .  .  .  .  .  .  .  .  .
O-Disp:   .  .  .  .  .  .  .  .  .  .
D-RX:    .  .  .  .  .  .  .  .  .  .
D-TX:    .  .  .  .  .  .  .  .  .  .
D-Link:   .  .  .  .  .  .  .  .  .  .
D-Disp:   .  .  .  .  .  .  .  .  .  .
InvDW:    0  0  0  0  0  0  0  0  0  0
RunDis:   0  0  0  0  0  0  0  0  0  0
LDWSYN:   0  0  0  0  0  0  0  0  0  0
PHYRES:   0  0  0  0  0  0  0  0  0  0

```

Figure 3-11 Disk PLS Tier 1 Status Screen Example

For other **DISK** parameters, see the section “Automatic Drive Rebuild” on page 71 in this guide.

Tier View

Tiers (also known as RAID groups) are the basic building blocks of the InfiniteStorage 11000. In an 8+1 mode, a tier contains 10 drives: eight (8) data drives (Channels **A** through **H**), one(1) parity drive (Channel **P**), and an optional spare drive (Channel **S**). In an 8+2 mode, a tier contains 10 drives, but the setup is different: eight (8) data drives (Channels **A** through **H**), and two (2) parity drives (Channel **P** and **S**). Drives that have the same SAS ID across all ten channels are put on the same tier. Tiers are automatically added to the system when the disks are detected. A tier will automatically be removed if it is not in use by any of the LUNs and all of the disks in the tier are removed or moved to another location.

TIER displays the current status and configuration of the tiers in the system (see Figure 3-12 on page 51).

```

11000[2]: tier

                                Tier Status
                                Space
                                Capacity Available Disk
Tier Owner (Mbytes) (Mbytes) Status LUN List
-----
  1 + 1 1121384 1111384 ABCDEFGHPS 0
  2 + 1 1121384 1101384 ABCDEFGHPS 1 6
  3 + 1 1121384 1110384 ABCDEFGHP. 2 4
  4 + 1 1121384 1103384 ABCDEFGHPS 3 5 7 8 9 10 11 12 13

Disk 3S Failed.

Automatic disk rebuilding: Enabled
Automatic SMART replacement: Disabled
Maximum active rebuilds: 4
Maximum active verifies: 2
System rebuild extent: 32 Mbytes
System rebuild delay: 30
Journaling: Disabled

System Capacity 4485536 Mbytes, 4426536 Mbytes available.

```

Figure 3-12 Tier CLI Command Screen Example

The tiers' total and available capacities are shown under the "Capacity" and "Space Available" columns respectively.

The **TIER** command shows the status of each disk on the tier as follows:

- A letter <ABCDEFGHIPS> represents a healthy disk at that location.
- A space indicates that the disk is not present or detected.
- A period (.) denotes that the disk was failed by the system.
- The symbol "?" indicates that the disk has failed the diagnostic tests or is not configured correctly.
- The character "r" indicates that the disk was failed by the system and replaced by a spare disk.
- The symbol "!" indicates that the disk is in the wrong location.

Note: The rate of rebuild and format operations can be adjusted with the commands, **Tier Delay=x** and **Tier Extent=x**.

Tier Configuration

TIER CONFIG displays the detailed tier configuration information for all of the tiers (see Figure 3-13 for an example).

```

11000[2]: tier config

Tier Configuration

Tier Owner  Total  Disk  Installed  Healthy  Sp Sp  Spare  Spare  Repl
            LUNs  Status  Disks    Disks    F  R    H  A  Owner  Used  Spare
            -----
1 + 1      1  ABCDEFGHPS  ABCDEFGHPS  ABCDEFGHPS          8+2 Mode no spares
2 + 1      2  ABCDEFGHPS  ABCDEFGHPS  ABCDEFGHPS          8+2 Mode no spares
3 + 1      2  ABCDEFGHP.  ABCDEFGHPS  ABCDEFGHP  S      8+2 Mode no spares
4 + 1      9  ABCDEFGHPS  ABCDEFGHPS  ABCDEFGHPS          8+2 Mode no spares
121
*****

Disk 3S Failed.

Automatic disk rebuilding:  Enabled
Automatic SMART replacement: Disabled

```

Figure 3-13 Tier Configuration CLI Command Screen Example

The headings for the **Tier Configuration** screen indicate the following values or conditions for the tiers.

Total LUNs lists the number of LUNs that currently reside on the tier.

Note: The health indication for the spare channel under the **'Healthy Disks'** heading is an indication of the health of the spare disk (if any) that is currently being used to replace a disk on the listed tier. The health indication for the spare channel that is physically on the listed tier is found under the **'Sp H'** heading.

These headings indicate the respective conditions on the tier.

- F** failed disk (if any).
- R** replaced disk (if any).
- Sp H** whether the spare disk that is physically on the tier is healthy.
- Sp A** whether the spare disk that is physically on the tier is available for use as a replacement.

Spare Owner current owner of the physical spare tier, where ownership is assigned when the spare is used as a replacement.

Spare Used on tier (if any) on which this physical spare is being used as a replacement.

Repl Spare from tier (if any) whose spare disk is being used as a replacement. Tiers are in 8+1 mode by default.

Note: RES-# will display under the **Spare Owner** heading while a replacement operation is underway to indicate that unit '#' currently has the spare reserved.

TIER CONFIG=ALL displays tier configuration and replacements for both 8+1 and 8+2 modes, see Figure 3-14 for an example.

```
11000[2]: tier config=all

                                Tier Configuration

Tier  Owner  Total  Disk  Installed  Healthy  Sp  Sp  Spare  Used  Spare
Tier  Owner  LUNs  Status  Disks     Disks   F  R    H  A  Owner  on  Spare
-----
  1 +  1    1  ABCDEFGHPS ABCDEFGHPS ABCDEFGHPS      8+2 Mode no spares
  2 +  1    2  ABCDEFGHPS ABCDEFGHPS ABCDEFGHPS      8+2 Mode no spares
  3 +  1    2  ABCDEFGHP. ABCDEFGHPS ABCDEFGHP  S    8+2 Mode no spares
  4 +  1    9  ABCDEFGHPS ABCDEFGHPS ABCDEFGHPS      8+2 Mode no spares
121
*****

Disk 3S Failed.

Automatic disk rebuilding:  Enabled
Automatic SMART replacement: Disabled
```

Figure 3-14 Tier Config=ALL CLI Command Screen Example

LUN View

The **LUN** command displays the current status of the LUNs (see Figure 3-15 on page 54). Possible status and definitions are listed in Table 3-4 on page 55.

```

11000[2]: lun
                Logical Unit Status
                Capacity  Block
LUN  Label      Owner  Status    (Mbytes)  Size  Tiers Tier list
-----
  0              1   Ready      10000    512    1   1
  1              1   Ready      10000    512    1   2
  2              1   Degraded   10000    512    1   3
  3              1   Ready      10000    512    1   4
  4 AUTO_LUN4    1   Degraded   1000     512    1   3

System verify extent: 16 Mbytes
System verify delay:  30

                System Capacity 4485536 Mbytes, 4426536 Mbytes available.
    
```

Figure 3-15 LUN CLI Command Screen

The **LUN LIST** command displays a list of all valid LUNs in the system. The screen shows the label, owner, status, capacity, and serial number of each LUN (see Figure 3-16).

Table 3-4 on page 55 provides a list of status possibilities for the LUNs.

```

11000[2]: lun list
                Logical Unit Status
                Capacity  Serial
LUN  Label      Owner  Status    (Mbytes)  Number
-----
  0              1   Ready      10000    00E636B60000
  1              1   Ready      10000    00E636C70100
  2              1   Degraded   10000    108E36DD0200
  3              1   Ready      10000    108E36F70300
  4 AUTO_LUN4    1   Degraded   1000     000239710400
  5 AUTO_LUN5    1   Ready      1000     0002CFDE0500
  6              1   Ready      10000    00E6005F0600

                System Capacity 4485536 Mbytes, 4426536 Mbytes available.
    
```

Figure 3-16 LUN List Command Screen Example

Table 3-4 LUN Status Possibilities List

STATUS	DESCRIPTION
Not Installed	The LUN does not exist.
Config Error	The LUN has a configuration error.
Unavailable	Cache data for the LUN is being purged. The System purges cache data when the LUN is being deleted or mirrored by the other unit in the system.
Not Ready	The unit is still booting up.
Stopped	The unit was stopped.
Unformatted	The LUN has not been formatted.
Critical	The LUN has a failed disk.
Critical [GHS]	The LUN has a failed disk and a disk that is replaced by a spare.
Degraded	The LUN has a failed disk (8+2 Only).
Degraded [GHS]	The LUN has a failed disk and a disk that is replaced by a spare (8+2 Only).
Ready [GHS]	The LUN has a replaced disk.
Ready	The LUN is OK.

LUN Configuration

LUN CONFIG displays the configuration information for all the valid LUNs in the system (see Figure 3-17).

```

11000[2]: lun config

                Logical Unit Configuration

LUN      Capacity  Block   LUN      Tier   Tier
         (Blocks) Size    Offset  Start  End    Tier list
-----
0         1388000  512     0         0     270FFF 1
1         1388000  512     0         0     270FFF 2
2         1388000  512     0         0     270FFF 3
3         1388000  512     0         0     270FFF 4
4          1F4000  512     0    271000 2AF7FF 3
5          1F4000  512     0    271000 2AF7FF 4
6         1388000  512     0    271000 4E1FFF 2
7          1F4000  512     0    2AF800 2EDFFF 4
8          1F4000  512     0    2EE000 32C7FF 4
9          1F4000  512     0    32C800 36AFFF 4
10         1F4000  512     0    36B000 3A97FF 4
11         1F4000  512     0    3A9800 3E7FFF 4
12         1F4000  512     0    3E8000 4267FF 4
13         1F4000  512     0    426800 464FFF 4

                System Capacity 4485536 Mbytes, 4426536 Mbytes available.

11000[2]:

```

Figure 3-17 LUN Config CLI Command Screen Example

LUN Reservations

LUN RESERVATIONS displays a list of all valid LUNs in the system and shows which LUNs currently have a SCSI reservation and which initiator holds the reservation (see Figure 3-18).

```

11000[2]: lun reservations

                Current SCSI LUN Reservations

LUN  Port Name      Key              Type  Port  Date
-----
0    No SCSI Reservation
1    No SCSI Reservation
2    No SCSI Reservation
3    No SCSI Reservation
4    No SCSI Reservation
5    No SCSI Reservation
6    No SCSI Reservation
7    No SCSI Reservation
8    No SCSI Reservation
9    No SCSI Reservation
10   No SCSI Reservation
11   No SCSI Reservation
12   No SCSI Reservation
13   No SCSI Reservation

11000[2]:

```

Figure 3-18 LUN Reservations CLI Command Screen Example

LUN RELEASE releases any SCSI reservations and registrations on a LUN.

Adding/Removing Storage Assets

The InfiniteStorage 11000 supports up to 128 tiers. New tiers can be added without affecting system operations.

DISK SCAN checks each disk channel in the system for any new disks. New tiers are automatically added to the system when the disks are detected. A tier is automatically deleted if it is not in use by any of the LUNs and all of the disks in the tier are removed or moved to another location.

Status of Drive Enclosures

The **SES** command displays the failures reported by the enclosure (see Figure 3-19), through the *SCSI Enclosure Services* (SES). It also provides a means to access SES specific functions such as disk, channel, and LUN. Drive failures are not displayed using the SES command; you must use the **TIER** command to view drive status.

```
11000 [1]: ses
EncID:50050CC0000033C8: Power Supply 1:DC Power Failure
```

Figure 3-19 Displaying the Current Disk Enclosure Failures Using SES

If your enclosures provide redundant SES communication paths, the error is reported twice. In Figure 3-19, **EncID** is the **Enclosure Logical Identifier** of the enclosure that reported the failure. The last four digits of the WWN are the last four digits of the enclosure's serial number.

SES ON saves the SES state to the parameter blocks, and starts up the SES monitors.

SES OFF saves the SES state to the parameter blocks, and shuts down the SES monitors.

Display SES Devices Information

SES SHOWDEVICES displays all the SES devices on all channels.

SES SHOWALL displays all configuration information for all the SES devices on all channels.

SES SHOW=<tier><channel> displays the configuration information and the status information returned from an SES Enclosure Status page for the SES device for the specified drive in the range of <1..128> and <ABCDEFGHPS>.

Visual Indication of Drive

SES IDDISK=<tier><channel> provides a visual indication of the specified drive (<1..128><ABCDEFGHPS>). The status LED of the drive blinks until the command **SES ID=OFF** is issued. The **SES ID=OFF** command restores the system to its original visual state.

Visual Indication of Tier

SES IDTIER=<tier> provides a visual indication of the specified tier <1..128>. The status LED of the drives blinks until the command **SES ID=OFF** is issued, which restores the system to its original visual state.

Visual Indication of Channel

SES IDCHANNEL=<channel> provides a visual indication of the specified channel <ABCDEFGHPS>. The status LED of the drives blinks until the command **SES ID=OFF** is issued, which restores the system to its original visual state.

Tier Mapping for Enclosures

The InfiniteStorage 11000 controllers support various drive enclosures. When the system is first configured, it is necessary to select a tier mapping mode so that the position of the tiers in the system are changed to conform with the layout of your drive enclosures. The tier mapping information also allows the controller to properly light the enclosure fault LEDs.

TIER MAP displays the current mapping mode for the disks in the array. **TIER CHANGEMAP** changes the current tier mapping for the disks in the array.

Important: The **CHANGEMAP** command should only be used when the system is first configured. Changing the mapping mode will alter all the tier information, making LUN information inaccessible.

System Network Configuration

These commands do the following:

NETWORK displays the current network interface settings.

NETWORK USAGE displays the address resolution protocol map, ICMP (ping), general network, and IP, TCP, and UDP layer statistics.

NETWORK IP=<new address> changes the IP address. (The system must be restarted before the changes will take effect).

NETWORK NETMASK=<aaa.bbb.ccc.ddd> changes the netmask.

NETWORK GATEWAY=<aaa.bbb.ccc.ddd> sets the current gateway in the network routing table to the supplied Internet address. The gateway is where IP datagrams are routed when there is no specific routing table entry available for the destination IP network or host.

Note: **GATEWAY=<no Internet address>** clears out the current gateway.

NETWORK PRIVATE displays the MAC address for the private network device.

Telnet

NETWORK TELNET=ON|OFF enables/disables the Telnet capability on the InfiniteStorage 11000. The system must be restarted before the changes will take effect.

Note: To only temporarily affect Telnet session availability during a concurrent power-cycle, refer to the TELNET command in the section “Firmware Update Procedure” on page 110 in this chapter.

Important: Telnet connections are “clear text.” If Telnet connections are used, you may expose InfiniteStorage 11000 passwords to third parties. For greater security, turn off Telnet access if it is not required.

NETWORK TELNETPORT=<port number> changes the Telnet port number for the current InfiniteStorage 11000 controller. The system must be restarted before the changes will take effect. Valid port numbers are **0** to **32768**; however, the results may be unpredictable if the port number chosen is already in use (on this unit) by either the GUI or SYSLOG facilities. The default port number is **23**.

SNMP & Syslog

NETWORK SNMP=ON|OFF enables and disables the SNMP functionality. The system must be restarted before the changes will take effect.

NETWORK LIMIT_SNMP=ON|OFF specifies whether the SNMP functionality will only report component-level information, or all levels of information, The default setting is **OFF**.

NETWORK TRAPIP=<aaa.bbb.ccc.ddd> changes the destination IP address for SNMP trap packets. The system must be restarted before the changes will take effect.

NETWORK SYSLOG=ON|OFF enables and disables the Syslog capability.

Note: **NETWORK SYSLOG** should be enabled, since it is the best way to find out what occurred in the event of a problem. However, since some problems can produce a large amount of output, it is a good idea to have your **syslog** program configured to rotate based on log size rather than date.

NETWORK SYSLOGIP=<aaa.bbb.ccc.ddd> changes the destination IP address for syslog packets, Both InfiniteStorage 11000s in the couplet pair will share the same syslog destination IP address but each InfiniteStorage 11000 can specify a different destination port.

NETWORK SYSLOGPORT=<port number> changes the destination port number for syslog packets for the current InfiniteStorage 11000. Both controllers in the couplet pair will share the same syslog destination IP address but each controller can specify a different destination port. Valid ports are **0** to **32768**. However the results may be unpredictable if the port number chosen is already in use (on this unit) by either the TELNET or GUI facilities. Default port number is **514**.

Note: The InfiniteStorage 11000 sends syslog messages via the local 7 (23) facility.

Refer to the section on remote management in this chapter for information on how to set up Telnet and SNMP functionality on your host computer.

API Server Connections

NETWORK API_SERVER=ON|OFF enables/disables the API server capability. The GUI Management Tool relies on an active and enabled API Server for its communications with the system. The system must be restarted before the changes will take effect.

Note: To affect the API Server connection availability only temporarily during the current power-cycle, see the section on API server connections in this chapter.

NETWORK API_PORT=<port number> specifies the API Server port number for the current InfiniteStorage 11000. The system must be restarted before the changes will take effect. Valid ports are 0 to 32768. The results may be unpredictable if the port number chosen is already in use (on this unit) by either the TELNET or SYSLOG facilities. The default port number is 8008.

Displaying and Editing the Routing Table

The **ROUTE** command displays the current routing table of the system (see Figure 3-20 on page 63) and allows the administrator to change it. The routing table describes how the InfiniteStorage 11000 can communicate with the hosts on other networks.

```

11000[2]: route

Gateway:

Permanent Routing Table:
=====
destination      gateway
-----
-- None --      -- None --
-----

Current Routing Tables:
=====

ROUTE NET TABLE
Destination      Gateway          Flags  Refcnt Use      Interface
-----
192.168.0.0      192.168.10.50   0x101  4      0          gei0
-----

ROUTE HOST TABLE
Destination      Gateway          Flags  Refcnt Use      Interface
-----
127.0.0.1        127.0.0.1       0x5    0      15475     lo0
-----

```

Figure 3-20 Route CLI Command Screen Example

ROUTE ADD=<aaa.bbb.ccc.ddd> GATEWAY=<aaa.bbb.ccc.ddd> adds gateways to the routing table. Up to 6 permanent routes can be added to the tables. For example, to indicate that the machine with Internet address 91.0.0.3 is the gateway to the destination network 90.0.0.0, enter: **ROUTE ADD=90.0.0.0 GATEWAY=91.0.0.3**

ROUTE DEL=<aaa.bbb.ccc.ddd> GATEWAY=<aaa.bbb.ccc.ddd> deletes gateways from the routing table.

ROUTE GATEWAY=<aaa.bbb.ccc.ddd> sets the current gateway in the network routing table to the specified Internet address. The gateway is where IP datagrams are routed when there is no specific routing table entry available for the destination IP network or host. If an empty gateway value is provided, then the current gateway is cleared.

Restart and Shutdown the InfiniteStorage 11000

Use the information in the following sections to restart or shutdown the controller(s) in your InfiniteStorage 11000 system.

System Restart

RESTART performs a restart on the InfiniteStorage 11000 on which the command is issued. This command prepares the system to be restarted. The system halts all I/O requests and saves the data to the disks before restarting. The restart process may take several minutes to complete.

Note: If cache coherency is enabled, restarting an InfiniteStorage 11000 unit will cause the partner controller to fail the restarting unit. Once the reboot is complete, you will have to heal the InfiniteStorage 11000 unit.

RESTART DELAY=X (where “X” is minutes) delays a restart of a unit between 0 and 255 minutes.

RESTART DUAL restarts both units.

RESTART KILL stops a timed restart that is in progress.

System Shutdown

SHUTDOWN shuts down the InfiniteStorage 11000 unit.

If you need to power down the InfiniteStorage 11000, use **SHUTDOWN** prior to shutting off power. This will cause the InfiniteStorage 11000 controller to flush its cache, abort all format and rebuild operations, and proceed with an orderly shutdown.

All hosts actively using the InfiniteStorage 11000 should be safely shutdown and all users logged out before using this command. The InfiniteStorage 11000 will halt all I/O requests and save the data to the disks.

Note: Use **SHUTDOWN** whenever you power down the InfiniteStorage 11000 for maintenance. **SHUTDOWN** flushes any data left in the cache and prepares the InfiniteStorage 11000 for an orderly shutdown. For couplet InfiniteStorage 11000 configuration, issue **SHUTDOWN** to both InfiniteStorage 11000s.

To perform a hard restart of the unit by cycling the power, use: **SHUTDOWN RESTART=X**, where X is a value between 1 and 1023 seconds before the unit powers up again. If the number is not specified, the default is 15 seconds.

Note: If **SHUTDOWN RESTART** is used in conjunction with the **DUAL** parameter, the restart will only affect the controller unit that it was issued on (not both units).

SHUTDOWN DELAY=X delays a shutdown of a unit between 0 and 255 minutes (where x is minutes delayed).

SHUTDOWN DUAL shutdowns both units.

SHUTDOWN KILL stops a timed shutdown that is in progress.

Setting the System's Date and Time

Valid date settings are between years 2000 and 2104. In dual mode, settings should always be done on Unit 1. Changes will automatically be applied to both units. Settings are automatically adjusted for leap years.

System Date

DATE displays the current system date. To change the system date, type **date mm dd yyyy** <Enter> where **mm** represents the two digit value for month, **dd** represents the two digit value for day, and **yyyy** represents the four digit value for year. For example, to change the date to March 14, 2009, type **date 03 14 2009** <Enter>.

System Time

TIME displays the current system time. To change the system time, type **time hh:mm:ss** <Enter> where **hh** represents the two digit value for hour (00 to 24), **mm** represents the two digit value for minutes, and **ss** represents the two digit value for seconds. For example, to change the system time to 2:15:32 p.m., type **time 14:15:32** <Enter>

Note: 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.

Saving the InfiniteStorage 11000's Configuration

The **SAVE** command can be used to save the system configuration to non-volatile memory (see Figure 3-21 for an example).

```
11000 [1]: save
Saving system parameters. Done.
```

Figure 3-21 Saving the System Parameters Example Screen

Backup copies of the system configuration are also saved on the disks. The system will automatically save and update the backup copies when changes are made to the system configuration or status.

The **SAVE STATUS** command, in addition to saving the parameter blocks to non-volatile memory and on the disks, displays the current status of the system parameters (see Figure 3-22).

```
11000[2]: save status

                System Parameters Status
PB
Number  Revision    Updates    Last Update
-----  -
0       202          354       16:53:37 09/12/2008
1       202        16995     19:33:31 09/10/2008
2       201           11       15:31:54 08/19/2008
3       201           11       15:31:54 08/19/2008
```

Figure 3-22 Save Status CLI Command Screen Example

Normally, the system must determine which copy of the parameter blocks is more recent, the one on the disks or the internal copy. When the system reboots, it will load the more recent copy.

Restoring the System's Default Configuration

The **DEFAULTS** command may be used to restore the system to its default configuration.



Warning

The DEFAULT command will delete all LUN configuration and data unconditionally. Do not issue this command without guidance from your service or support organization.

The system will halt all I/O requests, delete all the LUNs and restore all the parameters back to their default values. This is a **destructive** operation which will delete all the data stored in the system.

The system will ask if you want to erase all the configuration information stored on the disks. This will prevent the system from retrieving the backup copies of the configuration settings from the disks after the system is restarted. After the default settings have been loaded, the system will ask if you want to begin reconfiguration by scanning for the disks. New LUNs can be created after the disks have been added back to the system.

LUN Management

The InfiniteStorage 11000 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 (LUNs) 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.

Configuring the Storage Array

The storage array may consist of up to 120 tiers, depending on individual disk enclosure's numbering scheme. The tiers can be combined, used individually, or split into multiple LUNs. A LUN can be as small as part of a tier or as big as the whole system. LUNs can be shared or dedicated to individual users. Up to 1024 LUNs are supported in total. LUNs are "owned" by the InfiniteStorage 11000 via which they are created.

You can add and remove LUNs without affecting system operations. Use the **LUN** command to display the current Logical Unit Status (see Figure 3-23 on page 68).

Note: In dual mode, LUNs will be "owned" by the InfiniteStorage 11000 unit on which they are created. Hosts will only see the LUNs on the InfiniteStorage 11000 to which they are connected, unless cache coherency is used.

```

11000[2]: lun

                Logical Unit Status

LUN  Label          Owner  Status          Capacity  Block
      (Mbytes)      Size  Tiers Tier list
-----
  0                               10000  512   1  1
  1                               10000  512   1  2
  2                               10000  512   1  3
  3                               10000  512   1  4
  4 AUTO_LUN4         1     Degraded        1000    512   1  3
  5 AUTO_LUN5         1     Ready           1000    512   1  4
  6                               10000  512   1  2
  7 AUTO_LUN7         1     Ready           1000    512   1  4
  8 AUTO_LUN8         1     Ready           1000    512   1  4
  9 AUTO_LUN9         1     Ready           1000    512   1  4
 10 AUTO_LUN10        1     Ready           1000    512   1  4
 11 AUTO_LUN11        1     Ready           1000    512   1  4
 12 AUTO_LUN12        1     Ready           1000    512   1  4
 13 AUTO_LUN13        1     Ready           1000    512   1  4

System verify extent: 16 Mbytes
System verify delay: 30

System Capacity 4485536 Mbytes, 4426536 Mbytes available.

```

Figure 3-23 LUN CLI Command Screen Example

Creating a LUN

The system can support up to 1024 LUNs. The required LUN information includes:

- **Capacity** (in MBytes) - default is to use all available capacity
- **Number of tiers** - default is to use all tiers
- **Block size** (in Bytes) - default is 512Bytes
- **Label** - may contain up to 16 characters

Use one of the followings commands to add a LUN to the system:

- **LUN ADD[=[x]]** adds a **64-bit LUN** to the system. The system prompts you for all the necessary information to create the LUN and indicates if the LUN was successfully added to the system. The LUN to be added can be specified by 'x', where 'x' is in the range 0..1023.

- **LUN [ADD64[=x]]** adds a **64-bit LUN** to the system. The system prompts you for all the necessary information to create the LUN and indicates if the LUN was successfully added to the system. The LUN to be added can be specified by 'x', where 'x' is in the range 0..1023.
- **LUN [ADD32[=x]]** adds a **32-bit LUN** to the system. The system prompts you for all the necessary information to create the LUN and indicates if the LUN was successfully added to the system. The LUN to be added can be specified by 'x', where 'x' is in the range 0..1023. The maximum capacity of a 32-bit LUN or LUN group is limited to 0xFFFF0000 host blocks (see Table 3-5).

Table 3-5 32-bit LUN Block Size/Max Capacity

Block Size	Maximum Capacity
512	2097120 MB
1024	4194240 MB
2048	8388480 MB
4096	16776960 MB

Formatting a LUN

A LUN must be formatted before it can be used.

To format a LUN, use **LUN FORMAT**. Specify the LUN <0..1023> when prompted. This performs a destructive initialization on the specified LUN by over-writing all the data on the LUN with zeroes.

The rate of format can be adjusted using the **DELAY** and **EXTENT** parameters of the **LUN** command.

Interrupting a LUN Format Operation

If you need to interrupt a format operation, for any reason, use these commands:

- **LUN PAUSE** pauses the current format operations.
- **LUN RESUME** releases the paused format operations.
- **LUN STOP** aborts all the current format operations.

Changing a LUN Label

To change the label of a LUN:

1. Type **lun label** <Enter>.
2. Select the LUN to change <**0..1023**> and press <Enter>.
3. Type in the new label and press <Enter>. A LUN label may contain up to 16 characters (see Figure 3-24).

```
11000[2]: lun label
Enter the LUN (0..1023) to label, or 'e' to escape:
0
Enter a new label (up to 16 characters) for LUN 0 :
UserGuide
```

Logical Unit Status							
LUN	Label	Owner	Status	Capacity (Mbytes)	Block Size	Tiers	Tier list
0	UserGuide	1	Ready	10000	512	1	1
1		1	Ready	10000	512	1	2
2		1	Degraded	10000	512	1	3
3		1	Ready	10000	512	1	4
4	AUTO_LUN4	1	Degraded	1000	512	1	3
5	AUTO_LUN5	1	Ready	1000	512	1	4
6		1	Ready	10000	512	1	2
7	AUTO_LUN7	1	Ready	1000	512	1	4
8	AUTO_LUN8	1	Ready	1000	512	1	4
9	AUTO_LUN9	1	Ready	1000	512	1	4
10	AUTO_LUN10	1	Ready	1000	512	1	4
11	AUTO_LUN11	1	Ready	1000	512	1	4
12	AUTO_LUN12	1	Ready	1000	512	1	4
13	AUTO_LUN13	1	Ready	1000	512	1	4

```
System verify extent: 16 Mbytes
System verify delay: 30

System Capacity 4485536 Mbytes, 4426536 Mbytes available.
```

Figure 3-24 LUN Label CLI Command Screen Example

Moving a LUN (Dual Mode Only)

To change the ownership of a LUN from one InfiniteStorage 11000 controller to its partner (when the units are in dual mode), Type **lun move=x** <Enter> where **x** is the Logical Unit number <0..1023>

(see Figure 3-25). If a LUN is on a tier that is shared by other LUNs, the controller will prompt and then move the other dependent LUNs as well.

```
11000 [1]: lun move=0

LUN 0 is owned by this controller.

Do you want to move ownership to the OTHER controller? (y/n):
```

Figure 3-25 Moving a LUN's Ownership

Deleting a LUN

LUN DEL=x (where **x** is the LUN <0..1023>) deletes a LUN from the system. You can only delete a LUN that is owned by the InfiniteStorage 11000 unit onto which you are logged.

SCSI Reservations

LUN RELEASE=x allows you to release all SCSI reservations on a LUN. The command **LUN RESERVATIONS** can be used to view the current SCSI reservations on all of the LUNs in the system. The LUN to be released can be specified by “**x**” where “**x**” is in the range <0..1023>.

LUN START lets you start all the LUNs that have been stopped by a SCSI START/STOP request. This parameter is not related to the **LUN STOP** command.

Automatic Drive Rebuild

The InfiniteStorage 11000's automatic drive failure recovery procedures ensure that absolute data integrity is maintained while operating in degraded mode. In the event of a drive failure, the InfiniteStorage 11000 will automatically initiate a drive rebuild using a spare drive if the “autorebuild” function has been enabled. Use the **TIER** command to display the current setting (see Figure 3-26 on page 72). The rebuild operation can take up to several hours to complete, depending on the size of the disk and rate of rebuild.

```

11000 [1]: tier
                                Tier Status
                                Capacity  Space Available
                                (Mbytes)  (Mbytes)
Tier Owner
-----
  1      280012      280012      ABCDEFGHPS
  2      280012      280012      ABCDEFGHPS
  3      280012      280012      ABCDEFGHPS

Automatic disk rebuilding is Enabled
System rebuild extent: 32 Mbytes
System rebuild delay: 60

System Capacity 840036 Mbytes, 840036 Mbytes available.

```

Figure 3-26 Tier Status and Automatic Disk Rebuilding Parameters Screen

TIER AUTOREBUILD=ON|OFF enables/disables the automatic disk rebuild function. A disk will only be replaced by a spare disk if it fails and Autorebuild is **ON** (ON being the default setting). This function should always be enabled so that data can be reconstructed on the spare drive when a drive failure occurs. After the failed drive is replaced, data will be automatically copied from the spare drive to the replacement drive.

Manual Drive Rebuild

DISK REBUILD=<tier><channel> initiates a rebuild on a specific drive. This operation will reconstruct data on the replacement drive and restore a degraded LUN to healthy status.

Drive Rebuild Verify

DISK REBUILDVERIFY=ON|OFF determines if the system will send **SCSI Write with Verify** commands to the disks when rebuilding failed disks. This feature is used to guarantee that the data on the disks is rebuilt correctly. Default is OFF. This feature will increase the time it takes for rebuilds to complete.

Manual Drive Replace

To replace the specified failed drive with a spare drive, enter: **DISK REPLACE=<tier><channel>**

A Replace operation is used to temporarily replace a failed disk with a healthy spare disk.

Interrupting a Rebuild Operation

To interrupt a Rebuild operation, use these commands:

- **TIER PAUSE** pauses the current rebuild operations.
- **TIER RESUME** releases the paused rebuild operations.
- **TIER STOP** aborts all the current rebuild operations.

SMART Disk Monitoring and Test

Self-Monitoring, Analysis, and Reporting Technology (SMART) is a technology designed for monitoring disk drives. It provides an interface for detecting and reporting various indicators of drive reliability. SMART support has been added to the InfiniteStorage 11000 in order to detect drives that are experiencing symptoms of failure and to initiate preventive steps to deal with these drives.

Use the SMART command to identify failing drives before they fail.

SMART configuration is stored in the parameter blocks and will be restored when the system is restarted.

SMART, if enabled, generates check conditions (Recovered Error Check condition) if any of the pre-programmed threshold limits have been exceeded on the disk drive. Currently on the InfiniteStorage 11000, on a SMART threshold exceeded condition (referred to as SMART TRIP), the disk drive is failed.

SMART Support

SMART support on the InfiniteStorage 11000 consists of two functionalities: monitoring and self-testing:

- SMART monitoring can be used on both SCSI and SATA drives.
- SMART self-testing is limited to SCSI drives.

Monitoring SMART Support

At the InfiniteStorage 11000 Command Line Interface (CLI) the **SMART** command is used to check the status of SMART monitoring and display various data available from the disk. A

SMART command without an argument reports on the current SMART monitoring and temperature warning status (enabled or disabled).

SMART Enable and Disable

The **SMART ENABLE** and **SMART DISABLE** commands are used to enable or disable SMART monitoring on the InfiniteStorage 11000 disk drives. When these commands are executed, the Information Exception Mode page for all disk drives is modified to enable or disable SMART monitoring. Thus, once an “enable” or “disable” command is executed, all the disk drives connected to the InfiniteStorage 11000 are affected.

Enabling SMART monitoring configures the method used to report SMART monitoring threshold events. When there is a SMART Threshold Exceeded condition on a drive(s), a recovered error (a SMART trip) is reported to the InfiniteStorage 11000. A SMART Threshold Exceeded message initiates the steps the InfiniteStorage 11000 takes to resolve the situation (for example: attempts to spare the disk reporting a SMART trip). The SMART trip, in the form of a “recovered error” message, will be repeatedly issued by the disk drive until the condition is resolved.

Status and Show

The **SMART DISKSTATUS** command displays disks that have SMART trips but have not been failed and/or replaced since their last reboot. **SMART STATUS=ALL** displays all disks that have had SMART trips since the last reboot (inclusive of spared-out and failed drives). The information displayed depends upon the type of drive (SCSI or SATA) and may contain temperature information, overall SMART status and results from self-tests.

SMART Self-test Support

SMART Self-testing is currently supported on SCSI drives only.

Self-test commands:

- initiate background self-tests and default vendor-specific self-tests (if supported)
- retrieve self-test logs
- display self-test results.

SMART self-test command arguments are described below.

Start Self-test

The **SMART SELFTEST=tc** command initiates a self-test on a drive(s), specified by tc (tier and channel). The **SMART SELFTEST=ALL** command initiates a self-test on all FC drives in the system.

A self-test type needs to be selected after issuing the command, or can be entered at the command line. The self-test launches and the event is logged on the Self-Test Log page of the drive(s). There are 5 tests to choose from:

- **Default test**
- **Background short test**
- **Background long test**
- **Foreground short test**
- **Foreground long test.**

All the tests are supported on the fiber channel drives, while only the Background short test is supported on the SATA drives.

Stop Self-Test

The **SMART ABORTSELFTEST=tc** command stops a previously initiated self-test on a drive(s) specified by tc (tier and channel). The **SMART ABORTSELFTEST=ALL** command stops a previously initiated self-test on a all drives. The test stops and the event is logged in the Self-Test Log page of the drive(s).

Self-test Status

The **SMART LOG=tc** command displays the self-test log of a drive(s) specified by tc (tier and channel). **SMART LOG=tc** retrieves the Self-Test Log page from the drive(s) and displays the most recent tests as well as their results and times.

SMART and Proactive Sparing

Proactive sparing feature is used to provide the user an ability to replace disks from a healthy state instead of having to fail the disk in order to perform the replace. Proactive sparing is used for disk performance errors and medium errors. Proactive sparing has two modes of operation that the end user can choose from:

- the manually triggered CLI command mode
- and the SMART automatic sparing mode.

The SMART automatic mode couples the automatic sparing mode with SMART information provided by the disks to enable proactive sparing. Since the controller's firmware can read SMART data, the user can choose to have SMART threshold crossings as defined internal to the drive, to trigger the drive to be queued for proactive sparing. The CLI command to support this feature is **TIER SMARTREPLACE= [on|off]**. This command is a sub-qualifier to the **TIER AUTOREBUILD** setting. That is, if **TIER AUTOREBUILD** is **off**, then the **SMARTREPLACE** setting is irrelevant.

The setup required for SMART and Proactive Sparing to work together is:

- TIER AUTOREBUILD=ON
- TIER SMARTREPLACE=ON
- SMART ENABLE

SMART and Drive Sparing

8+2 Channel Sparing

Channel spares are used to replace failed disks on 8+2. Channel spares are provisioned by the system administrator by declaring a tier to be used as spare disks with the CLI command, **TIER SPARE='t'**. All disks installed on this tier ('t') will be available to replace 8+2 disks on tiers with LUNs. Channel spare disks are used exclusively for replacement of 8+2 disks

The management of 8+2 spares is done through the same mechanism as 8+1 spares although the display is a little different due to the number of spares on the single tier:

- The **TIER AUTOREBUILD** parameter is applicable on disk failures.
- The replacement retries follow the same rules.
- The search for spares is done using the same mechanism; however, the channel sparing search for spares is limited to the same channel as the failed disk (that is, if 1A fails, ONLY channel A will be searched for a suitable spare to replace 1A).

Channel spare disks are subject to spare diagnostics when they are not in use as a replacement disk. All commands that involve the use of a spare disk will apply to the display/usage of the channel spares.

Before a tier may be designated as a spare tier, these requirements must be met:

- The tier has no LUNs;
- And if the tier is an 8+1, that the S channel is not in use as a replacement for another disk.

Management of the spares is facilitated by the following CLI commands to simplify the display and maintain the same look and feel to the CLI output:

- **TIER SPARE='t'** (converts the identified tier. 't', to a spare tier)
- **TIER CONFIG=SPARE** (displays only spare tier configurations)
- **TIER CONFIG=ALL** (displays all usable tiers, spare and non-spare)

The setup required for SMART and Channel Sparing to work together is:

- TIER SPARE='t'
- TIER AUTOREPLACE=ON
- TIER SMARTREPLACE=ON
- SMART ENABLE

8+1+1 Sparing

The concept of channel sparing is NOT applicable to 8+1 disks. 8+1 tier S channel disks are used exclusively for replacement of 8+1 disks.

The setup required for SMART and 8+1 Sparing to work together is:

- TIER AUTOREPLACE=ON
- TIER SMARTREPLACE=ON
- SMART ENABLE

Couplet Controller Configuration (Cache/Non-Cache Coherent)

There are two primary couplet InfiniteStorage 11000 configurations: **cache coherent** and **non-cache coherent**.

The **DUAL** command displays information about couplet system configuration (see Figure 3-27 on page 78).

```

11000[1]: dual

                        Dual Unit Configuration

                        Unit 1                          Unit 2
-----
Label                   11000[1]                      11000[2]
Status                   Healthy                        Healthy

Dual communication:     established.
Ethernet communication: established.
Cache coherency:       not enabled.
Cache coherency timeout: 0

11000[1]:

```

Figure 3-27 DUAL CLI Command Screen

Cache Coherent

In this configuration, each InfiniteStorage 11000 controller can access all LUNs. The couplet InfiniteStorage 11000 communication occurs over the internal **UART** and private Ethernet. If the controllers detect an Ethernet failure, controller 2 will be failed. (This means that an external event can cause an InfiniteStorage 11000 controller to fail even though it may be perfectly fine.) Therefore, it is mandatory that the InfiniteStorage 11000 Ethernet resides on a private Ethernet segment.

Note: Data cache is not copied from one InfiniteStorage 11000 controller to another. If a controller fails, all “dirty” data in cache will be lost. Thus if power failures are a concern, writeback cache should be disabled.

Non-Cache Coherent

In this configuration, the couplet InfiniteStorage 11000 controller communication occurs over the internal UART. Each InfiniteStorage 11000 owns LUNs and tiers. Spare drives are “owned” by individual InfiniteStorage 11000 controller units, according to tier ownership.

In healthy situations, one InfiniteStorage 11000 controller cannot access LUNs or tiers owned by the other. However, if the other controller is failed, the healthy controller will have access to all LUNs and tiers.

Users, via mapping, can be assigned any combination of LUNs. In a healthy environment, users will only see LUNs owned by the InfiniteStorage 11000 controller to which they are connected.

For example, a user is given access to internal LUNs 5, 6, and 7, which are mapped to external LUNs 0, 1, and 2, respectively. Controller 1 owns LUNs 0 and 1 while controller 2 owns LUN 2. The user is physically connected to controller 1, thus, they will only see LUNs 0 and 1. The user will not be able to access LUN 2. If the user was physically connected to controller Unit 2, the reverse would be true: only LUN 2 would be accessible. When a controller fails, the user will be given access to all mapped LUNs regardless of the physical connection.

Data cache is not copied from one controller to another. If one fails, all “dirty” data in cache will be lost. Thus if power failures are a concern, writeback cache should be disabled.

DUAL COHERENCY=ON|OFF enables/disables the cache coherency function. Default is dual coherency disabled which is the non-cache coherent configuration.

DUAL TIMEOUT=X allows you to set the cache coherency timeout for cache node requests in seconds. Valid range is <0...255>. Default is zero (0) seconds. The timeout value should be less than the host timeout value (**HOST TIMEOUT=X**). A timeout value of **0** allows for only one retry.

Note: In dual mode, LUNs will be “owned” by the controller unit on which they are created. Hosts will only see the LUNs on the controller to which they are connected, unless cache coherency is enabled.

Fail / Restore the Other Controller Unit in the Couplet Pair

To fail the other InfiniteStorage 11000 unit in the system (for example, in order to perform maintenance), type **dual fail** <Enter>. The healthy InfiniteStorage 11000 unit will take ownership of all the LUNs/tiers from the failed controller unit.

To restore the other controller unit in the system to healthy status after failure recovery, type **dual heal** <Enter>. Ownership of LUNs/tiers are transferred back to the formerly failed InfiniteStorage 11000 unit.

Labeling the InfiniteStorage 11000 Controller Unit(s)

You may change the label assigned to each InfiniteStorage 11000 unit. This allows you to uniquely identify each unit in the system. The CLI prompt for each controller is built by adding a colon (:) and a space at the end of the label. Each InfiniteStorage 11000 controller can have a label up to 31 characters long.

To change the label ():

1. Type: **dual label=1|2** <Enter>.
2. Select which unit you want to rename (see Figure 3-28).
3. When prompted, type in the new label for the selected unit. The new name is displayed.

Note: If you type **DEFAULT** for the new label, the label is restored to its default setting.

```
11000[2]: dual label

Enter the number of the unit you wish to rename.

    LABEL=1 for unit 1, 11000[1]
    LABEL=2 for unit 2, 11000[2]

Unit: 2

Enter a new label for unit 2, or DEFAULT to return to the default label.
Up to 31 characters are permitted.

Current unit name: 11000[2]
New unit name: SYSTEM[2]

                Dual Unit Configuration
                Unit 1                      Unit 2
-----
Label           11000[1]                   SYSTEM[2]
Status          Not connected              Healthy

Dual communication:  not established.
Ethernet communication: failed!
Cache coherency:    not established.
Cache coherency timeout: 0

SYSTEM[2]:
```

Figure 3-28 Dual Label CLI Command Screen Example

Singlet

The **DUAL SINGLET** command sets the system in the singlet mode. System recognizes only Unit 1. This command:

- disables cache coherency
- heals unit 1 if it is failed
- fails unit 2 before attempting to remove it.

To set the system in singlet mode not couplet mode, type: **dual singlet**_{<Enter>}.

Note: The system may automatically add unit 2 if it is connected to the system. Therefore, we advise you to power off and remove unit 2 from the system after the **Dual Singlet** command is completed.

Performance Management

The InfiniteStorage 11000 controller optimizes performance operations due to its extensive monitoring and reporting capability.

Optimizing I/O Request Patterns

The InfiniteStorage 11000 controller manages pre-fetch and cache efficiency through the LUN.

Display Current Cache Settings

The **CACHE** command displays the current cache settings for each LUN in the system (see Figure 3-29).

```
11000[2]: cache

                        Current Cache settings

  LUN   Write   Maximum   MF   Prefetch
        Caching Prefetch  Bit  Ceiling
-----
   0   Enabled   x       1   On   65535
   1   Enabled   x       1   On   65535
   2   Enabled   x       1   On   65535
   3   Enabled   x       1   On   65535
   4   Enabled   x       1   On   65535
   5   Enabled   x       1   On   65535
   6   Enabled   x       1   On   65535
   7   Enabled   x       1   On   65535
   8   Enabled   x       1   On   65535
   9   Enabled   x       1   On   65535
  10   Enabled   x       1   On   65535
  11   Enabled   x       1   On   65535
  12   Enabled   x       1   On   65535
  13   Enabled   x       1   On   65535

Writeback Limit: 75%

                2560.0 Mbytes of Cache Installed
                (2048 Segments of 1024 Kbytes)

11000[2]:
```

Figure 3-29 Cache CLI Command Screen Example

You can use the **LUN=x** option to specify which LUN to change. If no LUN is specified, changes will be applied to all the LUNs. Valid LUN values are 0 to 1023. The default value will apply changes to all LUNs.

Cache Segment Size

A large cache segment size may give better performance for large I/O requests and a small cache segment size may give better performance for small I/O requests. For optimal performance, the cache segment size should be larger than the average host I/O request size. You may use the **STATS LENGTH** command to determine the average host I/O request size. The cache segment size should not be changed during heavy I/O conditions because the system will temporarily halt all I/O requests while the changes are taking effect.

Use the **CACHE SIZE=x** command to set the cache segment size for the specified LUN in kilobytes (kbs). Valid segment sizes are 128, 256, 512, 1024, and 2048 kilobytes (kbs). The default value is 1024. This command should not be issued under heavy I/O conditions because the system will momentarily halt all I/O requests while the changes are taking effect.

Writeback Cache Settings

Writeback caching allows the system to increase the performance of handling write I/O requests by storing the data in cache and saving the data to the disks at a later time.

CACHE WRITEBACK=ON|OFF enables or disables writeback caching for the specified LUN. Default setting is ON.

CACHE WRITELIMIT=x specifies the maximum percentage of the cache that can be used for writeback caching. The system will force all writeback requests to be flushed to the disks immediately if the percentage of writeback data in the cache exceeds this value. Valid range is <0...100>. Default value is 75.

Prefetch Settings

When the system receives a request, it can read more data than has been requested. **PREFETCH** tells the system how much data to look ahead. This will improve performance if your system needs to perform sequential reads. For random I/O applications, however, use the smallest prefetch value.

CACHE PREFETCH=x sets the prefetch that will occur on read commands for the specified LUN. Valid range is 0 to 65535. Default setting is 1.

If the **MF** (Multiplication Factor) parameter is **OFF**, the system will only prefetch the number of blocks specified by **PREFETCH** after every read command. If the MF parameter is **ON**, then the system will multiply the transfer length of the command by the prefetch value to determine how much data will be prefetched. A prefetch value of less than 8 is recommended when the **MF** parameter is **ON**.

CACHE MF=ON|OFF enables/disables the MF bit on the specified LUN. Default is **ON**.

The **Maximum Prefetch Ceiling** parameter sets the maximum prefetch ceiling in blocks for prefetches on read commands. It sets an upper limit on prefetching when the MF parameter is **ON**. The system will automatically limit the amount of prefetching if the system is running low on resources.

CACHE MAX=x (where **X** is a range from 0...65535) sets the maximum prefetch ceiling in blocks for prefetches on Read commands for the specified LUN. Valid range is 0 to 65535. Default setting is 65535.

Cache Settings Reset

CACHE DEFAULTS loads the default settings for all of the cache parameters for the specified LUNs.

Disk Configuration Settings

The **DISK** command displays the current disk configuration settings. The writeback cache and disk timeout settings can be configured manually.

DISK TIMEOUT=x sets the disk timeout for an I/O request in seconds. Valid range is 1 to 512 seconds. Default setting is 68 seconds.

DISK CMD_TIMEOUT=x sets the Retry Disk timeout (in seconds) for an I/O request. The retry timeout value indicates the maximum amount of time that is allotted to receive a reply for each retry of an I/O request. If the I/O request does not complete within this time, it is aborted and potentially retried: if there is still time remaining in the overall disk timeout to allow for another retry, it is retried; if not, it completes with an error status.

Note: The **DISK CMD_TIMEOUT** value must be smaller than or equal to **DISK TIMEOUT**. Valid range is 1 to 512 seconds.

Audio/Visual Settings of the System

The **audio** and **visual (AV)** settings of the system and the disks can be tuned to provide better performance and a lower latency. The writeback and prefetch settings for each LUN are changed with the **CACHE** command.

The **AV** command displays information about the audio/visual settings of the system (see Figure 3-30 for an example).

```
11000[2]: av

          Current LUN Audio/Visual settings

          Read      Write      Maximum
          LUN  Label      FastAV Continuous Caching  Prefetch
-----
          0 UserGuide      Off      Off      On      x 1
          1                Off      Off      On      x 1
          2                Off      Off      On      x 1
          3                Off      Off      On      x 1
          4 AUTO_LUN4      Off      Off      On      x 1
          5 AUTO_LUN5      Off      Off      On      x 1
          6                Off      Off      On      x 1
          7 AUTO_LUN7      Off      Off      On      x 1
          8 AUTO_LUN8      Off      Off      On      x 1
          9 AUTO_LUN9      Off      Off      On      x 1
         10 AUTO_LUN10     Off      Off      On      x 1
         11 AUTO_LUN11     Off      Off      On      x 1
         12 AUTO_LUN12     Off      Off      On      x 1
         13 AUTO_LUN13     Off      Off      On      x 1

Disk Audio/Visual settings are: Disabled (Using disk defaults)
Recovery Time Limit:          65535

Ordered Tag Count:           0
Unit Attention:              Enabled
FASTAV Timeout:              0
RC Timeout:                  0
Fail Check Condition:        Disabled

11000[2]:
```

Figure 3-30 AV CLI Command Screen Example

AV FASTAV=ON|OFF enables/disables the disk fast audio/video read options for streaming data. When enabled, the system will start the data transfer for read operations before all of the disk commands have finished. This feature reduces the latency for read operations but the system will

be unable to check the integrity of the data. This parameter is saved on a LUN by LUN basis. Use LUN=X command to change the settings for a single LUN. Default setting is OFF.

Important: When **FASTAV** mode is enabled, the controller no longer checks data in real-time.

Changing the disk parameters can adversely affect the I/O operation of the system. This parameter should only be adjusted when the system is idle. Default setting is OFF.

FASTAVTIMEOUT=x sets the timeout before the FASTAV option activates on a host read command. The FASTAV mechanism is not used until the host command takes longer than the timeout value. A value of zero indicates that the system starts the data transfer as soon as a minimum number of drives are ready. This value is in 100 millisecond increments. The range for “x” is 0 to 255. The default is 50.

ORDEREDQUEUE=x enables the use of ordered tags when communicating with the drives. The value “x” indicates the number of disk commands that can be sent before an ordered tag must be sent to the disks. Valid range is 0 to 255. Default is 0.

UA=ON enables the initial Unit Attention condition when an initiator logs into the system; the system reports a Unit Attention condition on the first SCSI command after the initiator logs in. Default is ON.

UA=OFF disables the initial Unit Attention condition when an initiator logs into the system; the system automatically clears the unit attention condition when an initiator logs in.

RC=ON|OFF enables the **Read Continuous (RC)** option for Audio/Video streaming data; the system starts the data transfer for read operations after RCTIMEOUT is reached, even if the disks commands have not finished. Use this to reduce the latency for read operations in Audio/Visual environments where latency is more important than data integrity. This parameter is saved on a per-LUN basis. Use in combination with the LUN=x parameter to change the settings for a single LUN. Enabling this feature automatically enables FASTAV.

Caution: This feature allows the system to return invalid data to the initiator.

RCTIMEOUT=x Default setting: disables the Read Continuous option for Audio/Video streaming data. Note: This parameter is saved on a per-LUN basis. Use in combination with the LUN=x parameter to change the settings for a single LUN.

LUN=x sets the host command timeout for the Read Continuous option for Audio/Video streaming data. Set to 0 to disable the Read Continuous feature in the system. This value is in 100 millisecond increments. The range for 'x' is 0 to 255. The default is 8.

FAILCC=ON instructs the host ports to report a check condition for all SCSI commands when the unit is in a failed state. This command should only be used in AV environments when a check condition is required instead of taking the unit off the loop.

FAILCC=OFF This is the default setting. Host ports will NOT report a check condition for all SCSI commands when the unit is in a failed state.

Locking LUN in Cache

Locking a LUN in data cache will keep all of the data for the LUN in the cache for faster access. Once a LUN is locked, the data that is gathered to service read and write commands will stay permanently in the cache. The InfiniteStorage 11000 controller will continue to fill up the cache until 50% of the total cache is filled with data from locked LUNs, while the other 50% of the cache is reserved to service I/O for unlocked LUNs (reference Figure 3-31 through Figure 3-34).

Initial Cache

50% of Data Cache used to service Unlocked LUNs	50% of Data Cache used to service Locked LUNs
---	---

Figure 3-31 Initial Cache Example

For example, when a host issues a read command for data from LUN 1 that has been locked in cache, the following will occur:

- The controller reads data from disks, locks data in cache, and sends data to host
- Any reads of the same data will be serviced from cache, which provides faster access than reading from disks.

Cache allocation after I/O completes

Unlocked LUN data	Unallocated cache*	Data for LUN 1
-------------------	--------------------	----------------

* Unallocated cache can be used for unlocked LUNs, or locked LUNs. Once cache has been allocated to a locked LUN, however, it cannot be used by an unlocked LUN.

Figure 3-32 Cache Allocation Example

Once the size of the locked LUNs exceeds 50% of the total cache, the InfiniteStorage 11000 controller must create cache space to process a new I/O by removing older data from the locked portion of cache. The **Least Recently Used (LRU)** algorithm is used to determine which locked data to remove from cache.

For example, LUNs 0 to 3 are locked in cache and all 50% of the total cache has been filled by data from LUN 0, 1, and 2.

Initial Cache

Unlocked LUN data	Data for LUN 0	Data for LUN 1	Data for LUN 2
-------------------	----------------	----------------	----------------

Figure 3-33 Initial Cache Example

When a host issues a read command for data from LUN 3, the following will occur:

- The InfiniteStorage 11000 controller unit determines which data to remove from the locked portion of cache, using the LRU algorithm. The LRU algorithm is thus: If LUN 0 has not been accessed for 1 hour, LUN 1 has not been accessed for 30 minutes, and LUN 2 has not been accessed for 2 minutes, then LUN 0's data will be removed from cache because it is the least recently used data.
- The controller reads data from disks, locks data in cache, sends data to host.
- Any reads of same data will be serviced from cache (until data is removed from cache due to its being the least recently used data).

Cache allocation after I/O completes

Unlocked LUN data	Data for LUN 3	Data for LUN 1	Data for LUN 2
-------------------	----------------	----------------	----------------

Figure 3-34 Cache Allocated After I/O Example

Locking / Unlocking a LUN

To lock a LUN in the data cache, type **LUN LOCK=X**_{<Enter>} where “X” is the Logical Unit number <0..1023> (see Figure 3-35). **LUN UNLOCK=x** unlocks a LUN and releases its cache locked by the LUN.

```
11000 [1]: lun lock=0
```

Logical Unit Status							
LUN	Label	Owner	Status	Capacity (Mbytes)	Block Size	Tiers	Tier List
0		1	Cache Locked	10002	512	3	1 2 3
1		1	Ready	10002	512	3	1 2 3
2		1	Ready	10002	512	3	1 2 3
3		1	Ready	10002	512	3	1 2 3

System Capacity 277810 Mbytes, 237802 Mbytes available.

Figure 3-35 LUN Lock Status Example Screen

System Performance Statistics

The InfiniteStorage 11000 controller monitors pre-fetch and cache efficiency, request distribution, transaction, and transfer rates by port.

The **STATS** command displays the **Performance Statistics** for the host ports, disk channels, and cache memory (see Figure 3-36 on page 90). It shows the read and write performance of each of the host ports.

```

11000[2]: stats

                System Performance Statistics
Read  MB/s:    All Ports    Port 1    Port 2    Port 3    Port 4
Write MB/s:    0.0         0.0      0.0      0.0      0.0      0.0
Total MB/s:    0.0         0.0      0.0      0.0      0.0      0.0

Read  IO/s:    0          0         0         0         0         0
Write IO/s:    0          0         0         0         0         0
Total IO/s:    0          0         0         0         0         0

Read Hits:     0.0%       0.0%     0.0%     0.0%     0.0%     0.0%
Prefetch Hits: 0.0%       0.0%     0.0%     0.0%     0.0%     0.0%
Prefetches:   0.0%       0.0%     0.0%     0.0%     0.0%     0.0%
Writebacks:   0.0%       0.0%     0.0%     0.0%     0.0%     0.0%
Rebuild MB/s: 0.0         0.0      0.0      0.0      0.0      0.0
Verify MB/s:  0.0         0.0      0.0      0.0      0.0      0.0

                Total      Reads      Writes
Disk IO/s:      0          0          0
Disk MB/s:      0.0        0.0        0.0
Disk Pieces:   24525870  12508051  12017819
BDB Pieces:    0

Cache Write Data:      0.0%
Rebuild/Verify Data:  0.0%   0.0%
Cache Data locked:    0.0%

11000[2]:

```

Figure 3-36 STATS CLI Command Screen

- Read Hits** shows the percentage of Read I/O requests where the data was already in the cache.
- Prefetch Hits** shows the percentage of Read I/O requests where the data was already in the cache due to prefetching
- Prefetches** shows the percentage of host Read I/O requests to the disks due to prefetching.
- Disk Pieces** shows the total number of disk I/O requests from the host ports
- BDB Pieces** is the number of host I/O blocking and deblocking requests.
- Cache Write Data** shows the percentage of the cache which contains writeback data that must be written to the disks.
- Cache Rebuild Data** shows the percentage of the cache in use for rebuild operations.
- Cache Data Lock** shows the percentage of the cache which is locked by the locked LUNs.

The bottom of the screen displays the Read and Write performance of the disks. The system will combine several host I/O requests into a single disk I/O request. The histogram at the lower right shows how often this is occurring for reads and writes.

STATS CLEAR resets all the statistics back to zero.

STATS DELAY displays a histogram of the time it takes for the host and disk I/O requests to complete in 100 msec intervals (see Figure 3-37).

```
11000 [1]: stats delay
```

Command Delay Statistics				
Time seconds	Host Reads	Host Writes	Disk Reads	Disk Writes
0.1	1690087	1446110	281633	253704
0.2	82900	79522	87112	45260
0.3	389	263	13243	7728
0.4	64	77	3319	3149
0.5	12	24	970	1435
0.6	5	7	336	672
0.7	0	8	92	344
0.8	0	4	38	136
0.9	0	3	13	84
1.0	0	9	8	8445
1.1	6	3	4	24
1.2	9	19	2	14
1.3	12	15	1	18
1.4	12	17	0	10
1.5	12	19	0	9
1.6	7	32	0	0
1.7	14	34	0	0
1.8	22	12	0	0
1.9	23	12	0	0
2.0	56	19	0	0
2.1	175	4	0	0
2.2	70	1	0	0

Figure 3-37 Command Delay Statistics Screen Example

STATS HOSTDELAY displays a histogram of the time delay between when the last data transfer is set ready and the host command completes (see Figure 3-38 on page 92). The host ready delay information is shown in 100msec intervals.

```
11000 [1]: stats hostdelay
```

Host Command Ready Delay Statistics

Time seconds	Port 1		Port 2		Port 3		Port 4	
	Reads	Writes	Reads	Writes	Reads	Writes	Reads	Writes
0.1	0	0	0	0	0	0	0	0
0.2	0	0	0	1	0	0	0	0
0.3	0	0	0	1	2	1	0	0
0.4	0	0	1	2	0	2	0	0
0.5	0	0	0	0	0	0	0	0
0.6	0	0	0	1	0	2	0	0
0.7	0	0	0	1	0	2	0	0
0.8	0	0	0	0	0	0	0	0
0.9	0	0	0	0	0	0	0	0
1.0	0	0	0	0	0	2	0	0
1.1	0	0	0	0	0	0	0	0
1.2	0	0	0	0	0	1	0	0
1.3	0	0	0	0	2	1	0	0
1.4	0	0	0	0	2	1	0	0
1.5	0	0	0	0	0	0	0	0
1.6	0	0	0	0	0	0	0	0

Figure 3-38 Host Delay Statistics Screen Example

STATS TIERDELAY=<tier> displays a histogram of the time it takes for the disk I/O request to complete for all the disks in the specified tier (see Figure 3-39). If no tier is specified, all valid tiers will be displayed.

```
11000 [1]: stats tierdelay
```

Tier 1 Delay Statistics

Time seconds	Disk Channels										
	A	B	C	D	E	F	G	H	P	S	
0.1	3407b	33108	339bd	3409f	572c5	34c0d	33640	30603	3391a	7ed5d	
0.2	480f4	4885c	4866a	48190	27b83	47910	484cc	4acc1	48196	21e	
0.3	2ca6	33d8	2def	2c1f	127	2928	324f	3a63	32a7	0	
0.4	d1	1bc	cd	c7	0	c0	185	10f	176	0	
0.5	2c	2b	26	12	0	23	27	33	36	0	
0.6	13	1b	14	12	0	e	13	1d	1d	0	
0.7	13	15	7	a	6	e	15	28	17	0	
:											
:											
1.8	0	0	0	0	0	0	0	0	0	0	
1.9	0	0	0	0	0	0	0	0	0	0	
2.0	0	0	0	0	0	0	0	0	0	0	

Hit enter to continue, ,Äöe,Äö to esc

Figure 3-39 Tier Delay Statistics Screen Example

STATS DISK displays a histogram of the disks in the system that have taken an unusually long time to complete an I/O request (see Figure 3-40). The count is incremented for a disk if that disk takes longer than the other disks to finish an I/O request. This command is used to determine if a disk in the array is slowing down system performance. Normally all the disks in a tier should have similar counts. A disk with a significantly higher count indicates that the disk may be slower than the other disks or have problems.

```
11000 [1]: stats disk
```

Delayed Disk Command Counts										
	A	B	C	D	E	F	G	H	P	S
0	0	0	0	0	0	0	0	0	0	0
2	3C5	392	34D	4DC	37C	361	3BD	3EE	48B	0
3	0	0	0	0	0	0	0	0	0	0
4	421	7F7	37F	396	7DB	3D2	5B6	3C6	55E	0
5	0	0	0	0	0	0	0	0	0	0
6	338	37E	37F	36C	30F	38B	8DF	5D1	58E	0
7	0	0	0	0	0	0	0	0	0	0
8	0	0	0	0	0	0	0	0	0	0
9	3F1	347	6D4	7DD	929	357	3B4	4D4	5FA	0
10	78C	3B3	412	2ED	642	40A	788	33B	43E	0
11	465	3EE	739	34C	2FC	A2F	358	310	382	0
12	0	0	0	0	0	0	0	0	0	0

Disks in the same tier should have similar results.

Figure 3-40 Host Command Offsets Screen Example

STATS DUAL displays the statistics for the dual mode messages (see Figure 3-41).

```

11000[2]: stats dual

Dual Message Statistics
Message          Total          Msgs/sec.
Lock requests           0             0
Release requests       0             0
Lock acknowledges     0             0
Lock releases          0             0

Lock requests received 0             0
Release requests received 0             0
Lock acknowledges received 0             0
Lock releases received 0             0

Total Messages sent    0             0
Total Messages received 0             0

11000 [2]:

```

Figure 3-41 Dual Message Statistics Screen Example

STATS LENGTH displays a histogram of the length of the host I/O requests in 16 kb intervals (see Figure 3-42 on page 94).

Command Length Statistics									
Length Kbytes	Port 1		Port 2		Port 3		Port 4		
	Reads	Writes	Reads	Writes	Reads	Writes	Reads	Writes	
> 0	0	0	0	0	0	0	0	0	0
> 2	0	0	0	0	0	0	0	0	0
> 4	0	0	0	0	0	0	0	0	0
> 8	0	0	0	0	0	0	0	0	0
> 16	0	0	0	0	0	0	0	0	0
> 32	0	0	0	0	0	0	0	0	0
> 48	0	0	0	0	0	0	0	0	0
> 64	0	0	0	0	0	0	0	0	0
> 80	0	0	0	0	0	0	0	0	0
> 96	0	0	0	0	0	0	0	0	0
> 112	0	0	0	0	0	0	0	0	0
> 128	0	0	0	0	0	0	0	0	0
> 144	0	0	0	0	0	0	0	0	0
> 160	0	0	0	0	0	0	0	0	0
> 176	0	0	0	0	0	0	0	0	0
> 192	0	0	0	0	0	0	0	0	0
> 208	0	0	0	0	0	0	0	0	0
> 224	0	0	0	0	0	0	0	0	0
> 240	0	0	0	0	0	0	0	0	0

Figure 3-42 Command Length Statistics Screen Example

STATS OFFSET displays a histogram of the offset of the host I/O requests into the cache segments (see Figure 3-43). Host I/O requests with offsets that are not in the 0x0 column may require blocking/deblocking which can slow down the performance of the system.

```
11000 [1]: stats offset
```

Host Command Offsets								
	x0	x1	x2	x3	x4	x5	x6	x7
0	720943	8	11	5	0	2	0	343AAD2
8	3FE8E9	5	10	1	0	2	0	3486F35
10	42754D	3	6	0	0	4	1	39B0635
18	4AA571	1	4	2	0	6	0	40677A9

Most commands should be in column 0 or 4 for the best performance.

Figure 3-43 Host Command Offsets Screen Example

STATS REPEAT=OFF|MBS|IOS allows you to enable/disable the repeating statistics display where:

- **MBS** displays MB/s,
- **IOS** displays IO/s, and
- **OFF** turns off (both) the repeating displays.

Resources Allocation

Background Format/Rebuild Operations

Format and rebuild operations are background processes; their rates can be adjusted to minimize their impact on system performance. **TIER** displays the current rebuild parameter settings for the system (see Figure 3-44).

```

11000[2]: tier
                                Tier Status
                                Space
Tier Owner  Capacity  Available  Disk
            (Mbytes) (Mbytes)   Status
-----
  1 + 1     1121384   1111384   ABCDEFGHPS  0
  2 + 1     1121384   1101384   ABCDEFGHPS  1 6
  3 + 1     1121384   1110384   ABCDEFGHP.  2 4
  4 + 1     1121384   1103384   ABCDEFGHPS  3 5 7 8 9 10 11 12 13
Disk 3S Failed.

Automatic disk rebuilding:  Enabled
Automatic SMART replacement: Disabled
Maximum active rebuilds:    4
Maximum active verifies:   2
System rebuild extent:     32 Mbytes
System rebuild delay:      30
Journaling:                 Disabled

System Capacity 4485536 Mbytes, 4426536 Mbytes available.

```

Figure 3-44 Tier Status Screen Example

The **TIER DELAY** parameter controls the amount of system wait time before rebuilding the next chunk of data. This parameter slows down the rebuild and format operations so they will not affect the performance of the system. **TIER DELAY=0** will remove many delays so the rebuild and format operations will go as fast as possible, but this could significantly affect the performance of the system.

Note: A delay value less than 1 (<1) is not recommended.

TIER DELAY=x is used to set the system rebuild/format delay. This value is in 100 millisecond increments. The range is 0 to 1000. The default setting is 30 milliseconds.

The **REBUILD EXTENT** parameter determines how much data to rebuild or format at one time. A small **EXTENT** value will slow down the rebuild and format operations so they will not affect the performance of the system. Increasing the **EXTENT** value will allow more data to be rebuilt in a single pass. The recommended setting is to use the default value of 32 MBytes (MBs) and only adjust **DELAY** to match your user load.

TIER EXTENT=X (where X is a value from 1 to 128) sets the system rebuild/format extent in MBs. The range is 1 to 128 MBs. Default is 32MBs.

Background LUN Verify Operations

LUN VERIFY displays the current setting for background verify on all LUNs (see Figure 3-45).

```
11000[2]: lun verify

                                LUNs Currently Being Verified

LUN   Label      Continuous  Progress  Passes  Last Verify Completion Time
-----
0 UserGuide      Yes          1266     10:46:28 08/15/2008
1          Yes          358      10:45:49 08/15/2008
2          Yes         40045    13:31:20 09/10/2008
3          Yes         41090    16:30:06 09/10/2008
4 AUTO_LUN4      Yes          1        13:31:24 09/10/2008
5 AUTO_LUN5      No           0
6          Yes          358      10:45:49 08/15/2008
7 AUTO_LUN7      No           0
8 AUTO_LUN8      No           0
9 AUTO_LUN9      No           0
10 AUTO_LUN10    No           0
11 AUTO_LUN11    No           0
12 AUTO_LUN12    No           0
13 AUTO_LUN13    Yes          1        16:30:01 09/10/2008

There are NO LUNs being verified currently.

System verify extent: 16 Mbytes
System verify delay: 30

11000[2]:
```

Figure 3-45 LUN Verify CLI Command Screen Example

LUN VERIFY=X turns on background verify for LUN X, where X is a Logical Unit <0..1023>.

LUN VERIFY=ON|OFF prompts you for a list of LUNs where the background verify will be turned either ON or OFF.

LUN VERIFY=ON will both turn on the background verify for the specified LUN(s), as well as start up the verify operation(s).

LUN VERIFY=OFF only turns off the **Background Verify** setting for the specified LUN(s). Therefore, any Verifys that are already active on the LUN(s) will not terminate until after the completion of that Verify's current iteration. To stop all verify operations immediately, use the **LUN STOP** command.

Note: It is recommended that you run LUN VERIFY in continuous mode, since it can help increase disk reliability.

LUN DELAY=X sets the system Verify Delay value to **X**, where **x** is a value from 0 to 1000. The Verify Delay value determines how long a verify operation will pause after it reaches the verify extent. This parameter slows down the verify operation so that it will not affect the performance of the system (except in the case where **X** is set to 0, as described below).

DELAY=X will remove all delays so that the verify operation will go as fast as possible; however, this will slow down the performance of the system. This value is in 100 millisecond increments. The range for **X** is 0 to 1000. Default is 40.

LUN EXTENT=X sets the system verify extent value **X** in Mbytes. The verify extent determines how much data can be verified before the verify operation must pause. This parameter slows down the verify operation so that it will not affect the performance of the system. Increasing the extent value will allow more data to be verified in a single pass. The range for **X** is 1 to 128 MBs. Default is 32 MBs.

Background TIER Verify Operations

TIER VERIFY verifies LUNs on a tier by tier basis. **TIER VERIFY** differs from **LUN VERIFY** in that the number of simultaneous Tier Verifys is limited to a value that is set by **TIER MAXVERIFIES** (default = 2) parameter. The valid range is 1 to 16. If a tier is marked for continuous verification, once the verification completes, the next sequential tier marked for verification, not presently being verified, will start.

TIER VERIFY Displays a summary of verifications. To enable Tier Verify (refer to the example screen in Figure 3-46):

1. At the prompt, type **TIER VERIFY=ON** <Enter>.
2. The system will ask which tier you wish to verify. Enter the tier number or type **a** for "All."
3. The system will ask if you want run the Tier Verify operation continuously or not. Type **y** to run continuously or **N** to run just once. The default is **N**.

```

11000[2]: tier verify=on
                Tier Status
                Space
Tier Owner      Capacity Available   Disk
                (Mbytes) (Mbytes)   Status
                -----
1 + 1          1121384   1111384 ABCDEFGHPS 0
2 + 1          1121384   1101384 ABCDEFGHPS 1 6
3 + 1          1121384   1110384 ABCDEFGHP. 2 4
4 + 1          1121384   1103384 ABCDEFGHPS 3 5 7 8 9 10 11 12 13

                Tier Verification Status:
TIER  Scheduled   Continuous Progress   Passes   Last Verify Completion Time
-----
1                    No                    0
2                    No                    0
3                    No                    0
4                    No                    0

System verify extent: 16 Mbytes
System verify delay: 30
Maximum active verifies: 2

Please enter a TIER ('a' for all TIERS owned by this unit or 'q' to quit): a
ALL valid TIERS owned by this unit selected
Do you want the verify to run continuously? (y/N): y

```

Figure 3-46 Tier Verify CLI Command Screen Example

To disable Tier Verify (see Figure 3-47 on page 100):

1. At the prompt, type **TIER VERIFY=OFF** <Enter>.

Note: This command only stops a Background Verify operation from running continuously on the TIER(s). The Verify on these TIERS will stop after the next iteration. To stop all Verify operations immediately, use 'TIER STOP'.

2. The system will ask which tier you wish to verify. Enter the tier number or type **a** for “All.” Tier Verify will be disabled off after the next iteration has completed.

```

11000[2]: tier verify=off
                Tier Verification Status:
  TIER  Scheduled  Continuous  Progress  Passes  Last Verify Completion Time
-----
    1                No                0
    2                No                0
    3                No                0
    4                No                0
System verify extent: 16 Mbytes
System verify delay: 30
Maximum active verifies: 2
This command only stops a Background Verify operation from running continuously on the
TIER(s). The Verify on these TIERS will stop after the next iteration.
To stop all Verify operations immediately, use 'TIER STOP'.

Please enter a TIER ('a' for all TIERS owned by this unit or 'q' to quit):

```

Figure 3-47 Tier Verify Disable CLI Command Screen Example

TIER VERIFY=X A specified tier will be verified if possible.

These API and CLI commands will affect the TIER verification process:

- The CLI commands **TIER PAUSE**, **TIER RESUME**, **TIER STOP**.
- The API command **TIER STATUSCHANGE**.

LUN operations (add, delete, move) can affect TIER VERIFY operations. As tiers are verified, only LUNs that are valid and formatted are verified.

If a tier is owned by the other unit, and it is healthy, the user is notified that the verification cannot occur due to the ownership of the tier. The user can then retry the verification on the other unit.

The verification of LUNs on a tier is performed in the order of addressing on the tier. Only valid and formatted LUNs can be verified.

Rebuild Journaling

The **rebuild journaling** feature is intended to speed the recovery from disk-side loss of communication problems. A loss of communication includes, but is not limited to, SAS expanders, hardware/software failures, and SAS cable failures, and SFP failures.

Currently, when the controller encounters a loss of communication with a drive or a group of drives on a SAS link, the software fails the drive(s) and continues operation. When cache coherency is enabled, if either controller encounters a loss of communication with a drive, the firmware will fail the drive. This allows a controller unit to maintain operation during disk-side events. Once the loss of communication is resolved, however, the controller must rebuild all the affected drives. For large installations, a loss of communication, such as a cable failure, can cause the controller to fail numerous disk drives. Once the loss of communication is resolved, the time to rebuild all the failed drives can take many weeks.

The rebuild journals contain bitmaps that indicate which portions of the disks in a tier have been updated with new data while a disk was failed or replaced. The system uses the information in the journals to reduce the rebuild time of drives that have not been swapped out. This can dramatically lower rebuild time, since only portions of the tier may have been updated while the drive was failed or replaced.

The granularity of the journal will be 4MB of data on a single disk or 32MB of host data. Thus a single host write will force the system to rebuild a minimum of 4MB of data on the disk. A new host write into a 4MB section that has already been journaled will not cause a new journal entry. The system will automatically update journals when disks are failed or replaced regardless of whether journaling is enabled.

To ensure that the journals are correct, the system carefully monitors the state of the journals and will automatically invalidate or disable the journals if it detects a condition where the journal cannot be used or journal information could potentially be lost.

The following summarizes the limitations that apply to journaling:

- Rebuild journaling will automatically be disabled if the failed disk is swapped with a new disk. The system will track the serial number of the disks when they are failed and will force a rebuild of the entire disk if the serial number changes.
- Rebuild journaling will not be used when a failed disk is replaced by a spare. The rebuild journal can be used when rebuilding a replaced disk that has not been swapped.
- The system will invalidate the journal on tiers that have failed or replaced disks on boot up. This is required because the system does not save the journal information.
- Rebuild journaling will be managed by the controller that owns the tier. If a controller is failed, then the journals on the tiers owned by that controller will be invalidated.
- The system tracks the original owner of a tier when a drive is failed so changing the ownership of the tier will disable use of the journal for rebuilds on that tier.

- Rebuild journaling will be disabled when rebuilding disks that are failed due to a change in the parity mode of the tier.
- Use of the rebuild journal will be temporarily disabled if the system is rebuilding a LUN that is a backup LUN in a mirror group.

To display the information about the rebuild journal, use the **TIER JOURNAL** command (Figure 3-48). To display the information for a specified tier, use the **TIER JOURNAL=t** command, where **t** is the specified tier. This screen will give detailed information about the status of the journal and a display of all the journal entries for the tier. This screen will also give more detailed information on the status of the journal and will indicate why it was disabled or invalidated.

```
11000[2]: tier journal

Tier Owner          Status    Rebuild  Disk    Capacity    Journal
                Status    OK      Status  (Mbytes)    Entries
-----
 1 + 1             Ready    Off     ABCDEFGHPS  1121384    n/a    17522
 2 + 1             Ready    Off     ABCDEFGHPS  1121384    n/a    17522
 3 + 1             Not valid Off     ABCDEFGHP.  1121384    n/a    17522
 4 + 1             Ready    Off     ABCDEFGHPS  1121384    n/a    17522

Journaling: Disabled

11000[2]:
```

Figure 3-48 Tier Journal CLI Command Screen

The Status field indicates the current status of the journal:

- Ready - The journal is waiting for updates.
- Active - A disk is failed and the journal has updates.
- All other statuses indicate why the journal cannot be used.

The Rebuild OK field indicates if a rebuild can use the journal:

- Off - Journaling not enabled. Use **JOURNAL=ON** to enable.
- Yes - Journaling can be used when rebuilding.
- No - Journaling cannot be used.

The rebuilds will only use the journals if the “Rebuild OK” field indicates “Yes”. In order to use journaling on rebuilds, the operation must be manually started using **DISK REBUILD=tc** where ‘t’ indicates the tier, and ‘c’ indicates the channel or **REBUILD=ALL** which will start a rebuild on all disks.

The **TIER JOURNAL=ON|OFF** command enables/disables use of the journals during rebuild operations. The system will automatically update the journals when disks are failed or replaced regardless of this setting. This parameter only indicates if the journal can be used during the rebuild. The default is OFF.

SES Device Monitoring Rate

The **SES device monitoring rate** can be adjusted to minimize its impact on system performance. **SES M_WAIT** displays the current setting in seconds (Figure 3-49).

```
11000 [1]: ses m_wait
SES timer m_wait = 6 seconds
```

Figure 3-49 SES Device Monitoring Rate Example

SES M_WAIT=x sets the SES device monitoring rate for the system in seconds. Valid range is 4 to 90. The default monitoring rate is 6 seconds.

Note: Improper use of the **SES M_WAIT command** can prevent the SES monitors from detecting an enclosure fault before the enclosure automatically shuts down.

Host Command Timeout

The **Host Command Timeout** parameter allows the system to free up resources and make them available to other users if the request from a particular user cannot be completed. This helps to improve performance in a SAN environment where there are a lot of users accessing the storage.

HOST TIMEOUT=X (where X is value range 1..512) lets the host command timeout for an I/O request in seconds. Valid range is 1 to 512 seconds. Default setting is 75 seconds.

Security Administration

The InfiniteStorage 11000 controller's dual-level, non-host based data security is maintained with scalable features including restricted management access and authentication against authorized listing. No security software is required on the host computers. (Refer to the Section, "Administrator and User Logins" for information regarding Telnet and serial port security.) Each authorized user will have its customized LUN identification scheme which applies to all host ports (see Figure 3-50).

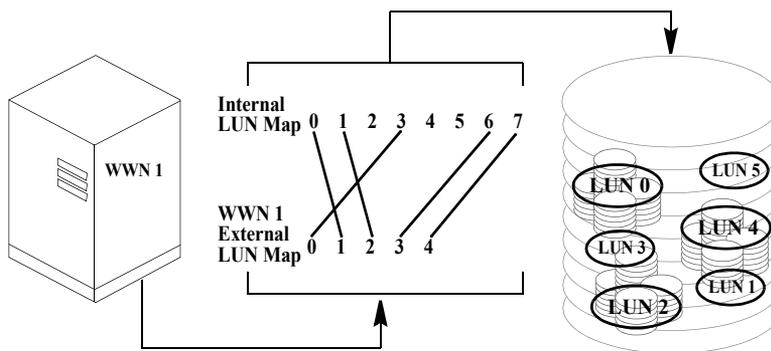


Figure 3-50 Mapping Internal LUNs to External LUNs

Read-only and read/write privileges can be specified for each LUN and for each user.

The “place holder” LUN feature allows the InfiniteStorage 11000 controller administrator to map a zero capacity LUN to a host or group of hosts (via zoning or user authentication). The administrator can then create a real LUN and map it to the host(s) to replace the “place holder” LUN in the future. In most cases, the host will not have to reboot since it already mapped to the “place holder” LUN.

Note: Support of place holder LUNs is dependent upon the operating system, the driver, Host Card Adapter (HCA), and host bus adapter.

Monitoring User Logins

The **AUDIT** function continuously monitors logins to the InfiniteStorage 11000 controller and provides alerts in the event of unauthorized login attempts (see Figure 3-51).

```
Host Int 15:04:07 User Logout Client1, port:4 S_ID:000004
Host Int 15:04:47 Authenticated Login Client10, port:3 S_ID:000002
```

Figure 3-51 User Login Messages

USER AUDIT=ON|OFF enables/disables the user auditing function. When enabled, the system will display a message when a user logs in or out. Default is OFF.

USER CONNECTIONS displays a list of all the currently connected users and the host port to which the user is connected.

Zoning (Anonymous Access)

This type of configuration provides the first-level protection. The LUN **identification scheme** can be customized for each host port. Any unauthorized user accessing the InfiniteStorage 11000 controller will be considered “anonymous” and granted the zoning rights for the host port to which they are connected.

The **ZONING** command displays the current settings for the host ports (Figure 3-52 on page 106). The LUN Zoning chart indicates which internal LUNs the users will be able to access (with Read-only and Read/Write privileges) and where the internal LUN will appear to the users.

```
11000[2]: zoning

Port Zoning Summary:

          LUN Zoning
Port  World Wide Name  (External LUN, Internal LUN)
-----
  1   25000001FF060002  000,000    004,004    013,013
  2   26000001FF060002  000,001
  3   27000001FF060002
  4   28000001FF060002

11000[2]:
```

Figure 3-52 Zoning CLI Command Screen

ZONING EDIT lets you change the settings for the host ports. You will be asked to select a host port to change and enter the mapping for each LUN (Figure 3-53 on page 107). The default configuration is to deny access to all the LUNs.

```

11000[2]: zoning edit

                                LUN Zoning
      Port  World Wide Name  (External LUN, Internal LUN)
-----
      1     25000001FF060002  000,000    004,004    013,013
      2     26000001FF060002  000,001
      3     27000001FF060002
      4     28000001FF060002
Enter the host port (1..4), or 'e' to escape.
3
Enter the new LUN zoning for host port 3.

Enter the unique LUN mapping, as follows:

G.1  GROUP.LUN number
P    Place-holder
R    Before GROUP.LUN to indicate Read-Only
N    Clear current assignment
<cr> No change
E    Exit command
?    Display detailed help text

External LUN 0: is not mapped. Enter new internal LUN: p
.
.
.
External LUN 4: is not mapped. Enter new internal LUN: e

*** Host Port 3: zoning has been updated! ***

                                LUN Zoning
      Port  World Wide Name  (External LUN, Internal LUN)
-----
      1     25000001FF060002  000,000    004,004    013,013
      2     26000001FF060002  000,001
      3     27000001FF060002  000,Place  001,Place  002,Place  003,Place
      4     28000001FF060002
1100[2]:

```

Figure 3-53 Zoning Edit CLI Command Screen

ZONING DEFAULT restores the zoning of a host port back to its default settings.

User Authentication

The InfiniteStorage 11000 controller creates correspondence between users (World Wide Name or GUIDs), storage LUNs, and permissions. The system can store configurations for up to 512 users in total, and the settings apply to all host ports.

Each authorized user will only have access to their own and “allowed-to-share” data determined by their customized LUN identification scheme. Administrators can also restrict users’ access to the host ports and their Read/Write privileges to the LUNs. Unauthorized users will be given the “host port zoning” rights as defined in Section ,“Zoning (Anonymous Access)”.

USER displays the current settings for all authorized users (Figure 3-54). Each user is identified by their 64-bit World Wide Name (or GUID) and is given a unique user ID number. The Ports column indicates which host ports, on each InfiniteStorage 11000 controller, the user is allowed. The LUN Zoning chart indicates which internal LUNs the user will have access to (with read-only and read/write privileges), and where the internal LUN is displayed to the user.

```
11000[1]: user

Host Port Users Summary:

      Unit Port-Map
ID  User Name      World Wide Name  1  2      Zoning Method
-----
0   annemarie      210000E08B9D4FE8 1  4 1  4  R000,001    001,001    ...

User auditing is enabled.
There are 4 users currently connected.
```

Figure 3-54 User CLI Command Screen

To configure/change the settings, use these commands:

- **USER ADD** adds a new user and defines the user’s access rights
- **USER EDIT** edits the access rights of an existing user
- **USER DELETE** deletes an existing user from the system.

See the subsection entitled ““User Authentication (Recommended for SAN Environment)” on page 30” for further information on how to add a new user.

Firmware Update Management

SGI periodically releases firmware updates to enhance features of their products. Contact your SGI or other authorized Technical Support organization to obtain the latest firmware files.

Displaying Current Firmware Version

The **VERSION** command displays version information of the InfiniteStorage 11000 controller's hardware and firmware (see Figure 3-55 for an example).

```
11000[2]: version

Silicon Graphics Storage Appliance Model 11000
  Firmware Version: 5.00_pre_18

  Firmware date: Sep 10 2008, 13:48:41
  IEEE ULA Number: 00060002
  Bootrom Version: 53-20072-001 Version 2.00
  Platform: 11000

11000[2]:
```

Figure 3-55 Version CLI Command Screen Example

Firmware Update Procedure

TFTP enables the administrator to download the new InfiniteStorage 11000 controller firmware from a TFTP server to the InfiniteStorage 11000 controller. A TFTP server, such as the directMONITOR console, must be running when using this command. This command “fails” the current controller and should not be used during active I/O. To update the firmware files:

1. Collect and save the output of the following commands before you update the firmware:

VERSION	AV	CACHE	DISK	DISK LIST
DUAL	HOST	HOST STATUS	LOG	LUN
LUN CONFIG	NETWORK	STATS	STATS DELAY	
STATS TIER	TIER	TIER CONFIG		
DELAY				

2. Copy the new firmware file to your TFTP server (such as the directMONITOR console).
3. Connect to the InfiniteStorage 11000 controller via Telnet or serial (CLI port). Enter **TFTP**
4. You will be asked to confirm action. Enter **y** to continue.
5. Enter the TFTP server's IP address: **TFTP <IP_address>**
6. Enter the firmware path and filename: **TFTP <filename>**
7. **For the couplet controller configuration**, connect and log into the other InfiniteStorage 11000 controller; repeat Steps 3-6 to update the firmware.
8. Enter **RESTART** to restart the unit(s).

Note: RESTART can be done at a later time.

9. (*For dual mode only*): After both controllers are back on-line, use the **DUAL** command to verify that both InfiniteStorage 11000 controller units are healthy. If either controller shows failed, login to the healthy controller and issue the **DUAL HEAL** command.

Remote Login Management

TELNET ENABLE allows the administrator to *temporarily enable* the establishment of a remote Telnet session. Use the **TELNET** command to display the current setting.

TELNET DISABLE allows the administrator to *temporarily disable* the establishment of a remote Telnet session.

Note: Telnet capability is reset to **ON** after a InfiniteStorage 11000 controller restart. To turn off Telnet access permanently, use the **NETWORK** command.

TELNET STATS allows the administrator to view various statistics maintained on remote Telnet sessions. These statistics are kept from the time that the system is powered on.

The administrator is strongly advised to perform any commands affecting the system's configuration from the CLI **UART** only (and not from a Telnet session), and to only perform such commands after issuing the **TELNET DISABLE** command, so that remote users cannot log into the system in the middle of an administrative command.

When a Telnet Session is Active

Whenever a remote Telnet session is active, the current RS-232 console switches to a CLI sub-shell which allows the administrator to enter a very limited sub-set of the CLI commands. The following message is displayed on the console when a Telnet session is initiated from a remote site.

Within the CLI subshell, the **TELNET** command allows the administrator to view information regarding the currently active Telnet session.

TELNET KILL=*m* lets the system administrator terminate the remote Telnet session. The **KILL** parameter may also be specified with **TELNET KILL=*m***, where ***m*** indicates the number of minutes that will be allowed to elapse before the remote Telnet session is terminated. The valid range is <0..15> minutes. Default is 1 minute. An administrative login is required before the command is processed.

The remote user is given a warning that the administrator has killed his session, and indicates to him the amount of time (if any) that he has remaining. An ***m*** value of **0** (zero) is an immediate KILL. The remote user will be notified, but most likely will be unable to read the entire warning message before the session ends.

Note: If a user is in the middle of running a CLI command at a remote Telnet site when the administrative KILL is issued, the command will continue on the CLI console.

System Logs

Message Log

All InfiniteStorage 11000 controller events are logged and saved in non-volatile memory. The log will automatically roll over when it is full.

LOG displays the log of previous system messages.

LOG CLEAR clears the log of all previous messages.

LOG CHECKCONDITION displays the Check Condition log.

LOG CHECKCONDITION=MORE will display additional information concerning the check condition.

LOG CHECKCLEAR clears the Check Condition log, enter **LOG CHECKCLEAR**.

LOG QUIET= ON|OFF. This Administrator command enables a “quiet mode” on the CLI where Message Log statements will still be logged, but not displayed.

LOG QUIET will display the current state of the Log Quiet mode. There should be the word “Quiet” at the CLI prompt when the Log Quiet mode has been enabled.

System and Drive Enclosure Faults

Use the **FAULTS** command to display a list of all current disk, system, and drive enclosure faults or failures (see Figure 3-56).

```
11000[2]: faults

    Current System Faults
    -----

Disk 3S Failed.

Not connected Unit 1 11000[1]

Disk channel B not ready.
Disk channel D not ready.
Disk channel F not ready.
Disk channel H not ready.
Disk channel S not ready.

    --- Faults detected! ---

Multi-bit ECC Error Shutdown: Enabled
Task Exception Shutdown:      Enabled

11000[2]:
```

Figure 3-56 Faults CLI Command Screen

To display the current SDRAM memory faults (*ECC- error controller counters*), use **FAULTS MEMORY** command. To clear the values in the memory faults (ECC) statistics, use **FAULTS MEMCLEAR** command.

For FC models, to display the current status of the host SFPs, use the **FAULTS SFP** command.

Note: A transmitter fault and a loss of signal on a disk channel or host port may indicate that there is no connection at the corresponding connector.

To display the number of LUN array parity errors detected by the system, use the **FAULTS ARRAYPARITY** command. The system saves the counts for each tier of all the LUNs. To clear the count of LUN array parity errors in the system, use the **FAULTS ARRAYPARITYCLEAR** command.

FAULTS BUSPARITY displays the number of bus parity and data path errors detected by the system.

FAULTS BUSPARITYCLEAR clears the count of errors.

You may set a parameter (**ECCSHUTDOWN**) that allows the system to automatically shutdown if it encounters an unrecoverable error. Use the **FAULTS ECCSHUTDOWN=on** command to enable automatic shutdown for unrecoverable ECC errors. This is the default setting. To disable and allow the system to continue to run in spite of unrecoverable ECC errors, use the **FAULTS ECCSHUTDOWN=off**.

The **EXCEPTIONSHUTDOWN** command parameter allows the system to automatically shutdown if it encounters a task exception.

FAULTS EXCEPTIONSHUTDOWN=ON enables automatic shutdown for task exceptions. This is the default setting.

The **FAULTS EXCEPTIONSHUTDOWN=OFF** disables automatic shutdown and allows the system to continue to run in spite of task exceptions.

Displaying the System's Uptime

UPTIME displays the total time the system has been operational--or “uptime”(also known as **Power on Hours**), as well as the total time since the last system restart (see Figure 3-57). The uptime is displayed as **YY:DDD:HH:MM** where **YY** is the number of years, **DDD** is the number of days, **HH** is the number of hours, and **MM** is the number of minutes. The Uptime Since Restart is displayed as **DDD:HH:MM:SS** where **DDD** is the number of days, **HH** is the number of hours, **MM** is the number of minutes, and **SS** is the number of seconds.

```
11000[2]: uptime

System Uptime (YY:DDD:HH:MM): 00:003:22:35
Since Restart (DDD:HH:MM:SS): 001:01:49:35
```

Figure 3-57 Uptime CLI Command Screen

Saving a Comment to the Log

COMMENT <text of message> allows you to echo a message to the screen. The message is saved in the LOG and is also sent to syslog if it is enabled. Any printable text can be entered.

Other Utilities

APC UPS SNMP Trap Monitor

APC_UPS displays the status of the **APC UPS SNMP** trap monitor (see Figure 3-58).

```
11000 [1]: apc_ups

APC UPS SNMP trap monitor is off.
No APC UPS faults detected via SNMP trap.
```

Figure 3-58 ACP UPS SNMP Trap Monitor Status

APC_UPS CLEAR_FAULTS will delete all pending APC UPS faults from the fault list. All APC UPS events that disabled writeback caching will be cleared.

API Server Connections

The **API** command displays the current status of the API connections (see Figure 3-59).

```
11000[2]: api
API Server connections are currently: Enabled
```

Figure 3-59 API CLI Command Screen

API DISABLE temporarily enables/disables the establishment of connections to the API server. When disabled, users at remote locations will be unable to establish a new API connection until an **API ENABLE** command is issued. This command only provides control over API connections during the current power cycle.

API STATS displays the collected statistics on API connections (see Figure 3-60).

API CLEARSTATS resets the collected statistics.

```
11000[2]: api stats

      API Server Connection Statistics
      =====

                Time      Date
                -----
System Boot-Up Completion      : 17:52:26 09/10/2008
System's Current                : 19:45:49 09/11/2008

API Server Initiation           : 18:13:24 09/10/2008
API Server Connections have been : Enabled
                                since : 18:13:24 09/10/2008

API Server services have been :
Enabled : 1 time.
Disabled : 1 time.
```

Figure 3-60 API Stats CLI Command Screen

Changing Baud Rate for the CLI Interface

The **CONSOLE** command displays the current serial console setting (see Figure 3-61) of the InfiniteStorage 11000 controller.

```
11000[1]: console
Serial Console baud rate is 115200 baud.
Console hardware flow control is OFF.
Console software flow control is OFF.
```

Figure 3-61 Console CLI Command Screen

CONSOLE BAUD changes the baud rate of the CONFIG port of the InfiniteStorage 11000 controller (see Figure 3-62).

```
11000 [1]: console baud

Select the new serial console baud rate from choices below:
 1 - 9600
 2 - 19200
 3 - 38400
 4 - 57600
 5 - 115200 <- Current setting
 e - escape out of this command

Enter selection:
```

Figure 3-62 Changing the Console Baud Rate

CLI/Telnet Session Control Settings

You may change the CLI's and Telnet's various session control settings. The **SETTINGS** command displays the current setting.

SETTINGS DEFAULTS resets all the CLI and Telnet session control settings to their default values.

SETTINGS LINES=<number of lines> sets the number of lines displayed at a time in a page of screen information. Pages provide a way to control the amount of information displayed to the user at one time. You will be prompted to either press a specified key in order to scroll from one page to the next, or (in certain circumstances), to terminate the display. Valid range is 0 to 512

lines, where 0 indicates that no paging is to be performed on the output information. Default setting is 0.

SETTINGS PROMPTINFO=ON enables extra status information in the CLI prompt. **OFF** disables extra status information in the CLI prompt. When enabled, the CLI prompt indicates whether the system is booting (**BOOTING**), failed (**FAILED**) or waiting for a restart (**RESTART NEEDED**). **ON** is the default.

Disk Diagnostics

The **DISK DIAG=tc** command performs a series of diagnostics tests on the specified disk. The disk is specified by its physical: tier (t) in the range <1..128>, and channel (c) in the range <ABCDEFGHPS>.

Disk Reassignment and Miscellaneous Disk Commands

The **DISK REASSIGN=tc 0xh** command allows for the reassigning of defective logical blocks on a disk to an area of the disk reserved for this purpose. The disk is specified by its: tier (t) in the range <1..128>, and channel (c) in the range <ABCDEFGHPS> 0xh is the hexadecimal value of the LBA (Logical Block Address) to be reassigned.

The **DISK LLFORMAT=tc** command allows the user to perform a low level format of a disk drive. The disk is specified by its: tier (t) in the range <1..128>, and channel (c) in the range <ABCDEFGHPS>.

The **DISK AUTOREASSIGN=ON** command is the default setting. When enabled bad blocks are reassigned when a medium error occurs on a healthy tier, the **DISK AUTOREASSIGN=OFF** command disables this feature and bad blocks are NOT reassigned when a medium error occurs on a healthy tier.

The **DISK MAXCMDS=x** command sets the maximum command queue depth to a tier of disks in the range of 1 to 32 commands per tier. The default is 32 commands.

SPARE Commands

Use the **SPARE** command to display information about the spare disks in the system or to change the configuration settings for background diagnostics in the system. The information displayed pertains to the current spare configuration settings as well as task status (see Figure 3-63 on page 118).

```

11000[2]: spare

                               Spares Installed

Vendor      Product ID      Test Status
-----

No Spares Found

EXTENT      : 8 MBytes
DELAY       : NONE
COVERAGE    : 1%
PATTERN     : UNIQUE
Diagnostics are ENABLED

```

Figure 3-63 Spare CLI Command Screen

The **SPARE** CLI commands are for background diagnostics. The intent of these commands is to test otherwise idle spare disks at least one (1) time per month to validate that they are continuing to function properly, and are truly available to be swapped in as a replacement disk. It is testing of the “hot” spares. They are intended to run in the background and SPARE operations are always at lower priority than any other kind of I/O in the system.

The **SPARE INFO=tc** command displays the information and status about a specific spare disk in the system. The disk is specified by its physical tier and channel locations, “tc.” The “t” indicates the tier in the range <1..128>, and the 'c' indicates the channel in the range <ABCDEFGHPS>.

The **SPARE CLEAN=tc** command erases any previous test data stored on the disk indicated. The disk is specified by its physical tier and channel locations, “tc.” The “t” indicates the tier in the range <1..128>, and the 'c' indicates the channel in the range <ABCDEFGHPS>.

The **SPARE COVERAGE=x** command sets the spare diagnostic coverage of the blocks being tested as a percent of the total number of blocks available for test. Note that increasing the coverage to higher numbers means that more blocks on the disk will be tested for better coverage, but it also will take a longer time for the test to complete. This parameter can be tuned to provide an optimal test time for a single disk in the system such that all spares are tested in a reasonable amount of time. The parameter is limited to a discrete set of values. The valid parameters for “x” are [1, 5, 10, 20, 40, 80, 100] Percent. Default is 1 Percent.

The **SPARE EXTENT=x** command sets the spare diagnostic extent in Mbytes. The diagnostic extent determines how much data can be tested before the test must sleep. This parameter slows down the test operations so they will not affect the performance of the system. Increasing the

extent will allow more data to be tested in a single pass. Any changes applied to extent will affect tests in progress as well as future testing. The valid range for 'x' is 1..32 Mbytes. Default is 8 MB.

The **SPARE DELAY=x** command sets the system spare diagnostics delay. The test delay determines how long a test operation will pause after it reaches the test extent. This parameter slows down the spare test so it will not affect the performance of the system. Any changes applied to delay will affect tests in progress as well as future testing. This system spare diagnostic delay value is given in 100 millisecond increments. The valid range for 'x' is 0..100. The default is 0.

The **SPARE PATTERN=x** command sets the system spare diagnostics pattern. The test pattern determines the pattern written to the disks during the test. The system supports these patterns:

- **UNIQUE** Includes unique information including timestamp
- **AA** 0xAA is written to each byte
- **55** 0x55 is written to each byte
- **FF** 0xFF is written to each byte
- **00** 0x00 is written to each byte
- **COUNTUP** A pattern of counting up is written to each byte
- **COUNTDOWN** A pattern of counting down is written to each byte

The default is **UNIQUE**. Note that the tests in progress are not affected by this parameter setting. Changing the pattern only applies to tests started after the parameter was modified.

The **SPARE START** command starts the spare diagnostics task if it is not running. Note that this will start diagnostics on both units in a dual system, as this is a system parameter.

The **SPARE STOP** command aborts any ongoing diagnostic operations. Note that this command will stop them and then the task will be idle until the **SPARE RESTART** command is executed. Note that this will stop diagnostics on both units in a dual system as this is a system parameter.

The **SPARE PAUSE** command pauses but does not stop any ongoing diagnostic operation, only on the unit from which the command is run. If a test is being run from the other unit in a dual, the pause command will NOT affect that test.

The **SPARE RESUME** command releases any paused diagnostic operations and allows them to continue only on the unit from which the command is run. If a test has been paused on the other unit in a couplet, the **SPARE RESUME** command will NOT affect that test.

Remote Management of the InfiniteStorage 11000 controller

The InfiniteStorage 11000 controller can be managed locally through the RS-232 interface, or remotely via Telnet. The Administrative Utility is the same regardless of the management interface (RS-232 or Telnet).

The InfiniteStorage 11000 controller can also be operated through a java-supported graphical user interface (GUI) management tool, which comes with its own user guide, **The Management Tool User Guide**, that instructs the user in installing and operating this GUI.

The InfiniteStorage 11000 controller supports SNMP and allows the system to be remotely monitored.

Network Connection

Connect the Telnet port on the back of the InfiniteStorage 11000 controller to your Ethernet network (see Figure 3-64). Then set the IP addresses, login names, and passwords as described in the following sections.

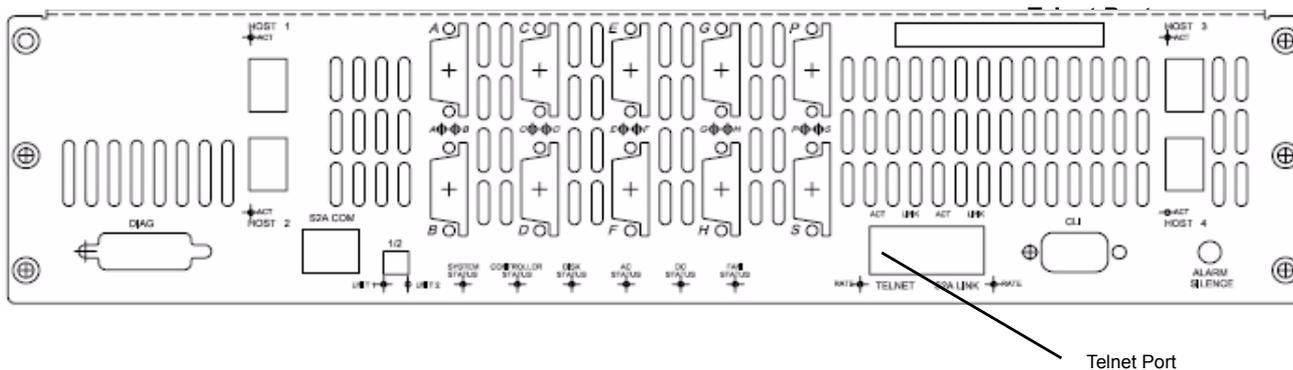


Figure 3-64 Telnet Port on the Controller

Note: Currently, the InfiniteStorage 11000 controller does not support network configuration protocols such as DHCP or BOOTP.

Network Interface Set Up

For first time set up, you will need to connect to the CLI (RS-232) port in order to change the IP address and/or network settings.

To set up the network interface:

1. Use the **NETWORK** command to display the current settings (see Figure 3-65).

```

11000[2]: network
                                Network Configuration
-----
MAC Address:                      00:01:ff:06:00:02
Link Status:                      100 Mbps Full duplex

                                Unit 1          Unit 2
-----
IP Address:                       192.168.10.48      192.168.10.50
Subnet Mask:                      255.255.0.0        255.255.0.0
Default Gateway:

-----
Telnet:                           Enabled      23           23
API Server:                       Enabled      8008         8008
Syslog:                           Enabled      514          514
Syslog IP Address:
SNTP:                             Enabled
SNTP IP Address:                  192.168.5.1
SNMP:                             Disabled
SNMP IP Address:
Limited SNMP:                     Disabled

11000[2]:

```

Figure 3-65 Network CLI Command Screen

2. Use the **network IP=<new IP address>** command to change the InfiniteStorage 11000 controller's IP address for your network environment.
3. Use the **network netmask=<new netmask>** command to change the netmask of the InfiniteStorage 11000 controller (if needed).
4. Use the **network telnet=ON** command to enable the Telnet capability (if needed).

Note: Telnet connections are clear text. If Telnet connections are used, you may expose InfiniteStorage 11000 controller passwords to third parties. For higher security, we recommend that you disengage Telnet access if it is not required.

5. Decide whether the SNMP functionality should be enabled. To enable or disable SNMP, use the appropriate version of the command **network SNMP=on|off**.

Note: If you are using **directMONITOR** (the external system console option), the **SNMP** function should be enabled.

6. If the SNMP function is enabled, enter the IP address of the computer to be used to monitor the SNMP traps: **network trapip=<computer's IP address>**
7. Decide whether the Syslog capability should be enabled. To enable (**ON**) or disable (**OFF**) the Syslog, enter: **network syslog=on|off**

Note: If you are using **directMONITOR**, the syslog function should be enabled.

8. If the SYSLOG function is enabled, enter the destination IP address for the Syslog packets: **network SYSLOGIP=<destination IP address>**
9. Ensure your destination computer supports the SYSLOG feature. For example, on UNIX systems, the SYSLOG application must be properly installed and running.
10. The default destination port number for Syslog packets is **514**. To change it, enter: **network SYSLOGPORT=<port number>**
 - Set up the routing table. This table describes how the InfiniteStorage 11000 controller communicates with the hosts on other networks. Use the **ROUTE** command to display the current settings (see Figure 3-66 on page 123).
 - **ROUTE GATEWAY=<aaa.bbb.ccc.ddd>** sets the current gateway in the network routing table to the specified Internet address.
 - **ROUTE DEL=<aaa.bbb.ccc.ddd> GATEWAY=<aaa.bbb.ccc.ddd>** deletes gateways from the routing table.
 - **ROUTE ADD=<aaa.bbb.ccc.ddd> GATEWAY=<aaa.bbb.ccc.ddd>** adds new gateways to the table.
 - If you have couplet InfiniteStorage 11000 controllers, connect to (or Telnet to if this is not the initial set up) and log into the other controller. Repeat the appropriate steps to set its network parameters.

```

11000[2]: route
Gateway:
Permanent Routing Table:
=====
destination      gateway
-----
-- None --      -- None --
-----

Current Routing Tables:
=====
ROUTE NET TABLE
Destination      Gateway          Flags  Refcnt  Use      Interface
-----
192.168.0.0      192.168.10.50   0x101  4        0        gei0
-----

ROUTE HOST TABLE
Destination      Gateway          Flags  Refcnt  Use      Interface
-----
127.0.0.1        127.0.0.1       0x5    0        15475    lo0
-----

```

Figure 3-66 Route CLI Command Screen Example

11. On boot up, verify the network connections, enter **NETWORK PING=<IP address of a network system>**.
 - Ping a host on the same subnet as the InfiniteStorage 11000 controller.
 - Ping another host on a different subnet.
12. Make sure your firewall is running (which will block traffic and keep hosts from talking to each other), ensure it is set up to allow the InfiniteStorage 11000 controller to pass information on Port 23 (for Telnet), Port 161 (for SNMP), and Port 162 (for SNMP traps).

Login Names and Passwords

The InfiniteStorage 11000 controller's two levels of security, administrative and general purpose user access, also applies to remote management. If you login as an administrator, you can access all the management and administrative functions. You can obtain status information and make changes to the system configuration. At the user access level, you are only allowed to view the status and configuration information of the system.

The login names and passwords can be changed using the **PASSWORD** command (see Figure 3-67), via RS-232 or Telnet. By default, the administrator name is **admin** and its password is **password**. Similarly, the default user name is **user** and its password is **password**. If a user forgets the password, entering **PASSWORD DEFAULTS** while logged in as “admin” will restore all passwords and user names to the default values.

```
11000[2]: login
Enter a login name: admin
Enter the password: *****

    Successful Telnet session login.
        New owner      : admin.
        New security level: Administrative.

11000[2]: password
Enter current password:
*****
Enter a new name to replace 'admin', or return to leave unchanged:
Administrative user name 'admin' unchanged.

Enter new password:
*****
Re-enter the new password:
*****

Enter a new name to replace 'user', or return to leave unchanged:

General user name 'user' unchanged.

Enter new password:

Password for general user 'user' unchanged.
Committing changes.

11000[2]: password defaults
All login names and passwords are now restored to their default values.
```

Figure 3-67 Password and Password Defaults CLI Commands Screen

Note: Only one Telnet session is permitted at a time. Once a Telnet session is initiated, the RS-232 console switches to a CLI sub-shell (see the section on “Firmware Update Procedure” on page 110 for more information). The Telnet client should be using port 23 and have its local echoing function disabled.

SNMP Set Up on Host Computer

The following procedure explains how to set up the host computer to monitor the InfiniteStorage 11000 controller's events (SNMP traps).

Important: It is not possible to configure the InfiniteStorage 11000 controller through SNMP.

Refer to this guide for information regarding network and SNMP configuration on the InfiniteStorage 11000 controller. Contact your sales or service provider to obtain the **Management Information Base (MIB)** files for the InfiniteStorage 11000 controller.

1. Ensure the host computer, which will receive the traps, has an SNMP browser (such as HP OpenView) properly installed and configured. Refer to your browser's documentation for instructions on how to load the MIB files.
2. Set up the host computer to listen to Ports 161 and 162.
3. Load the following InfiniteStorage 11000 controller MIB. Note these MIB files:
 - `sdd.mib`
 - `sdd_trap.mib`

Note: When loading the InfiniteStorage 11000 controller MIB, `sdd_trap.mib` must be loaded last.

4. Start a query. You should be able to see various controller information such as **tempLevel**, **powerNumber**, **fanNumber**, **powerTable**, **tierNumber**, and **tierTable**.
5. To verify that the SNMP function has been set up correctly:
 - Unplug one of the power supplies.
 - Check that a power failure message appears in your SNMP browser.

InfiniteStorage 11000 Controller Implementation of SNMP

InfiniteStorage 11000 Controller MIB (The MIB is read only.)

- Temperature status {normal, warning, critical}
- Fan Status ¥1 {healthy, failure}
- Power Supply Status ¥4 {healthy, failure}
- Tier Status ¥125
 - Disk A Status {healthy, failure, missing}
 - Disk B Status {healthy, failure, missing}
 - Disk C Status {healthy, failure, missing}
 - Disk D Status {healthy, failure, missing}
 - Disk E Status {healthy, failure, missing}
 - Disk F Status {healthy, failure, missing}
 - Disk G Status {healthy, failure, missing}
 - Disk H Status {healthy, failure, missing}
 - Disk P Status {healthy, failure, missing}
 - Disk S Status {healthy, failure, missing}

Traps

- Generic Traps
- coldStart (occurs when the InfiniteStorage 11000 controller is restarted)
- Enterprise Specific Traps
 - Any changes in InfiniteStorage 11000 controller MIB.
If any variable in the InfiniteStorage 11000 controller MIB changes, a trap will be sent. For example, when a power supply fails or is replaced, a trap is sent. Trap includes ASCII string declaring new status of the controller element.
 - InfiniteStorage 11000 controller status alerts.
Trap includes ASCII string declaring the InfiniteStorage 11000 controller as failed or healed.

Troubleshooting the SGI InfiniteStorage 11000 Controller and Drive Enclosures

Component Failure Recovery

This chapter provides information regarding error recovery on the InfiniteStorage 11000 controller and drive enclosures.

To display a list of all current controller system and drive enclosure faults, enter command: **faults**

The InfiniteStorage 11000 controller contains redundant and serviceable fans and power supply units. A single component failure, therefore, will not shut down the system. However, in the unlikely event of a 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.

Controller Power Supply Failure

A power supply failure message will be displayed on your console if a power supply fails. The status LED on the power supply module will also turn off (see Figure 4-1 on page 128).

Important: Always use the correct power supply replacement part for the InfiniteStorage 11000 controller. Power supplies used in SGI InfiniteStorage 10000 and 6700 storage products **cannot** be used in the InfiniteStorage 11000 product.

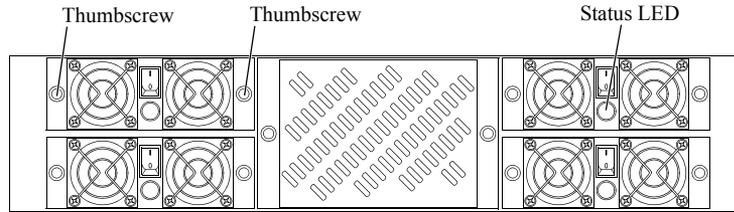


Figure 4-1 Power Supply and Cooling Modules on the Controller Front

Important: Make sure you have the correct replacement power supply available before removing the failed module.

1. Remove the cover panel (if the panel is installed).
2. Locate the failed power supply module which is indicated by an **off** Status LED.
3. Turn off the failed power supply module's power switch and the power switch of its partner supply (which is located directly above or below the failed power supply module within the same controller unit).
4. Disconnect the AC cord from the rear of the unit on the corresponding side where the power supply module is to be replaced.
5. Loosen the two thumbscrews on the failed power supply module and slide the module out of the bay.
6. On the new module, check that the power switch is off.
7. Slide the module into the bay. Make sure it is fully inserted. Install and tighten the two thumbscrews to secure it. The thumbscrews should be tightened with an appropriate screwdriver after both installation and subsequent access to the power supply or fan module.
8. Reconnect the AC cord into the rear of the unit. Verify that the plug is fully seated into the receptacle.
9. Turn on the power switch on the new power supply module and its partner supply. Check that the Status LED is green on both supplies, indicating that the modules are operating normally.
10. Replace the cover panel (if necessary).

Controller Fan Failure

A fan failure error message will display on your console if the fan module fails. Follow these steps to replace the module.

Caution: The controller should not be operated without the fan module for more than 15 seconds. Make sure you have the replacement fan unit available before removing the failed fan unit.

1. If a cover panel has been installed, take it off by removing its two thumbscrews.
2. Remove the two thumbscrews from the fan module (see Figure 4-2).
3. Slide the module out of its bay.

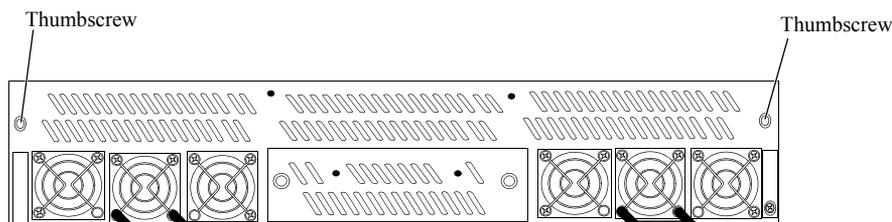


Figure 4-2 Controller Fan Module Removal

4. Slide the new module into the bay, making sure it is fully inserted.
5. Install the thumbscrews to secure it.
6. Replace the cover panel (if necessary).

Recovering from Drive Failures

When a drive failure occurs on the InfiniteStorage 11000, the tier containing that drive will begin operating in degraded mode. This means that the tier 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 tier. If another disk drive fails on the same tier before the data on the first drive is rebuilt (to a replacement disk or hot spare), the tier will go offline.

If one disk in a tier fails, the data or parity information on the failed disk will be reconstructed from the parity disk and data disks of that tier. An entire channel may fail without data loss.

It is recommended that the controller's Automatic **Disk Rebuild** function be enabled at all times (**TIER AUTOREBUILD=ON**).

Note: If more than one disk is failed in a single tier, only the first disk to fail will be replaced by the spare disk. Any disks which fail after the first in each tier will not be reconstructed using a hot spare.

Single Drive Failures

A single drive failure in any tier does not result in the data loss. The LUN(s) on that tier will continue to operate in degraded mode. If a spare drive is available, the controller will automatically rebuild the data on the spare drive if “autorebuild” is enabled. System operation is not affected while recovery is taking place.

When a drive failure occurs, the controller displays an error message on your console, indicating which drive has failed and which spare drive (if available) is replacing the failed drive. If your InfiniteStorage 11000 drive enclosure was ordered from SGI as part of a storage solution, the **Fault LED** on the failed drive will turn **amber**. The error is also 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 change the rebuild rate to match the user load.

To obtain additional information, use the **TIER** or **LOG** command.

TIER displays the status of the disks. The failed drive is denoted by an “**r**”. If a Rebuild is taking place, the percentage of completion is also displayed.

DISK INFO will indicate what spare drive is replacing the failed drive.

Returning the system to a fault-tolerant state

When you find a failed drive on a tier, replace it as soon as possible so that the tier can return to its optimal state. If you replace the failed drive while a rebuild is taking place, the system will finish rebuilding data on the spare drive first. When a rebuild is complete and the failed drive has been replaced, the system will automatically copy the data from the spare drive to the new drive, and return the spare drive to its standby state. The tier is then fully restored.

If no suitable spare drive exists when a drive fails, replace the failed drive as soon as possible. Once the replacement drive is inserted and verified by the system, a rebuild will start automatically if Autorebuild is ON. If not, initiate a rebuild by doing the following:

1. Enter: **diskscan**
The system checks each channel and looks for newly inserted drive(s).
2. Enter: **disk rebuild=<tier><channel>**.
This command initiates a disk build by tier and channel.

Manually Replacing a Failed Disk with a Spare Disk

You may manually replace a specified failed disk with a spare disk using **DISK REPLACE=<tier><channel>**. A Replace operation is used to temporarily 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 speed of the replace operation can be adjusted with the **DELAY** and **EXTENT** parameters (see Chapter 3 in this guide).

Changing the Rate of Rebuild

The commands **TIER DELAY** and **TIER EXTENT** control the percentage of processor time allocated to rebuild compared to I/O handling. Rebuild is done in steps. **DELAY** will control the amount of wait time before rebuilding the next chunk of data. **EXTENT** determines the size of data used to rebuild.

To increase the rate of rebuild, reduce the **DELAY** value (refer to Chapter 3 for more information). We recommend that you use the default **DELAY** and **EXTENT** settings unless you have a specific need to change them.

Interrupting the Rebuild Operation

To abort the rebuild, enter **tier stop** to stop all current rebuild operations.

Multiple Drive and Channel Failures

If multiple drives on the same drive channel fail simultaneously, the problem may be a channel failure rather than a series of drive failures. For example, if drives 1D, 2D, and 3D fail at the same time, the problem may be in Channel D rather than in the drives. If the drives fail as the result of a channel failure, data on the drives may not be lost. Any single channel failure can be recovered.

Before you replace any drives:

1. Use the **TIER** command to check the current disk status and see if the drives failed are all on the same channel.
2. Use the **DISK** command to check the status of the channel.
3. The Recovery steps are:
4. Contact your Technical Support organization to determine the cause of error and what steps may need to be taken to recover from the error.
5. Isolate the cause of channel failure (for example, loose cable connections).
6. Rebuild all the drives on that channel individually by issuing the command:
`disk rebuild=<tier><channel>`

For example, **disk rebuild=1d**, **disk rebuild=2d**, **disk rebuild=3d**.

Several rebuild operations are processed at the same time. The other rebuild jobs are queued up and processed in the same order as they were entered.

Component Failure on Enclosures

The InfiniteStorage 11000 controller implements the **SCSI Enclosure Services (SES)** protocol to communicate with its drive enclosures. If your enclosure provides SES communications, then enclosure status information of the enclosure, including power supply, fan, and presence of drive, will be obtained and evaluated. If a change in status is found, an SES message is displayed on your console (see the figure for an example).

```
EncID: 50001FF1:3E081000 Power Supply 1: OK
```

Figure 4-3 Example SES Message

If your enclosures provide redundant SES communication paths, the message will be reported twice. “EncID” is the Enclosure Identifier of the enclosure that reported the failure. The last four digits of the WWN is the last four digits of the enclosure’s serial number.

Drive Enclosure Troubleshooting Introduction

The SGI InfiniteStorage 11000 Drive Enclosure includes an Enclosure Services Processor and associated monitoring and controller 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 PCMs. Each unit is independently monitored.

Initial Start-up Problems

Faulty Cords

First check that you have wired up the subsystem correctly. If:

- cords are missing or damaged,
- plugs are incorrect,
- cords are too short,

Call your supplier for a replacement.

Alarm Sounds On Power Up

Refer to the sections later in this chapter on “Dealing with Hardware Faults” and also “Audible Alarm.”.

Green “Signal Good” LED on I/O Module Not Lit

Check to ensure that the cables have not been reversed during installation.

Computer Doesn't Recognize the Drive Enclosure Subsystem

1. Check that the SAS interface cables from the drive enclosure to the host computer, or I/O controller, are fitted correctly.
2. Check that all drive carrier modules have been correctly installed.
3. Check that there is a valid SAS signal present at the I/O connector. If there is no signal present, ensure the cable has been properly inserted.

4. Check the I/O module setup as follows:
 - Ensure the I/O module has been correctly installed and all external links and cables are securely fitted.
 - Ensure the maximum cable length has not been exceeded.

Status LEDs

Green LEDs are always used for good or positive indication. Amber LEDs indicate there is a critical fault present within the module. See specific LED tables for further information.

HDD (Hard Disk Drive)

When a HDD is faulted, Table 4-1 describes the fault LED behavior.

- Under Normal conditions, the LEDs should all be illuminated constant GREEN.
- If a problem is detected, the color of the relevant LED will change to AMBER.

Table 4-1 HDD LEDs

Location	Color	Identifier	Behavior
Enclosure front	Amber	Enclosure fault	This LED is ON with a drive fault
Enclosure front	Amber	Drive fault	This LED is ON with a drive fault

PCM (Power Cooling Module) Status LEDs

The PCM LEDs are shown in Table 4-2 on page 135.

- Under Normal conditions, the LEDs should all be illuminated constant GREEN
- If a problem is detected, the color of the relevant LED will change to AMBER.

Table 4-2 PCM LEDs

Location	Color	LED Identifier	Behavior
FRONT Enclosure	Amber	Enclosure fault	This LED is ON with a PCM fault.
REAR Enclosure PCM	Amber	PCM fault	This LED is ON with an AC input, DC output, fan or other PCM fault.
REAR Enclosure PCM	Green	AC OK	This LED is ON with a DC or fan fault. This LED is OFF with an AC input failure.
REAR Enclosure PCM	Green	DC OK	This LED is ON with an AC or fan fault. This LED is OFF with an DC output failure

DEM (Drive Expander Module)

The DEM fault LEDs are explained in Table 4-3.

- Under Normal conditions, the LEDs should all be illuminated constant GREEN
- If a problem is detected, the color of the relevant LED will change to AMBER.

Table 4-3 DEM LEDs

Location	Color	LED Identifier	Behavior
FRONT Enclosure	Amber	Enclosure fault	ON with a DEM fault.
FRONT Enclosure	Amber	DEM fault	ON with a DEM fault.
Internal (DEM) Enclosure	Amber	DEM fault	ON with a DEM fault.
Internal (DEM) Enclosure	Green	DC OK	ON when the DC voltage regulation is within limits. OFF with a DC output failure.
Internal (DEM) Enclosure	Green	Expander MIPS ready	ON when the expander internal processor is booted and operating correctly. OFF when the expander internal processor is NOT booted or ready.
Internal (DEM) Enclosure	BLUE	DEM identify	ON indicates that this DEM is being sent an identify command by the SEP.

I/O Module

The I/O fault LEDs are explained in Table 4-4.

- Under Normal conditions, the LEDs should all be illuminated constant GREEN
- If a problem is detected, the color of the relevant LED will change to AMBER.

Table 4-4 I/O LEDs

Location	Color	LED Identifier	Behavior
FRONT Enclosure	Amber	Enclosure fault	ON with an I/O fault.
REAR Enclosure	Green	I/O OK	OFF when an I/O module detectable fault is present. INDEPENDENT of SAS link fault condition and stays lit during a SAS link fault.
REAR Enclosure	Amber	I/O fault	ON when a detectable fault is present in the I/O module. INDEPENDENT of SAS link fault condition and stays lit during a SAS link fault.
REAR Enclosure	Green	SAS Link Activity	ON when a valid link is present on any of the 4 links of the 4-wide port. OFF when none of the 4 links of the 4-wide port have a valid connection.
REAR Enclosure	Amber	SAS Link Fault	ON when a fault is present on any of the 4 links of the 4-wide port. OFF when none of the 4 links of the 4-wide port have a fault.

Front Panel Drive Activity Indicators

The Front Panel Drive Activity Indicators show the aggregated status of all the modules. This panel and its associated LEDs are shown in Figure 4-4. The Display Panel LEDs are defined in Table 4-5.

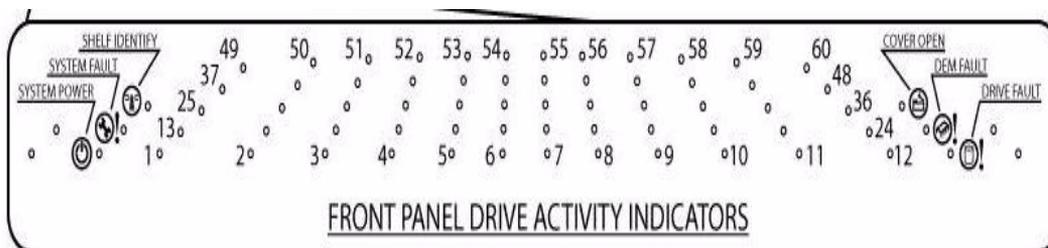


Figure 4-4 Front Panel Drive Activity Indicators Panel

Table 4-5 Display Panel LEDs

ICON	LABEL	COLOR	Definition	Normal Behavior
	Shelf Identify	BLUE	enclosure Identity	ON -SES is sending an identity command. OFF -SES is NOT sending an identity command
	System Fault	AMBER	Enclosure Fault	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	Enclosure Powered ON	ON - DC power is present OFF - DC power is NOT present LED does NOT flash under normal operating conditions.

Table 4-5 (continued) Display Panel LEDs

ICON	LABEL	COLOR	Definition	Normal Behavior
	Cover Open	AMBER	Enclosure cover is open	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	DEM fault	OFF -all DEMs operating correctly. ON -at least one DEM has failed; service action required
	Drive Fault	AMBER	Drive fault	ON -one or more HDDs are faulted. SES must determine exact HDD. OFF -no detectable drive faults

Audible Alarm

When the Audible Alarm sounds, use the appropriate SES CLI command.

Top Cover Open

The Audible Alarm sounds anytime the top cover is open. The alarm mutes automatically once the cover is closed.

SES Command

The Audible Alarm can be set through SES communication. When set in this manner, the alarm remains on until any one of the following conditions is met:

- Enclosure is power cycled.
- Alarm is turned off through SES command.

General Drive Enclosure Troubleshooting

The following sections describe common problems (and possible solutions) you might encounter with your drive enclosure system.

Thermal Control

The drive enclosure uses extensive thermal monitoring and takes a number of actions to ensure component temperatures are kept low and also to minimize acoustic noise (see Table 4-6). Air flow is from the front, over the drive section, through the midplane, through the I/O modules to rear of the enclosure.

Table 4-6 Thermal Monitor Actions

Symptom	Cause	Action
<p>If the ambient air is cool (below 25 °C) and the fans are observed to increase in speed, then some restriction on airflow may be causing additional internal temperature rise.</p> <p>Note: This is not a fault condition.</p>	<p>The first stage in the thermal control process is for the fans to automatically increase in speed when a thermal threshold is reached. This may be caused by higher ambient temperatures in the local environment and may be perfectly normal.</p> <p>Note: This threshold changes according to the number of drives and PCMs fitted.</p>	<ol style="list-style-type: none"> 1 Check the installation for any airflow restrictions at either the front or rear of the enclosure. A minimum gap of 25mm at the front and 50mm at the rear is recommended. 2 Check for restrictions due to dust build-up; clean as appropriate. 3 Check for excessive re-circulation of heated air from rear to the front. Use in a fully enclosed rack installation is not recommended. 4 Check that all Blank modules are in place. 5 Reduce the ambient temperature.

Thermal Alarm

Table 4-7 provides actions to be taken when symptoms of thermal problems are present on the system.

Table 4-7 Thermal Alarm Actions

Symptom	Cause	Action
<ul style="list-style-type: none"> • Display Panel PCM/Cooling Fault LED. • An AMBER LED on one or more PCMs. • Audible Alarm Sounding. • Air temperature exiting PCM above 35°C. 	<p>If the internal temperature measured in the airflow through the enclosure exceeds a pre-set threshold, a thermal alarm will sound.</p>	<ol style="list-style-type: none"> 1 Check local ambient environment temperature is below the upper 40°C specification. 2 Check the installation for any airflow restrictions at either the front or rear of the enclosure. A minimum gap of 25mm at the front and 50mm at the rear is recommended. 3 Check for restrictions due to dust build-up. Clean as appropriate. 4 Check for excessive re-circulation of heated air from rear to the front. Use in a fully enclosed rack installation is not recommended. 5 If possible, shutdown the enclosure and investigate the problem before continuing.

Thermal Shutdown

Important: An enclosure will shut down when a critical temperature threshold is exceeded in order to prevent permanent damage to the disk drives.

Dealing with Drive Enclosure Hardware Faults

Ensure that you have obtained a replacement module of the same type *before* removing any faulty module.

Warning: If the drive enclosure subsystem is powered up and you remove any module, replace it immediately. If the subsystem is used with modules or module blanks missing for more than a few minutes, the enclosure can overheat, causing power failure and data loss. Such use will invalidate the storage system warranty.

- Replace a faulty drive with a drive of the same type and equal or greater capacity.
- All the supplied plug-in PCMs, electronics modules, and blank modules must be in place for the air to flow correctly around the cabinet.

Continuous Operation During Replacement

Depending on how the subsystem is set up, if a disk unit fails, it can normally be replaced without interrupting the use of the system.

Replacing a Module

Caution: Whenever replacing a module NEVER leave an EMPTY bay in the rear of the enclosure, obtain a replacement before removing the problem part.

Please refer to Chapter 6 in this guide for information on the initial installation of the plug-in modules in the drive enclosure.

Important: Observe all conventional ESD precautions when handling drive enclosure modules and components. Avoid contact with Backplane components and module connectors, etc.

Power Cooling Modules

If a power supply unit or its fan is faulty, you must replace the whole PCM. You must not take any longer than 5 minutes to replace this unit to prevent the enclosure from over-heating.

Caution: Do not remove covers from the Power Cooling module (PCM). There is a danger of electric shock. Return the PCM to your supplier for repair.

Warning: To ensure your system has warning of a power failure please disconnect the power from the power supply, by either the switch (where present) or by physically removing the power source, prior to removing the PCM from the enclosure/shelf. Do not remove the faulty PCM unless you have a replacement unit of the correct type ready for insertion. The system must not be run without all units in place.

As there should always be two (2) PCMs installed, you can continue working while replacing the faulty module.

To remove an AC PCM/Cooling Module- Handle Model:

Caution: Do not remove the faulty PCM/Cooling module unless you have a replacement unit of the correct type ready for insertion.

If a power supply unit or its fan is faulty, you must replace the whole PCM/Cooling module.

As there should always be two PCMs installed, you can continue working while replacing the faulty module.

1. Make sure you identify the faulty PCM correctly from the two modules installed.
2. Lift latch which secures the power supply cord.
3. Switch off and disconnect the power supply cord.
4. Lift the PCM handle to the open position to cam the PCM out of the enclosure.
5. Grip the handle and withdraw the PCM.

To remove an AC PCM - Thumb Screw Model:

Warning: Do not remove the faulty PCM unless you have a replacement unit of the correct type ready for insertion.

If a power supply unit or its fan is faulty, you must replace the whole PCM.

As there should always be two PCMs installed, you can continue working while replacing the faulty module.

1. Make sure you identify the faulty PCM correctly from the two modules installed.
2. Lift latch which secures the power supply cord.
3. Switch off and disconnect the power supply cord.
4. Remove the two thumbscrews on the right and left sides of the unit.
5. Firmly grip the handle on the bottom of the unit and withdraw the PCM.

To install an AC PCM - Handle Model:

1. Check for damage, especially to the rear connector on the PCM.

Caution: Handle the module carefully and avoid damaging the connector pins. Do not install the module if any pins appear to be bent.

2. With the PCM handle in the open position, slide the module into the enclosure.
3. Cam the module home by manually closing the PCM handle. A click should be heard as the handle latches engage.
4. Connect the power supply cord to the power source, secure the latch over the cord, and switch the power supply ON.

Note: The alarm will sound until the new PCM is operating correctly.

To install an AC PCM - Thumbscrew Model:

1. Check for damage, especially to the rear connector on the PCM.

Caution: Handle the module carefully and avoid damaging the connector pins. Do not install the module if any pins appear to be bent.

2. Slide the module into the enclosure and push until you hear a click as the latch is engaged.
3. Secure the two thumbscrews on the left and right sides of the unit.
4. Connect the power supply cord to the power source; secure latch, and switch the power supply ON.

Note: The alarm will sound until the new PCM is operating correctly.

I/O Module

Warning: Do not remove this module unless a replacement can be immediately added. The system must not be run without all units in place.

To remove the I/O Module:

1. Release the two latches on the bottom of the unit by simply pulling each latch out and away from unit.
2. Pull the latches forward to cam the module out of the enclosure.
3. Grip the unit securely and withdraw the module.

To install the I/O Module:

1. With the latches in the open position, slide the EBOD module into the enclosure until the latch engages automatically.
2. Cam the module home by manually closing the latches. A click should be heard as the latch engages.

Replacing the Drive Carrier Module

Caution: Observe all conventional ESD precautions when handling drive enclosure modules and components. Avoid contact with backplane components and module connectors, etc.

Instructions are also clearly labeled on the inside cover of the enclosures, see Figure 4-5.

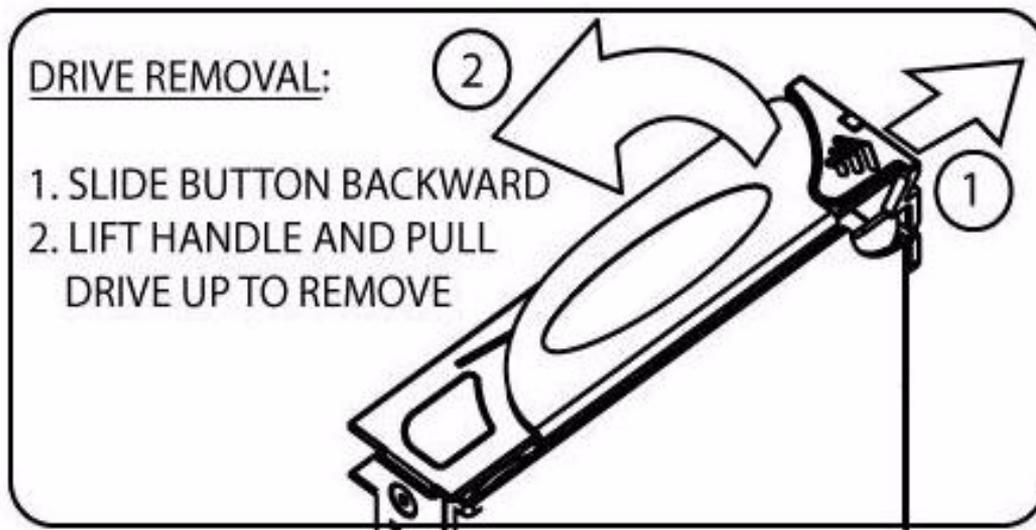


Figure 4-5 Drive Carrier - Handle Released

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 Insert the Drive Carrier Module:

1. Release the carrier handle by sliding the latch backwards.
2. Insert the carrier into the enclosure.

3. Slide the drive carrier, gently, all the way into the enclosure.
4. Cam the drive carrier home. The camming foot on the base of the carrier will engage into a slot in the enclosure.
5. When the carrier is fully home, close the handle. You should hear a click as the latch engages and holds the handle closed.
6. Close the enclosure.

Replacing the DEM

Important: The DEM card should only be replaced by trained personnel, contact your service provider as needed.

SGI InfiniteStorage 11000 Drive Enclosure Overview

The Drive Enclosure

The SGI InfiniteStorage 11000 drive enclosure storage system is an ultra dense (4U-high) 60 disk drive enclosure. Figure 5-1 and Figure 5-2 on page 148 show front and rear views of a drive enclosure, respectively. A top-down view of the module's major components and their locations are shown in Figure 5-3 on page 148.

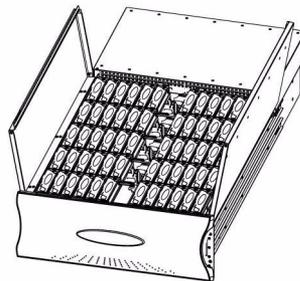


Figure 5-1 The Drive Enclosure System - front open view

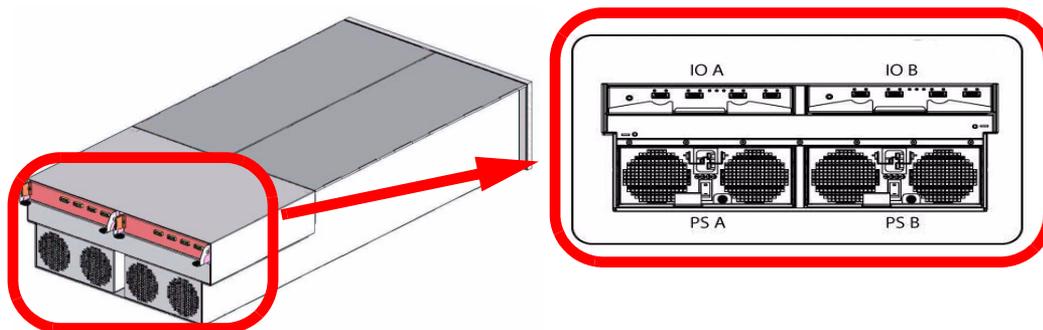


Figure 5-2 The Drive Enclosure System - rear isometric and rear views

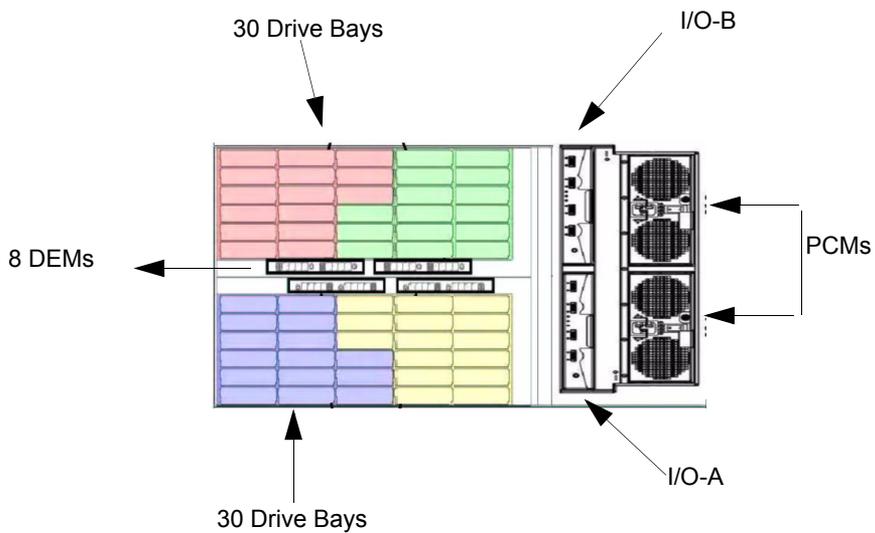


Figure 5-3 Drive Enclosure Module Locations

Enclosure Core Product

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

- **Enclosure Chassis** with integral Front Panel Drive Status Indicator (see Figure 5-9 on page 156).
- Two (2) 1865W **Power Cooling (PCM)** plug-in modules (see Figure 5-4 on page 150).
- Two (2) plug-in **Input/Output (I/O)** modules (see Figure 5-5 on page 152).
- A midplane separates the front and back of the chassis and provides the interconnect system between the PCMs, I/O modules, and the baseboard.
- Up to 60 top loadable hard disk drives (HDDs) in a 5x12 matrix.
- SAS and SATA drive intermix allowed.
The InfiniteStorage 11000 supports both SATA and SAS interface drives. SATA and SAS interface drives can be mixed within a configuration. However, the mixing of SAS and SATA drives within a tier is not recommended or supported.
- Eight (8) SAS disk expander modules (DEMs).

Enclosure Chassis

The chassis assembly contains 60 drive bays at the front, each of which accommodates a plug-in drive carrier module. The 60 drive bays are arranged in five rows of twelve drives (5x12). At the rear, the chassis assembly contains two (2) PCMs and two (2) I/O modules.

The chassis is fitted with 19-inch rack mounting features which enables it to be fitted to four (4) 19-inch racks and uses four (4) EIA units of rack space.

The Plug-in Modules

A drive enclosure requires the following modules for normal operation:

- Power Cooling Module (PCM)
- Input/Output Module (I/O)
- Drive Carrier Module

- Drive Expander Module (DEM)

Power Cooling Module (PCM)

Two (2) auto ranging AC 1865W Power Cooling modules (Figure 5-4) are supplied already mounted in the rear of the enclosure as part of the subsystem core product.

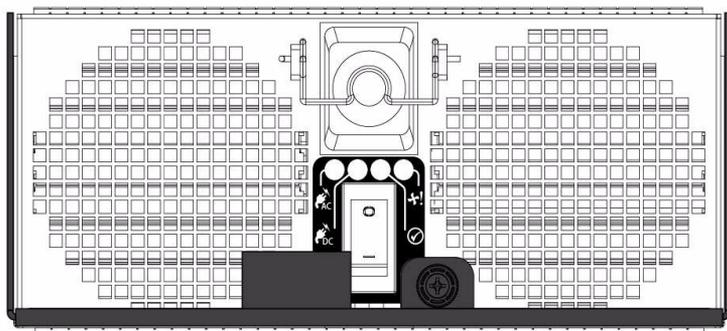


Figure 5-4 Power Cooling Module (1)

The drive enclosure must always be operated with two Power Cooling Modules fitted (2 PCMs). Module replacement should only take a few minutes to perform but must be completed within 10 minutes from removal of the failed module. Four (4) LEDs mounted on the PCM (Table 5-1 on page 151) indicate the status of the PCM and the fans. PCM voltage operating ranges are nominally 200V to 240V AC.

Table 5-1 PCM LEDs

PCM LEDs				
		Faults / LEDs	Enclosure Audible Alarm	
	GREEN	AC Input FAILURE	off	Off
	GREEN	DC Output FAILURE	off	Off
	OFF	Fan Fault / Power Supply Fault	AMBER	Off
	OFF	ID	BLUE	Off

Input/Output (I/O) Module

The drive enclosure storage subsystem includes an enclosure with rear facing bays which house two (2) I/O modules (see Figure 5-5). Processors housed on the I/O modules provide enclosure management and interface to devices on the Backplane, PCM, and Display Panel in order to monitor internal functions.

The plug-in I/O modules have been designed for integration into a drive enclosure storage subsystem, providing external SAS cable interfacing with up to 60 SAS or SATA disk drives.

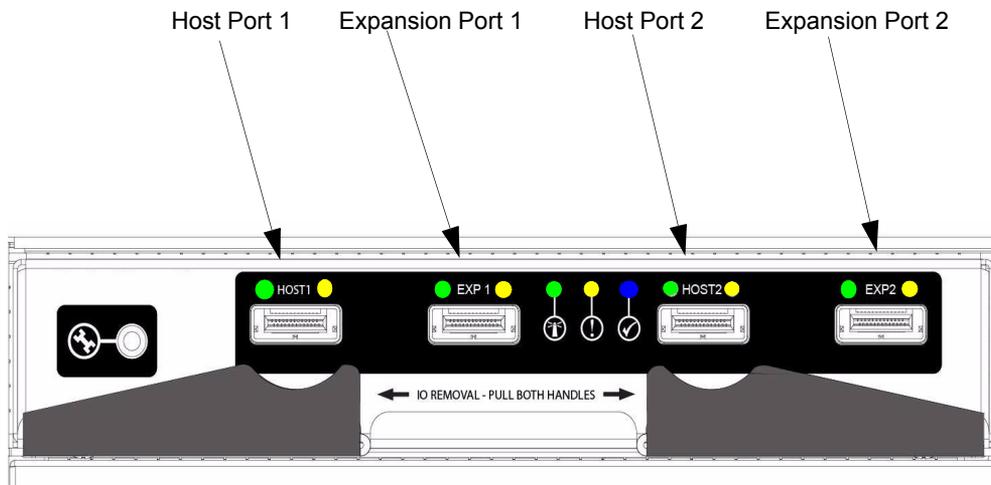


Figure 5-5 I/O Module

Table 5-2 defines the LED indicators incorporated on the I/O module.

Table 5-2 I/O Module LEDs

LED		Color	STATUS	Definition	Enclosure Audible Alarm
Host Port 1 Expansion Port 1 Host Port 2 Expansion Port 2		Green	ON OFF	Incoming signal is GOOD No connection or incorrect connection	Off
Host Port 1 Expansion Port 1 Host Port 2 Expansion Port 2		Amber	ON OFF	Fault No faults	Off
 I/O Module OK	I/O Module OK	Green	ON OFF	properly booted and functioning correctly internal fault	Off
 I/O Module Fault	I/O Module Fault	Amber	OFF ON	No Faults present I/O Fault	Off
 I/O Module Identity	I/O Module Identity	Blue	ON OFF	Receiving SES Command Not receiving SES Command	Off

Drive Carrier Module and Status Indicator

The Drive Carrier Module comprises a hard disk mounted in a carrier (see Figure 5-6). Each drive bay can house a single Low Profile 1.0 inch high, 3.5 inch form factor disk drive in its carrier.

The handle provides the following functions:

- Camming of carrier into and out of drive bays.
- Positive 'spring loading' of the drive/baseplane connector.
- The handle assembly also incorporates a Drive Status LED (see Figure 5-6).

Note: The enclosure system design allows for drive bays to be left empty without the need for fitting dummy drive carriers.

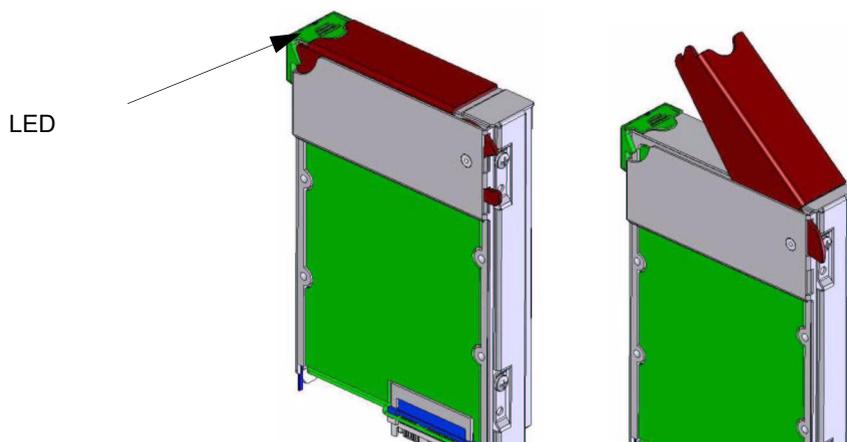


Figure 5-6 Drive Carrier Module: Closed and Opened

DEM Card

The drive enclosure contains eight (8) top-loadable DEMs. The **DEM**s (see Figure 5-7 on page 155) provide the SAS connectivity between the I/O module and the HDDs located within the enclosure.

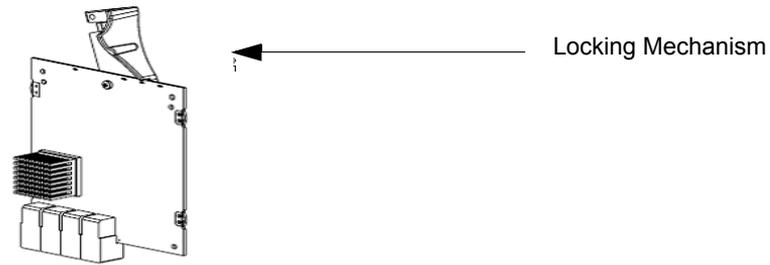


Figure 5-7 DEM

Note: The DEM is a serviceable PCB which may be replaced by trained personnel.

Each DEM connects to a single port of the HDD based on the location within the enclosure within which it is installed (see Figure 5-8). Each provides power control signals to each drive slot. It controls HDD identify/service LEDs and monitors the status from partner DEM. Each pair set DEM supports 15 drives; there are 8 DEMs in a fully configured system.

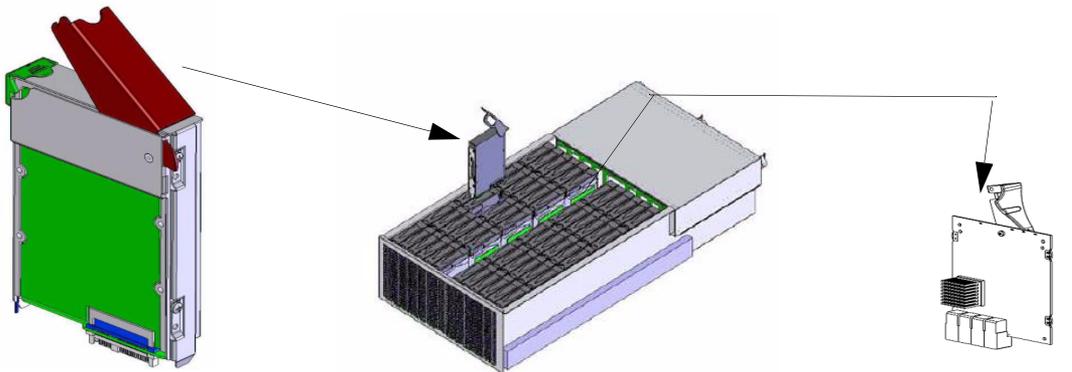


Figure 5-8 Drive Carrier and DEM Locations

Indicators

The drive enclosure 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 HDD, has a visible AMBER LED which indicates a fault.

Note: In some components, the failure LED is internal to the enclosure and visible only when the cover is open.

The state of each LED within the enclosure is available through the in-band SES functionality at all times. The host can manually set the state of each LED through an in-band SES page.

Front Panel Drive Activity Indicators

The Front Panel Drive Activity Indicators show the aggregated status of all the modules. This panel and its associated LEDs are shown in Figure 5-9. Drive activity LEDs flash during data I/O when transition cards are fitted to drives. These LEDs indicate drive presence. The Front Display Panel LEDs are defined in Table 5-3 on page 157.

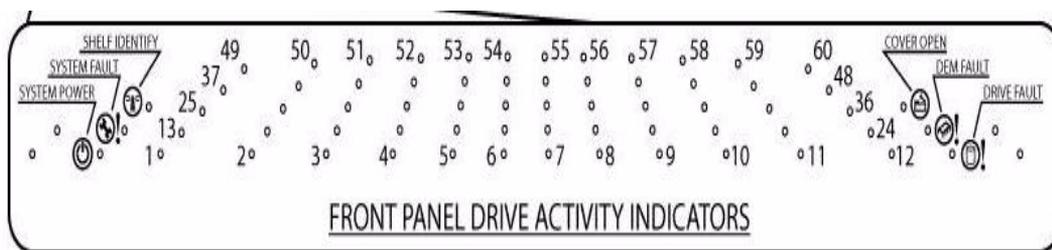


Figure 5-9 Front Panel Drive Activity Indicators Panel

Table 5-3 Display Panel LEDs

ICON	LABEL	COLOR	Definition	Normal Behavior	Enclosure Audible Alarm
	Shelf Identify	BLUE	enclosure Identity	ON -SES is sending an identity command. OFF -SES is NOT sending an identity command	Off
	System Fault	AMBER	Enclosure Fault	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	Off
	System Power	GREEN	Enclosure Powered ON	ON - DC power is present OFF - DC power is NOT present LED does NOT flash under normal operating conditions.	Off
	Cover Open	AMBER	Enclosure cover is open	OFF -both cover pieces securely closed and latched in place. ON -either of the cover pieces is NOT securely closed and latched in place	Off
	DEM Fault	AMBER	DEM fault	OFF -all DEMs operating correctly. ON -at least one DEM has failed; service action required	Off
	Drive Fault	AMBER	Drive fault	ON -one or more HDDs are faulted. SES must determine exact HDD. OFF -no detectable drive faults	Off
Individually numbered	HDDs	GREEN	HDD activity	ON -indicates HDD activity OFF -no HDD activity	Off

Internal Indicators

Note: In some components, the failure LED is internal to the enclosure and visible only when the cover is open. The Internal LED indicators are explained in Table 5-4.

Table 5-4 INTERNAL LED Indicators

Description	Location	Color	Normal Behavior	Enclosure Audible Alarm
DEM DC	DEM internal to enclosure	GREEN	ON -1.2VDC regulator circuit correctly functioning OFF - faulty 2.1VDC regulator circuit	Off
DEM ID	DEM internal to enclosure	BLUE	ON -receiving SES identity command OFF -NOT receiving SES identity command	Off
HDD ID	HDD internal to enclosure	BLUE	ON -receiving SES identity command OFF -NOT receiving SES identity command	Off

Rear of Enclosure Activity Indicators

The PCMs and the I/O Modules are located on the rear of the enclosures. The LEDs on the rear of the enclosure are explained in Table 5-5.

Table 5-5 LEDs on the Rear of the Enclosure

Description	Location	COLOR	Normal Behavior	Enclosure Audible Alarm
PCM DC ok	PCM rear of enclosure	GREEN	ON -DC output of PCM within tolerances. OFF -failed PCM	Off
PCM AC ok	PCM rear of enclosure	GREEN	ON -AC input to PCM within tolerances. OFF -failed PCM	Off
PCM Fault	PCM rear of enclosure	AMBER	ON - PCM fault detected OFF - no detected PCM faults.	Off
PCM ID	PCM rear of enclosure	BLUE	ON -receiving SES identity command OFF -NOT receiving SES identity command	Off

Table 5-5 (continued) LEDs on the Rear of the Enclosure

Description	Location	COLOR	Normal Behavior	Enclosure Audible Alarm
I/O Module ok	I/O Module rear of enclosure	GREEN	ON -properly booted and functioning correctly. OFF -I/O Module internal fault	Off
I/O Fault	I/O Module rear of enclosure	AMBER	ON - I/O module fault detected OFF - no detected I/O module faults.	Off
I/O Module ID	I/O Module rear of enclosure	BLUE	ON -receiving SES identity command OFF -NOT receiving SES identity command	Off
SAS Link	I/O Module rear of enclosure	GREEN	ON -a valid SAS link established on at least 1 of the 4 SAS links of the 4-wide SAS port. OFF -no SAS links	Off
SAS Link Fault	I/O Module rear of enclosure	AMBER	ON -a detectable fault on at least 1 of the 4 SAS links of the 4-wide SAS port. OFF -no detectable faults	Off

Visible and Audible Alarms

The functional modules have associated status LEDs. The Display Panel (reference Table 5-3 on page 157) shows a consolidated status for all modules:

- Constant GREEN or BLUE LEDs indicate good or positive indication.
- Constant AMBER LEDs indicate that there is a fault present within that module.

Note: The Display Panel also incorporates an *Audible Alarm*. The only time the enclosure will sound an alarm is when the top cover of the enclosure is opened. There is a warning beep when the cover is first opened. If the top cover is left open for more than five minutes, the audible alarm will be continuous.

Important: The Display Panel is an integral part of the enclosure chassis assembly and is not field replaceable.

Drive Enclosure Technical Specifications

Dimensions

Enclosure	inches	millimeters
Height	6.97	177
Width front	16.56	420.6
Width rear	17.56	446
Depth	42	1067
front surface of bezel to end of cable management arms		
LENGTH	36	914
rack mounting surface to rear connectors surface		

Weight

Item	Quantity per System	LBS (kg)	Extended Weight
Chassis Includes Base PCB & Midplane	1	83	83
PCM	2	12	24
I/O Module Drives	2	4	8
	60	2	120
Miscellaneous	1	5	5
Total Weight			240

AC INPUT PCM

	Value
Input Voltage	190 - 264 VAC
Input Frequency	47 - 63 Hz and 400 hz
Maximum Input Current	13.1A RMS @ 190 VAC
Peak Inrush Current	50A @ 230VAC, 25°C, 5 msec max
Minimum Efficiency measured at 30% and 50% load and nominal line (208 VAC)	90%
Maximum Average Output Power	1750 W
Maximum Peak Output Power, 1 second	1922W

DC INPUT PCM

Input Parameter	Value
Input Voltage	36 - 72 VAC
Minimum Efficiency measured at 30% and 50% load and nominal line (48 DVC)	90%
Maximum Average Output Power	1750 W
Maximum Peak Output Power, 1 second	1922W

DC OUTPUT PCM

	3.3 VDC Output	5 VDC Output	12 VDC Output
Normal Voltage	3.3 VDC	5.10 VDC	12.10 VDC
Full Load	5.0A	55.0A	120.0A
Peak Load for 1 Second	5.0A	65.0A	130.0A
Minimum Load (A)	0.0A	0.0A	0.0A

PCM Safety and EMC Compliance

Safety Compliance UL 60950
 IEC 60950
 EN 60950

EMC Compliance CFR47 Part 15B Class A
 EN55022
 EN55024

Power Cord

(minimum requirements)

Cord Type SV or SVT, 18 AWG minimum, 3 conductor
Plug 250V, 10A
Socket IEC 320 C-14, 250V, 15A

Note: Power cord is not included in the standard drive enclosure package.

Environmental Compliance

The enclosure and all of its sub-components are compliant to the European Union RoHS (Restriction of Hazardous Substances) Directive (Directive 2002/95/EC) with no exceptions or exemptions

Drive Enclosure Installation

This chapter provides information about installing or removing components drive enclosures from your SGI storage system.

Introduction

Caution: When connecting up the drive enclosure subsystem, use only the cables supplied.

Planning the Installation

Before you begin installation, you should become familiar with the configuration requirements of your drive enclosure system and the correct positions of each of the optional plug-in modules (Table 6-1 on page 164).

Table 6-1 Drive Enclosure Configuration

Module	Location
Drive Bays	<p>You can install from 1 - 60 drives in the enclosure.</p> <p>In 1x mode, (connected to 10 enclosures), you can install any number of drives in the enclosure, but every enclosure should have the same number of drives.</p> <p>In 2x mode, (connected to 5 enclosures), you must install drives in multiples of two (2).</p>
Power Cooling Modules (PCM)	<p>Two (2) PCMs must be fitted. Full power and cooling redundancy is provided while a faulty module is replaced. Install the PCMs in lower rear bays A and B.</p> <p>Rear bays are numbered from A and B starting from the left when viewed from the back.</p>
I/O Module	<p>One or two I/O modules should be Installed in the upper rear bays A & B.</p> <p>If only one module is installed, it should be fitted in Module Location A and a blank plate must be fitted over the unused bay.</p>

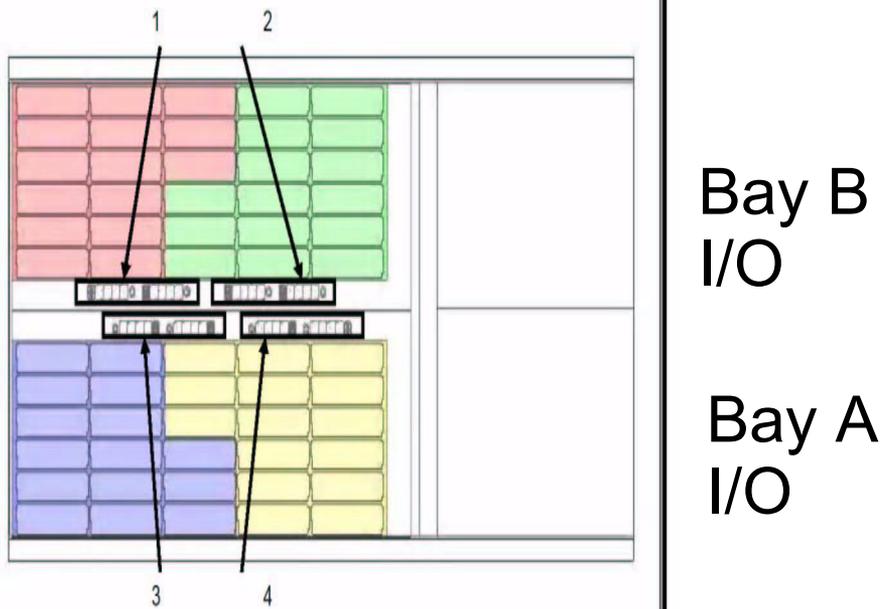


Figure 6-1 DEM Pair Locations

Enclosure Bay Numbering Conventions

Warning: Operation of the Enclosure with *ANY* of the plug-in modules missing from the rear of the enclosure will disrupt the airflow and the drives will not receive sufficient cooling. It is *ESSENTIAL* that all (rear) apertures are filled before operating the unit.

The drive enclosure subsystem is housed in a 60-drive bay enclosure, arranged in four (4) groups. Each group comprises two banks of 15 bays, that is, (as viewed from the front): 12 bays across the enclosure by 5 bays deep. There are two numbering schemes: 1x60 and 2x30. The drive bays are numbered in accordance with the tables shown in Figure 6-2 on page 166 and Figure 6-3 on page 167, when viewed from above.

Note: Drive 61 is an optional drive.

DRIVE 1	DRIVE 13	DRIVE 25	DRIVE 37	DRIVE 49
DRIVE 2	DRIVE 14	DRIVE 26	DRIVE 38	DRIVE 50
DRIVE 3	DRIVE 15	DRIVE 27	DRIVE 39	DRIVE 51
DRIVE 4	DRIVE 16	DRIVE 28	DRIVE 40	DRIVE 52
DRIVE 5	DRIVE 17	DRIVE 29	DRIVE 41	DRIVE 53
DRIVE 6	DRIVE 18	DRIVE 30	DRIVE 42	DRIVE 54
AUX DRIVE	DRIVE 18	DRIVE 30	DRIVE 42	DRIVE 54
	DRIVE 19	DRIVE 31	DRIVE 43	DRIVE 55
	DRIVE 20	DRIVE 32	DRIVE 44	DRIVE 56
	DRIVE 21	DRIVE 33	DRIVE 45	DRIVE 57
	DRIVE 22	DRIVE 34	DRIVE 46	DRIVE 58
	DRIVE 23	DRIVE 35	DRIVE 47	DRIVE 59
	DRIVE 24	DRIVE 36	DRIVE 48	DRIVE 60

Figure 6-2 1 X 60 Drive Numbering Table.

DRIVE 1	DRIVE 7	DRIVE 13	DRIVE 19	DRIVE 25
DRIVE 2	DRIVE 8	DRIVE 14	DRIVE 20	DRIVE 26
DRIVE 3	DRIVE 9	DRIVE 15	DRIVE 21	DRIVE 27
DRIVE 4	DRIVE 10	DRIVE 16	DRIVE 22	DRIVE 28
DRIVE 5	DRIVE 11	DRIVE 17	DRIVE 23	DRIVE 29
DRIVE 6	DRIVE 12	DRIVE 18	DRIVE 24	DRIVE 30
AUX DRIVE	DEM 2B	DEM 2A	DEM 1B	DEM 1A
DRIVE 1	DRIVE 7	DRIVE 13	DRIVE 19	DRIVE 25
DRIVE 2	DRIVE 8	DRIVE 14	DRIVE 20	DRIVE 26
DRIVE 3	DRIVE 9	DRIVE 15	DRIVE 21	DRIVE 27
DRIVE 4	DRIVE 10	DRIVE 16	DRIVE 22	DRIVE 28
DRIVE 5	DRIVE 11	DRIVE 17	DRIVE 23	DRIVE 29
DRIVE 6	DRIVE 12	DRIVE 18	DRIVE 24	DRIVE 30
	DEM 4A	DEM 4B	DEM 3A	DEM 3B

Figure 6-3 2 x 30 Drive Numbering Table

Enclosure Installation Procedures

Important: SGI InfiniteStorage 11000 drive enclosures should only be installed in SGI InfiniteStorage 11000 racks. Mounting and installing these drive enclosures in any other rack is not authorized or supported by SGI.

Warning: The drive enclosure with all its component parts installed is too heavy for a single person to easily install into a Rack cabinet.

Caution: Ensure that you have fitted and checked a suitable anti-static wrist or ankle strap and observe all conventional ESD precautions when handling drive enclosure modules and components. Avoid contact with Backplane components and module connectors, etc.

Note: Drive enclosures are supplied and delivered populated with Backplane, Baseplane, Front Panel, DEMs, and with PCMs installed. The Drive Carrier Modules are supplied as a separate package.

A typical rack mounting installation is shown in Figure 6-4

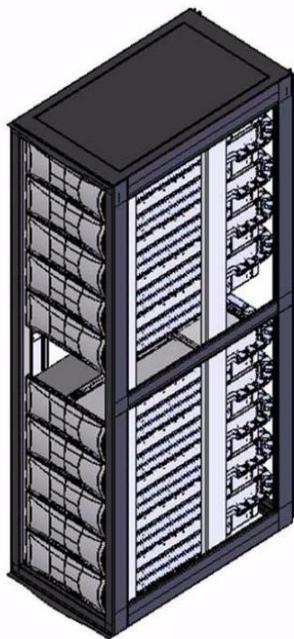


Figure 6-4 Full Rack without Controllers.

I/O Module Configurations

Controller Options

The drive enclosure has a standard SAS JBOD IO controller, which employs a 36-port expander with zoning. The zoning is implemented via a preset piano-switch located on the enclosure baseboard. Host ports 1 and 2 present drives 1 through 30 and host ports 3 and 4 present drives 31 through 60. Each host port responds with drive information applicable only to the drives presented on the respective zone and enclosure status information for the entire enclosure.

SAS DEM

The drive enclosure contains eight (8) top-loadable DEMs. The DEMs provide the SAS connectivity between the I/O module and the HDDs located within the enclosure. Each DEM connects to a single port of the HDD based on the location within the enclosure in which it is installed. Each of the eight DEMs installed in the enclosure has a unique I²C and SAS address based on a proprietary SAS address generation scheme.

SATA Interposer Features

The drive enclosure supports an active/active SATA MUX interposer which allows simultaneous access to the SATA HDD from both I/O modules installed in the enclosure. The MUX is implemented as a dongle module installed in the HDD carrier assembly. Additionally, the SATA interposer allows the power cycling of an individual HDD via SES commands.

Grounding Checks

The product must only be connected to a power source that has a safety electrical earth connection.

Warning: If more than one product is fitted in a rack, the earth connection to the rack is even more important, because the rack will then have a high “EARTH LEAKAGE CURRENT” (“TOUCH CURRENT”).

The earth connection to the rack must be checked before switching on, by an electrical engineer who is qualified to the appropriate local and National standards to perform the check.

Drive Enclosure Operation Overview

Before You Begin

Before powering up the enclosure, please ensure that all the drive modules are firmly seated in their correct bays.

Power On / Power Down

Caution: Do not operate the subsystem until the ambient temperature is within the specified operating range. If the drives have been recently installed, ensure they have had time to acclimatize before operating them.

To Power On the enclosure:

- Apply AC power to the enclosure.
- Turn the PCMs to ON.

All LEDs on the I/O Panel should be lit (Green) when the enclosure power is activated (and the disk drive motors should start).

Important: If AC power is lost for any reason, on restoration of power the enclosure will re-start automatically.

To power the enclosure down:

- Switch off the PCMs installed in the enclosure.

PCM LEDs

The PCM incorporates 4 LEDs, located above the On/Off switch. Please refer to Figure 7-1 and Table 7-1.

I/O Panel LEDs

The I/O Panel LEDs fault and status conditions are shown in Figure 7-1

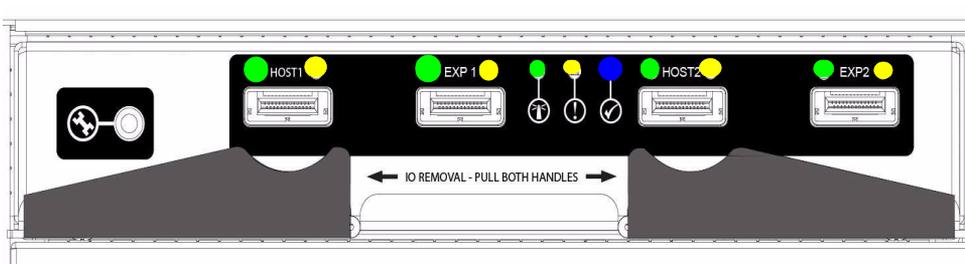


Figure 7-1 I/O Panel LEDs

Table 7-1 defines the LED indicators incorporated on the I/O module.

Table 7-1 I/O Module LEDs

LED	Color	STATUS	Definition
Host Port 1 Expansion Port 1 Host Port 2 Expansion Port 2	Green	ON OFF	Incoming signal is GOOD No connection or incorrect connection
Host Port 1 Expansion Port 1 Host Port 2 Expansion Port 2	Amber	ON OFF	Fault No faults

Table 7-1 (continued) I/O Module LEDs

	I/O Module OK	Green	ON OFF	properly booted and functioning correctly internal fault
	I/O Module Fault	Amber	OFF ON	No Faults present I/O Fault
	I/O Module Identity	Blue	ON OFF	Receiving SES Command Not receiving SES Command

Technical Specifications

This appendix contains technical specification information for the InfiniteStorage 11000.

Note: Specifications subject to change without notice.

Configuration, Performance, & Capacity

Host interface	FC-4
Drive interface	SAS
Management interface	RS-232 and Ethernet (Telnet)
Cache memory	2.5GB
Independent drive channels	10
Number of host ports	4
Number of host connections	Up to 512
Number of drives supported	Up to 1250 (1000 drives for data)
Number of LUN groups supported	Up to 1024
Hot spare capability	Yes (up to 125 spare modules)
Full duplex 10/100/1000 BaseT onboard	Yes

Reliability

SES (SCSI Enclosure Services) protocol support	Yes
Temperature monitoring	Yes
Redundant hot-swappable power supplies	2
Redundant hot-swappable cooling fans	3

Physical, Power & Environmental

Controller chassis (H×W×D, Weight)	3.5" × 19.0" × 25", 40 lbs
Electrical/AC	100-120V / 200-240V @ 47-63Hz
Power consumption (maximum current)	4.0A @ 110VAC, 1.9A @ 230VAC Couplet: 8.0A @ 110VAC, 3.8A @ 230VAC
Power consumption (average current)	3.0A @ 110VAC, 1.5A @ 230VAC Couplet: 6.0A @ 110VAC, 3.0A @ 230VAC
Operating environment (temperature / relative humidity)	5°C to 35°C / 20% - 80%, non-condensing
Non-operating environment (temperature/ relative humidity)	-10°C to 50°C / 20% - 80%, non-condensing
Thermal rating (single mode/dual mode)	1500 BTU / 3000 BTU
Certification	UL, CE, CUL, C-Tick, FCC

Drive Addressing Information

This appendix provides information that illustrates the drive addressing for the expansion disk enclosures. There are two drive numbering options:

- 1 x 60 (one set of drives 1 through 60); (Figures B-1 and B-3).
- 2 x 30 (two sets of drives 1 through 30); (Figures B-2 and B-4)

DRIVE 1	DRIVE 13	DRIVE 25	DRIVE 37	DRIVE 49
DRIVE 2	DRIVE 14	DRIVE 26	DRIVE 38	DRIVE 50
DRIVE 3	DRIVE 15	DRIVE 27	DRIVE 39	DRIVE 51
DRIVE 4	DRIVE 16	DRIVE 28	DRIVE 40	DRIVE 52
DRIVE 5	DRIVE 17	DRIVE 29	DRIVE 41	DRIVE 53
DRIVE 6	DRIVE 18	DRIVE 30	DRIVE 42	DRIVE 54
AUX DRIVE	DEM 2B	DEM 2A	DEM 1B	DEM 1A
	DEM 4A	DEM 4B	DEM 3A	DEM 3B
DRIVE 7	DRIVE 19	DRIVE 31	DRIVE 43	DRIVE 55
DRIVE 8	DRIVE 20	DRIVE 32	DRIVE 44	DRIVE 56
DRIVE 9	DRIVE 21	DRIVE 33	DRIVE 45	DRIVE 57
DRIVE 10	DRIVE 22	DRIVE 34	DRIVE 46	DRIVE 58
DRIVE 11	DRIVE 23	DRIVE 35	DRIVE 47	DRIVE 59
DRIVE 12	DRIVE 24	DRIVE 36	DRIVE 48	DRIVE 60

Figure B-1 1 x 60 Drive Layout Configuration

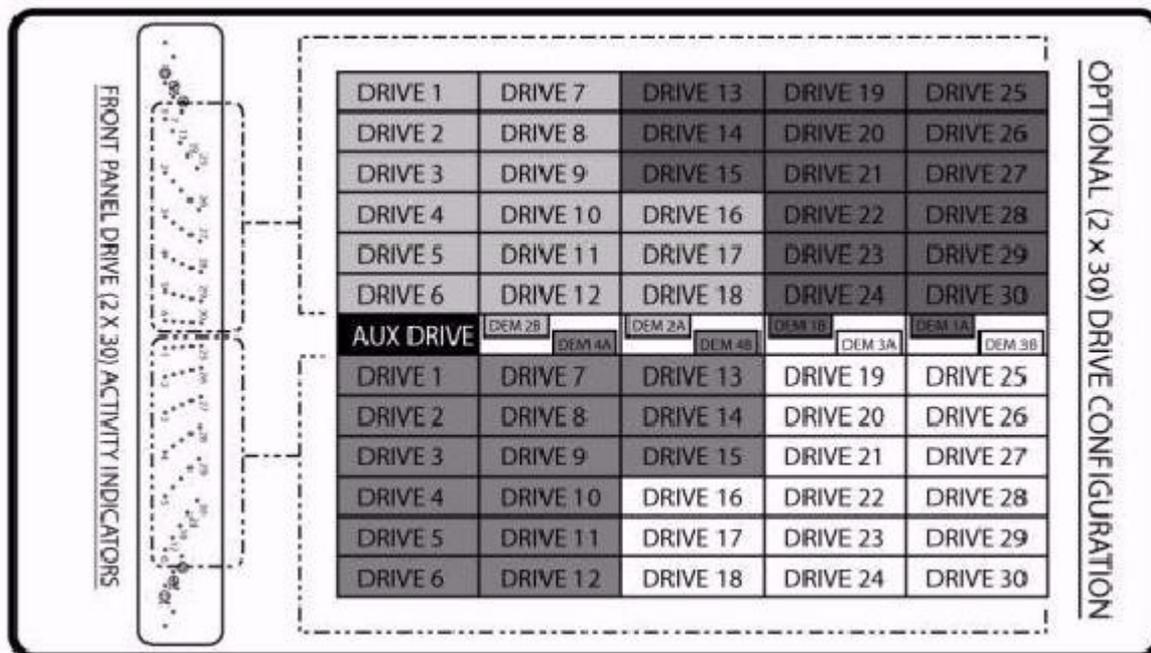


Figure B-2 2 x 30 Drive Layout Configuration

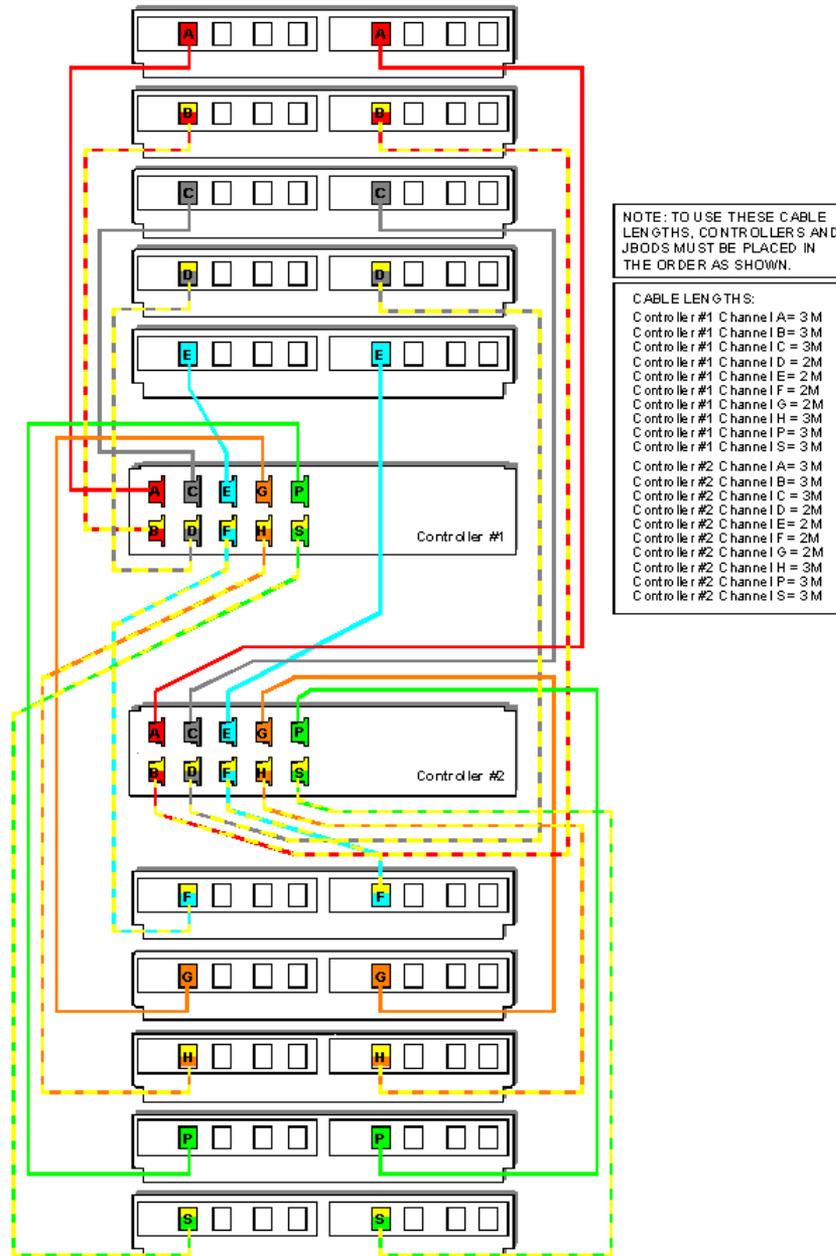


Figure B-3 Controller-Drive Enclosure Couplet (1x Configuration)

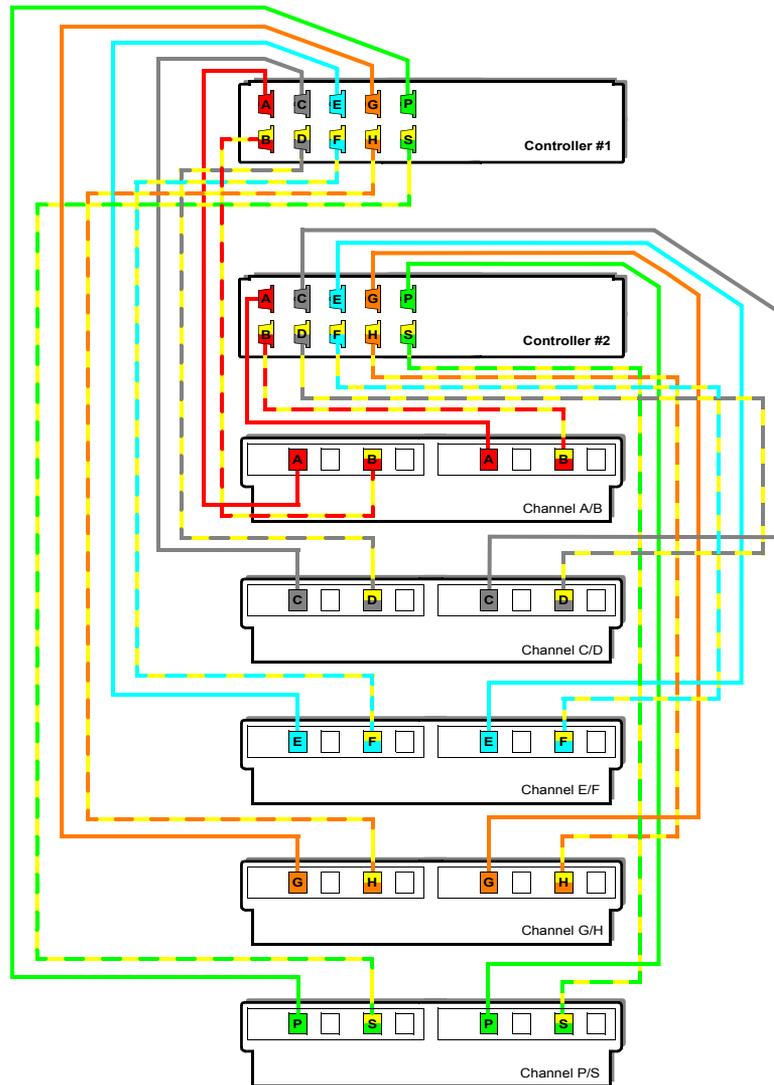


Figure B-4 Controller-Drive Enclosure Couplet (2x Configuration)

Drive Cabling Configuration Examples

This appendix provides information on the proper cabling of the following two configurations:

- 2 controllers and 5 drive enclosures (Figures C-1 and C-2)
- 2 controllers and 10 drive enclosures (Figures C-3 and C-4)

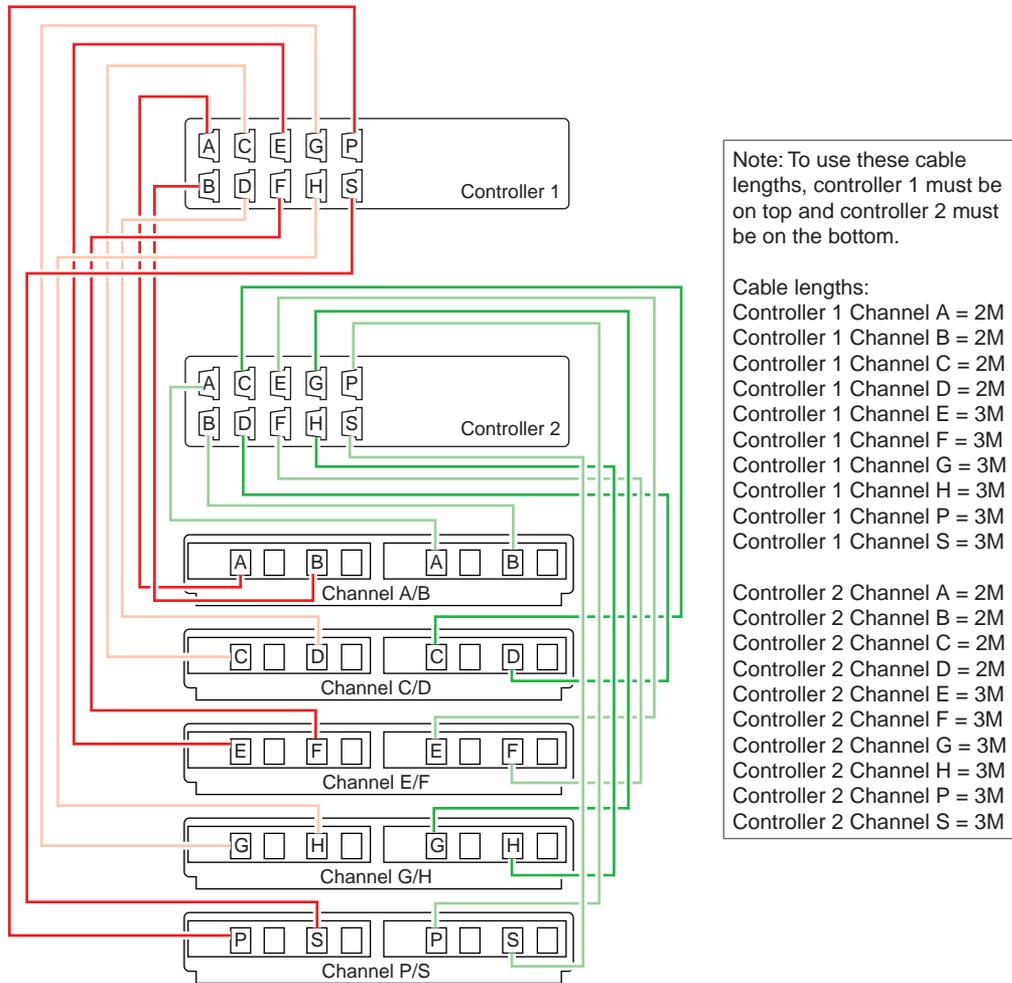


Figure C-1 Controller Cabling (2 Controllers and 5 Drive Enclosures)

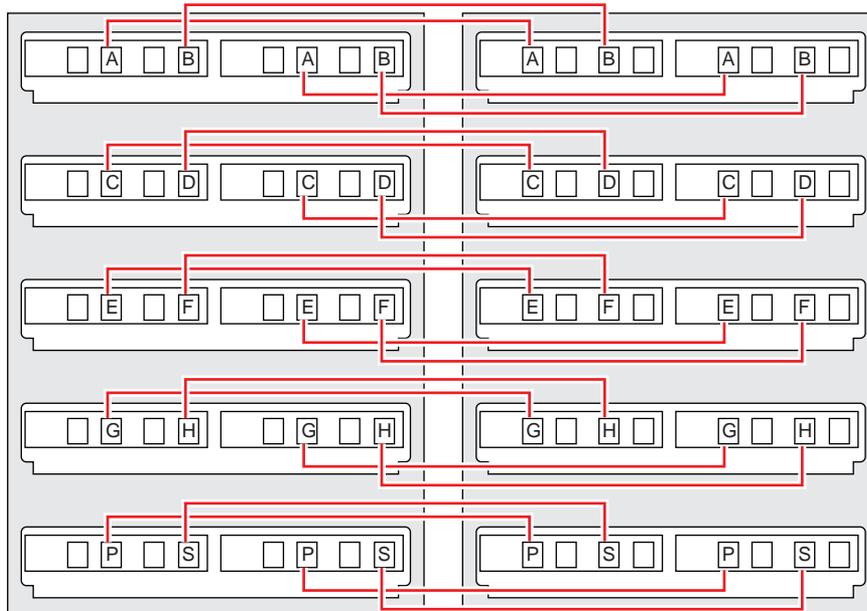


Figure C-2 Drive Enclosure Cabling (2 Controllers and 5 Drive Enclosures)

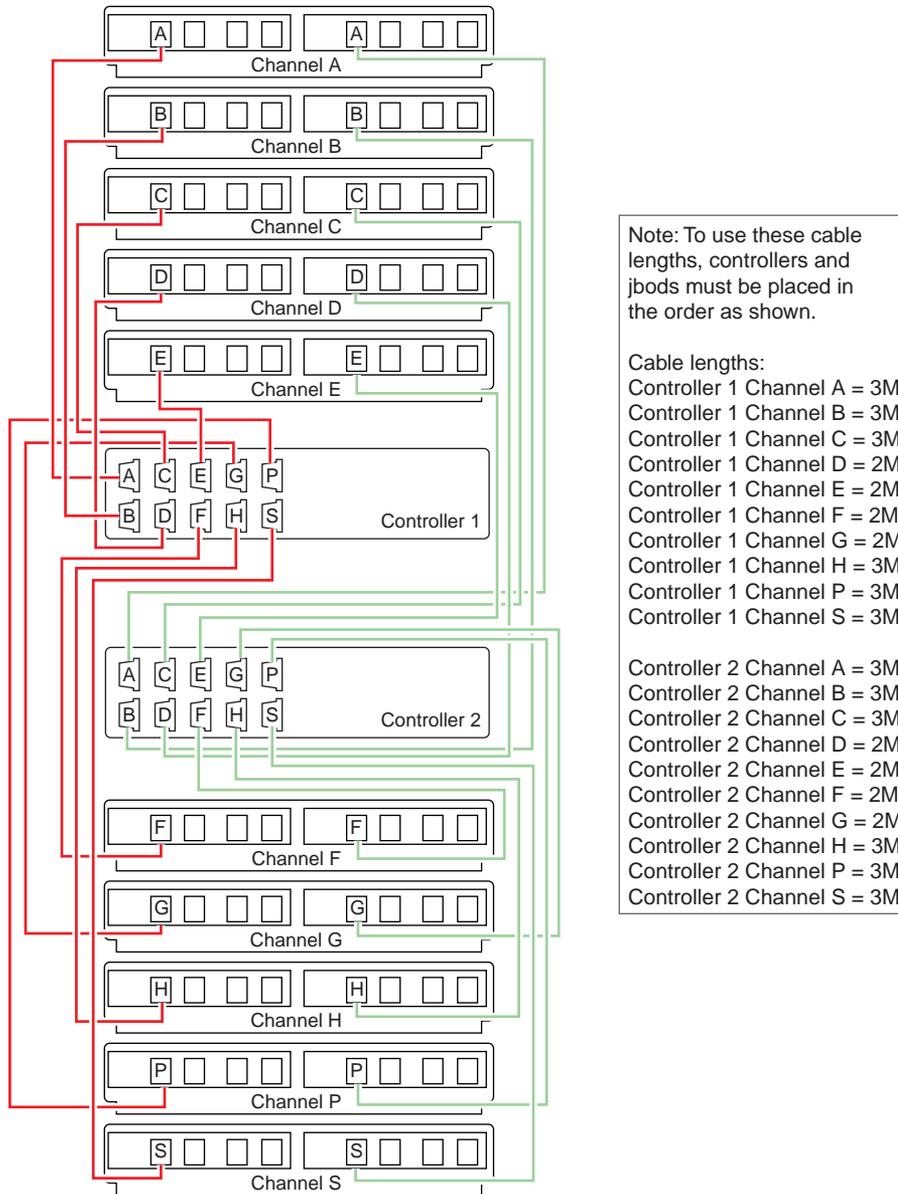


Figure C-3 Controller Cabling (2 Controllers and 10 Drive Enclosures)

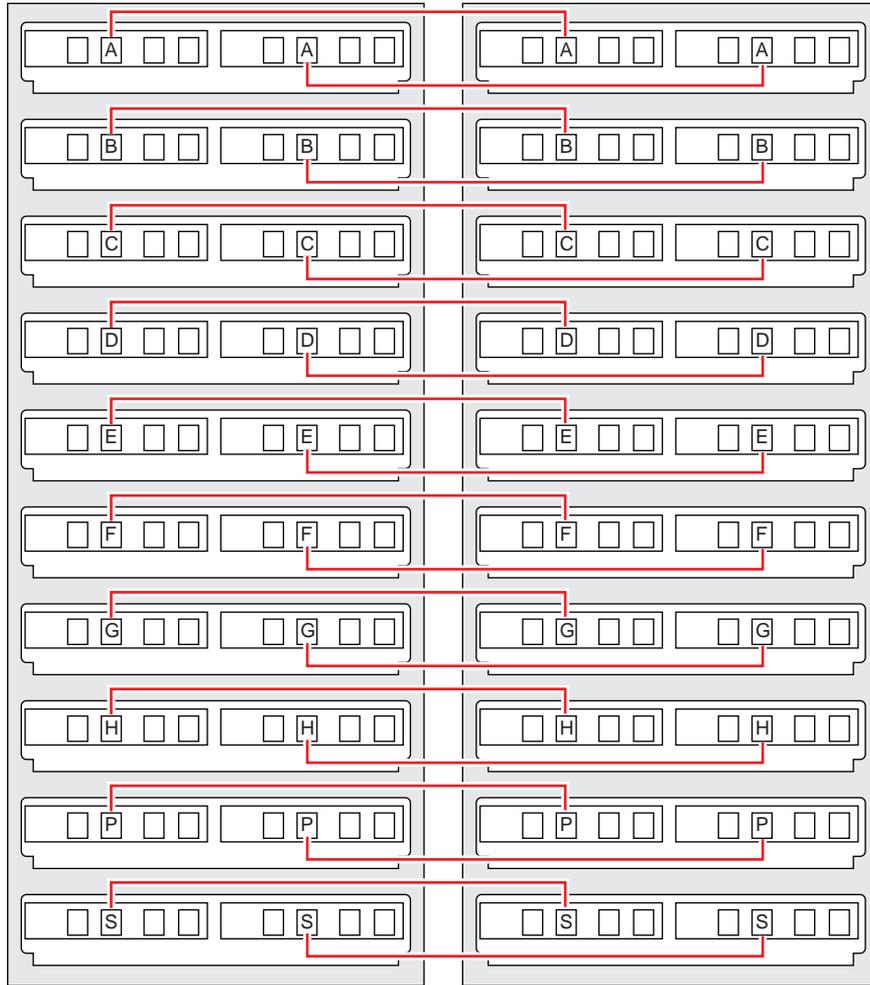


Figure C-4 Drive Enclosure Cabling (2 Controllers and 10 Drive Enclosures)

Safety Guidelines for Rack Installation

Follow these safety guidelines when installing and racking the InfiniteStorage 11000 products.

Elevated Operating Ambient Temperature

If the 11000 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. Therefore, consideration should be given to installing the system in an environment compatible with the system's maximum rated ambient temperature as given in **Appendix A**.

Air Flow

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

Mechanical Loading

The mounting of the system must conform to even and safe mechanical loading. Uneven mechanical loading could cause the rack to tip and fall.

Circuit Overloading

Closely monitor the connection of the system to the supply circuit. Be aware of the effect that the overloading of circuits may have on over-current protection and supply wiring.

Reliable Earthing

Ensure that the earthing (also known as “grounding”) of rack-mounted systems is correct. Closely monitor supply connections other than direct connections to the branch circuit (power distribution units, for example).