

SGI® InfiniteStorage RM610 and RM660
User's Guide

007-4715-002

COPYRIGHT

© 2004, 2005, Silicon Graphics, Inc. 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 Silicon Graphics, Inc.

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., 1500 Crittenden Lane, Mountain View, CA 94043-1351.

TRADEMARKS AND ATTRIBUTIONS

Silicon Graphics, SGI, and the SGI logo are registered trademarks of Silicon Graphics, Inc. All other trademarks mentioned herein are the property of their respective owners.

New Features in This Guide

This revision of the *SGI Guide InfiniteStorage RM610 and RM660 User's Guide* supports the SGI InfiniteStorage release.

Note: Information in the *SA2016 Drive Enclosure for the SGI InfiniteStorage RM610/RM660 User Guide*, 007-4750-001, now resides in Chapters 5, 6, 7, and 8 of this manual.

Major Documentation Changes

Changes in this guide for this release include the following:

- Added information about the SA2016 drive enclosure in Chapter 5, “Drive Enclosure for SGI InfiniteStorage RM610/RM660”.
- Added installation instructions for the SA2016 drive enclosure in Chapter 6, “Installation of the SA2016 Drive Enclosure”.
- Added operating instructions for the SA2016 drive enclosure plug-in modules in Chapter 7, “Operation of SA2016 Enclosure Plug-in Modules”.
- Added troubleshooting information in Chapter 8, “Troubleshooting and Problem Solving”.

Record of Revision

Version	Description
001	October 2004 Original publication
002	June 2005 Updated to support SA2016 SATA disk drive enclosure.

Contents

Figures	xxiii
Tables	xxxi
About This Guide.	xxxiii
What This Guide Contains	xxxiii
International Standards	xxxv
Potential for Radio Frequency Interference	xxxv
USA Federal Communications Commission (FCC)	xxxv
European Regulations	xxxvi
Safety	xxxvi
PSU Warning Label: Power Hazards	xxxvii
Rack System Precautions.xxxviii
Fibre Channel Host & Expansion Connectors	xxxix
ESD Precautions	xxxix
Data Security	xxxix
Special Tools and Equipment	xl
Obtaining Publications	xl
Conventions	xli
Reader Comments.	xli
1. Introducing the SGI InfiniteStorage RM 610/660	1
Introduction	1
The RM 610/RM660 Solution	1
The RM610/RM660	1
Features of the RM610/RM660	2

The RM660 System Hardware	4
Power Supply and Fan Modules.	5
I/O Connectors and Status LED Indicators.	7
Host and Disk Connectivity	8
The RM610 System Hardware	8
Power Supply and Fan Modules.	10
I/O Connectors and Status LED Indicators.	11
Host and Disk Connectivity	12
2. Installing the SGI InfiniteStorage RM610 and RM660	13
Installation Overview	13
Setting Up the SGI InfiniteStorage RM610/RM660.	13
Unpacking the Rack	14
Unpacking the RM610/RM660	16
Rack-Mounting the RM610/RM660 Chassis	17
Connecting the RM Controllers in Dual Mode.	21
Connecting the RM Controllers	22
Connecting the RM660	22
Connecting the RM610	23
Selecting AL_PA for Your Drives	24
Laying Out your Storage Drives	24
Laying Out your Storage Drives on the RM660	24
Laying Out your Storage Drives on the RM610	25
Connecting the RS-232 Terminal.	26
Powering On the RM Controller.	28

Configuring the RM610/RM660	28
Setup Planning	28
Configuration Options	30
Accessing the Administrative Utility	31
Login as Administrator	32
Setting System Time & Date	32
Setting Tier Mapping Mode	32
Verifying Connections for SF6016 Enclosures	33
Checking Tier Configuration	34
Cache Coherency and Labeling in Dual Mode	35
Configuring the Storage Arrays	36
Setting Security Levels	39
User Authentication (Recommended for SAN Environment)	40
Host Port Zoning (Anonymous Access)	43
3. Using the SGI InfiniteStorage RM610/RM660 Management and Administrative Facilities	45
Managing the SGI InfiniteStorage RM610/RM660	45
Management Interface	45
Locally via Serial Interface	45
Remotely via Telnet	47
Available Commands	47
Administrator and User Logins	49
Login	50
Logout	50
Password	50
Who Am I	51

Configuration Management 51
Configure and Monitor Status of Host Ports 51
Host ID. 52
Host WWN 53
Host Status. 54
Host LILP Payload 54
Host Array Parity Checking 55
Host Port Speed 55
Configure and Monitor Status of Storage Assets 55
Disk and Channel Information 55
Tier View 57
LUN View 60
Adding/Removing Storage Assets 62
Status of Drive Enclosures 63
Display Information of SES Devices 63
Visual Indication of Drive 63
Visual Indication of Tier 64
Visual Indication of Channel 64
Visual Indication of LUN 64
Verify Connections to SF6016 Enclosures 64
Tier Mapping for Enclosures 64
System Network Configuration 66
Telnet 68
SNMP & Syslog 68
API Server Connections 69
Displaying and Editing the Routing Table. 70
Restarting the SGI InfiniteStorage RM610/RM660 71
System Restart. 71
System Shutdown. 71
Setting the System's Date and Time 72
Saving the RM Controller's Configuration 72
Restoring the System's Default Configuration 73
LUN Management. 73

Configuring the Storage Array	74
Creating and Formatting a LUN	74
Interrupting a LUN Format Operation	75
Changing a LUN Label	75
Moving a LUN (Dual Mode Only)	76
Deleting a LUN	77
SCSI Reservations	77
Automatic Drive Rebuild	77
Manual Drive Rebuild	78
Drive Rebuild Verify.	78
Manual Drive Replace	79
Interrupting a Rebuild Operation	79
Couplet RM Controller Configuration (Cache/Non-Cache Coherent)	79
Non-Cache Coherent (firmware version 1.06 and above)	80
Cache Coherent (firmware version 2.02 and above)	80
Fail / Restore the Other RM Controller Unit in the Couplet Pair	81
Labeling the RM Controller Unit(s)	81

Performance Management 82
Optimization of I/O Request Patterns 82
Display Current Cache Settings 82
Cache Segment Size 83
Writeback Cache Settings 83
Prefetch Settings 84
Cache Settings Reset 84
Disk Configuration Settings 85
Audio/Visual Settings of the System 85
Locking LUN in Cache 88
Locking / Unlocking a LUN. 89
System Performance Statistics 90
Resources Allocation 96
Background Format/Rebuild Operations 96
Background LUN Verify Operations 97
SES Device Monitoring Rate. 98
Host Command Timeout. 98
Security Administration 99
Monitoring User Logins	100
Zoning (Anonymous Access).	101
User Authentication	102
Firmware Update Management	103
Displaying Current Firmware Version	103
Firmware Update Procedure	104
Remote Login Management	105
When a Telnet Session is Active	106
The System Logs	108
Message Log	108
System and Drive Enclosure Faults	109
Displaying System's Uptime	109
Saving a Comment to the Log	110

Other Utilities110
APC UPS SNMP Trap Monitor110
API Server Connections110
Internal Mirrored Groups (IMG)111
Changing Baud Rate for CLI Interface113
CLI/Telnet Session Control Settings113
Remote Management of the RM Controller114
Network Connection114
Network Interface Set Up115
Login Names and Passwords118
4. Supporting the SGI InfiniteStorage RM610/RM660119
Maintaining the SGI InfiniteStorage RM610 and RM660119
Component Failure Recovery119
Power Supply Failure119
Fan Failure120
Recovering from Drive Failures.121
Single Drive Failures.122
Multiple Drive Failures and Channel Failures124
Component Failure on Enclosures125
RM Controller Messages and Descriptions126
Fault Messages127
Warning Messages127
Timeout Messages128
Informational Messages128
Disk Boot-up Error Messages130
Disk Error Messages134
Host Error Messages139
Rebuild Error Messages141
SES (SCSI Enclosure Services) Messages142
Other Messages143

5.	Drive Enclosure for SGI InfiniteStorage RM610/RM660	145
	Introduction to the SA2016 System	145
	Enclosure Chassis	146
	The Plug-in Modules	147
	Power Supply/Cooling Module	148
	Multiple Power Supply/Cooling Modules	150
	Operators Panel	150
	Ops Panel Indicators and Switches	151
	LRC Input/Output Module	153
	Drive Carrier Module	156
	SATA Dual Path Transition Card	157
	Drive Status Indicators	157
	Anti-tamper Locks	158
	Dummy Carrier Modules	158
	Blank Modules	158
	Visible and Audible Alarms	159
	SA2016 System Technical Specification	159
	Dimensions	159
	Weight	160
	AC Power (450W PSU)	160
	PSU Safety and EMC Compliance	160
	Power Cord	160
	Environment	161
	Interfaces	162
	LRC I/O Module Specification	162
	Drive Carrier Module Specification	163
	Software Enclosure Services (SES) Support	163
6.	Installation of the SA2016 Drive Enclosure	165
	Planning Your Installation	165
	Enclosure Bay Numbering Convention	167

Enclosure Installation Procedures168
Pre-Requisites168
Rack Mounting Rail Kit168
Parts Check List169
Chassis Installation169
Parts Check List169
Procedure for Chassis Installation169
Power Supply/Cooling Module Installation169
Parts Check List170
Procedure for Power Supply/Cooling Installation170
LRC I/O Module Configurations172
Internal Loop Structures172
FC-AL Interface172
Connecting Multiple Enclosures173
Connection SA2016 Enclosures on the RM610.174
Connection SA2016 Enclosures on the RM660.178
Daisy-Chaining the SA2016 Enclosures on the RM610184
Daisy-Chaining the SA2016 Enclosures on the RM660186
LRC I/O Module Installation190
Parts Check List190
Procedure for LRC I/O Module Installation190
Drive Enclosure Device Addressing191
Drive Carrier Configuration194
Planning and Configuring Your Installation194
System Configuration194
Drive Configuration194
Drive Carrier Installation195
Parts Check List195
Procedure for Drive Carrier Installation195
Dummy Carrier Modules199
Engaging the Anti-tamper Locks199
Activating the Locks199
Deactivating the Locks200

Power Cord Connection	200
Parts Check List	200
Procedure	200
Grounding Checks	200
7. Operation of SA2016 Enclosure Plug-in Modules	203
Power On	203
Power Supply/Cooling Module LEDs	204
Ops Panel LEDs.	205
Starting the Drives	206
Disk Drives LEDs	207
Power Down an Enclosure	207
8. Troubleshooting and Problem Solving.	209
Initial Start-up Problems	209
Faulty Cords	209
Alarm Sounds On Power Up.	210
Green “Signal Good” LED on LRC Not Lit	210
Computer Doesn’t Recognize the SA2016 Subsystem	210
LEDs	211
Power Supply/Cooling Module	211
Ops Panel	211
Ops Panel LEDs	213
LRC I/O Module LEDs	213
Audible Alarm	214
Audible Alarm Mute	214
LED Test Mode	215
Troubleshooting	215
System Faults	215
Power Supply/Cooling Faults	216
Thermal Control	217
Thermal Alarm	218
Thermal Shutdown	218

Drive Carrier Module Faults219
Dummy Carrier Modules219
Dealing with Hardware Faults219
Continuous Operation During Replacement220
Replacing a Module220
Power Supply/Cooling Modules220
Removing a Power Supply/Cooling Module.221
Inserting the Power Supply/Cooling Module221
Ops Panel224
LRC I/O Module224
Removing the the LRC I/O Module225
Inserting the the LRC I/O Module.226
Drive Carrier Module227
Removal and Replacement227
Spare Parts and Ancillary Items.228
A. Technical Specifications for SGI InfiniteStorage RM660 and RM610229
SGI InfiniteStorage RM660 Technical Specification229
SGI InfiniteStorage RM610 Technical Specifications231
B. Safety Guidelines for SGI InfiniteStorage RM610/RM660 Rack Installation233
C. Using the SFx016 Drive Enclosures.235
The Drive Enclosure235
Enclosure Bay Numbering Convention236
Operators Panel Module237
Power Supply/Cooling Module240
I/O Modules241
Drive Modules (FC-AL)244
Drive LEDs245
Anti-Tamper Locks245
Dummy Drive Modules and Blank Modules246
Visible and Audible Alarms247

Setting Up the Drive Enclosures	247
Installing the Drive Enclosures in Rack	248
Verify Configuration Switch Settings on Ops Panel Modules	253
Cable Connections and Enclosure ID Settings on the RM660.	256
Connecting SF6016 Enclosures on the RM660	256
Connecting SF4016 Enclosures on the RM660	259
Daisy-Chaining the SF4016 Enclosures on the RM660	263
Connecting SF2016 Enclosures on the RM660	265
Daisy-Chaining the SF2016 Enclosures on the RM660	268
Cable Connections and Enclosure ID Settings on the RM610.	270
Connecting SF6016 Enclosures on the RM610	270
Connecting SF4016 Enclosures on the RM610	273
Daisy-Chaining the SF4016 Enclosures on the RM610	275
Connecting SF2016 Enclosures on the RM610	277
Daisy-Chaining the SF2016 Enclosures on the RM610	279
Installing Drive Modules	281
Powering on the Drive Enclosures	284
Verifying Connections for SF6016 Enclosures	284
Enclosure AL_PA and Tier Mapping Charts	295

Troubleshooting & Problem Solving308
RM660 Doesn't Recognize the Drive Enclosure308
LED Test Mode308
Status LED Indicators309
Power Supply/Cooling Module309
Ops Panel309
Troubleshooting311
System Faults312
Power Supply/Cooling Faults312
Thermal Control313
Thermal Alarm314
Thermal Shutdown315
FC-AL Drive Module Faults315
Dummy Drive Modules316
Auto Start Failure316
Dealing with Hardware Faults316
Continuous Operation During Replacement317
Replacing a Power Supply/Cooling Module317
Replacing the Ops Panel319
Replacing the I/O Module319
Replacing a Drive Module320
SFx016 Enclosure Technical Specifications321
Dimensions321
Weight321
Power Supply Unit (PSU)322
Power Cord322
Environment322
Interfaces323
FC-AL LRC I/O Module Specification324
Drive Module Specification324
Software Enclosure Services (SES) Support324
Spare Parts and Accessories325

D. Complete List of Commands at OEM Level	327
Available Commands:	327
ALARM.	329
APC_UPS	330
API	331
AV	332
CACHE.	337
COMMENT	340
CONSOLE	341
DATE	342
DEBUG	344
DEFAULTS.	345
DISK.	346
DISK_CLI	352
DUAL	354
FAULTS.	356
HELP	358
HOST	359
LICENSE	362
LOG.	364
LOGIN	367
LOGOUT	368
LUN.	369
MIRROR	373
NETWORK.	375
PASSWORD	379
RESTART	380
ROUTE	381
SAVE	383
SES	384
SETTINGS	390
SHOWALL.	391
SHUTDOWN	398

STATS399
TELNET403
TFTP405
TIER406
TIME409
UPTIME410
USER411
VERSION413
WHOAMI414
ZONING415
Glossary417
Index425

Figures

Figure 1-1	The RM660 Front and Rear Views	5
Figure 1-2	Power Supply/Cooling Module	6
Figure 1-3	Fan Module	7
Figure 1-4	I/O Ports on Rear Panel of the RM660	7
Figure 1-5	The RM610 Front & Rear Views	9
Figure 1-6	Power Supply/Cooling Module	10
Figure 1-7	Fan Module	11
Figure 1-8	I/O Ports on Rear Panel of RM610	11
Figure 2-1	Remove Anchoring Brackets	15
Figure 2-2	Creating the Ramps for the Pallet	15
Figure 2-3	Attach Trims to Bottom of Rack	16
Figure 2-4	Opening the Rack Door	16
Figure 2-5	RM610/RM660 8x00 Mounting Hardware	18
Figure 2-6	Attaching the Mounting Brackets to Front of the RM660	18
Figure 2-7	Attaching the Mounting Brackets to the Front of the RM610	19
Figure 2-8	Square Hole Rack with Cage Nuts Installed.	19
Figure 2-9	Rack with Rear Mounting Bracket Installed.	20
Figure 2-10	Securing Front of the RM660 in the Rack	21
Figure 2-11	Securing Front of the RM610 in the Rack	21
Figure 2-12	I/O Connectors on Back of the RM660	22
Figure 2-13	I/O Connectors on Back of the RM610	23
Figure 2-14	Channel and Tier Numbering of Drives on SF4016 Enclosures	25
Figure 2-15	Channel and Tier Numbering of Drives on SF4016 Enclosures	26
Figure 2-16	CLI Port on the RM660.	27
Figure 2-17	CLI Port on the RM610.	27
Figure 2-18	Striping Across Tiers	29
Figure 2-19	CLI Port on the RM660.	31

Figure 2-20	CLI Port on the RM610 31
Figure 2-21	Tier Mapping Mode Setup Screen. 33
Figure 2-22	Current Tier Configuration Screen 34
Figure 2-23	Dual RM Controller Configuration Screen 35
Figure 2-24	Labeling an RM Controller Unit 36
Figure 2-25	Logical Unit Status - Formatting 38
Figure 2-26	Logical Unit Status - Ready 38
Figure 2-27	LUN Zoning Screen. 40
Figure 2-28	Configuring the User Access to LUNs 42
Figure 2-29	Current Security Settings Screen 43
Figure 2-30	Zoning Configuration Screen 44
Figure 3-1	CLI Port at Back of the RM660. 46
Figure 3-2	CLI Port at back of the RM610 46
Figure 3-3	Login Screen 50
Figure 3-4	Password Configuration Screen 51
Figure 3-5	“Who Am I” Screen 51
Figure 3-6	Host Ports Status Screen 52
Figure 3-7	Host Port Status Screen. 54
Figure 3-8	Host Port LILP Payload Screen 54
Figure 3-9	RM660 Disk Channel Status Screen 55
Figure 3-10	Disk Status Screen 56
Figure 3-11	Disk Defect List Screen 57
Figure 3-12	Tier Configuration Screen 58
Figure 3-13	Displaying the Tier Configuration Information 59
Figure 3-14	Tier Configuration Information (1) 60
Figure 3-15	Tier Configuration Information (2) 60
Figure 3-16	Logical Unit Status Screen 61
Figure 3-17	LUN List Screen 61
Figure 3-18	LUN Configuration Screen. 62
Figure 3-19	LUN Reservations Screen 62
Figure 3-20	Displaying the Current Disk Enclosure Failures. 63
Figure 3-21	Displaying the Current Tier Mapping Mode. 65
Figure 3-22	Displaying the Current Mapping AL_PA Values for Tiers 65

Figure 3-23	Tier Mapping Mode Setup Screen	66
Figure 3-24	Current Network Configuration Screen	67
Figure 3-25	Current RM Controller Routing Table	70
Figure 3-26	Saving System Parameters Screen	72
Figure 3-27	Current System Parameters Status Screen	73
Figure 3-28	Logical Unit Status Screen	74
Figure 3-29	Changing a LUN Label Screen	76
Figure 3-30	Moving a LUN Screen	76
Figure 3-31	Automatic Disk Rebuilding Parameter	78
Figure 3-32	Couplet RM Controller Configuration Screen	79
Figure 3-33	Labeling an RM Controller Unit	82
Figure 3-34	Current Cache Settings Screen	83
Figure 3-35	RM660 Disk Configuration Setting Screen	85
Figure 3-36	Current Audio/Visual Settings	86
Figure 3-37	Logical Unit Status - LUN Locked in Cache.	90
Figure 3-38	System Performance Statistics Screen	91
Figure 3-39	Command Delay Statistics Screen	92
Figure 3-40	Host Delay Statistics Screen	93
Figure 3-41	Tier Delay Statistics Screen	93
Figure 3-42	Host Command Offsets Screen	94
Figure 3-43	Dual Message Statistics Screen	95
Figure 3-44	Command Length Statistics Screen	95
Figure 3-45	Host Command Offsets Screen	96
Figure 3-46	Displaying the Current Rebuild Parameters	96
Figure 3-47	SES Device Monitoring Rate	98
Figure 3-48	Mapping Internal LUN's to External LUN's	99
Figure 3-49	User Login Messages	100
Figure 3-50	User Connections Screen	100
Figure 3-51	Users' LUN Mapping Information Screen	101
Figure 3-52	Current Zoning Configuration	101
Figure 3-53	Edit Zoning Configuration Screen	102
Figure 3-54	Current LUN Mapping Screen	103
Figure 3-55	Version Information Screen	104

Figure 3-56	Downloading RM controller Firmware	105
Figure 3-57	Telnet Statistics	106
Figure 3-58	Telnet Session Initiated	106
Figure 3-59	Telnet Session Information	107
Figure 3-60	Terminating a Telnet Session	107
Figure 3-61	Telnet Session Being Terminated	108
Figure 3-62	Current System Faults	109
Figure 3-63	Display System Uptime.	109
Figure 3-64	Displaying APC UPS SNMP Trap Monitor Status	110
Figure 3-65	Displaying Status of API Connections	110
Figure 3-66	API Server Connection Statistics	111
Figure 3-67	Current Internal Mirrored Groups	112
Figure 3-68	Displaying the Serial Console Setting.	113
Figure 3-69	Changing the Baud Rate	113
Figure 3-70	Current Session Control Settings	113
Figure 3-71	Telnet Port on the RM660	114
Figure 3-72	Telnet Port on the RM610	115
Figure 3-73	Current Network Configuration Screen	115
Figure 3-74	Current RM Controller Routing Table	117
Figure 4-1	Power Supply/Cooling Modules in the Front of the RM660	119
Figure 4-2	Power Supply/Cooling Modules in the Front of the RM610	120
Figure 4-3	Fan Module	121
Figure 4-4	Checking the Disk Status	122
Figure 4-5	Obtaining the Disk Information	123
Figure 4-6	Checking the Tier Status	124
Figure 4-7	Disk Channel Information	125
Figure 4-8	Example of SES Messages	126
Figure 4-9	Displaying the Current Disk Enclosure Failures.	126
Figure 5-1	The SA2016 System	145
Figure 5-2	Enclosure Chassis (Front)	147
Figure 5-3	Enclosure Chassis (Rear)	147
Figure 5-4	Power Supply Cooling Module	149
Figure 5-5	Ops Panel	151

Figure 5-6	LRC I/O Module154
Figure 5-7	LRC Front Panel (as viewed from rear of Enclosure)155
Figure 5-8	SA2016 System Enclosure with LRC I/O Modules Installed156
Figure 5-9	Drive Carrier Module157
Figure 5-10	Anti-tamper Lock158
Figure 6-1	Module locations167
Figure 6-2	AC Power Supply/Cooling Module - Handle in Open Position171
Figure 6-3	Installing an AC Power Supply Cooling Module (1)171
Figure 6-4	Installing an AC Power Supply Cooling Module (2)172
Figure 6-5	SA2016 Enclosures Connection to RM610 Unit #1176
Figure 6-6	SA2016 Enclosures Connection RM610 Unit #2.177
Figure 6-7	SA2016 Enclosures Connected to RM660 Unit #1 (1)180
Figure 6-8	SA2016 Enclosures Connected to RM660 Unit #1 (2)181
Figure 6-9	SA2016 Enclosures Connected to Couplet RM660 Unit #2 (1).182
Figure 6-10	SA2016 Enclosures Connected to Couplet RM660 Unit #2 (2).183
Figure 6-11	Daisy-Chaining the SA2016 Enclosures to RM610 Unit#1185
Figure 6-12	Daisy-Chaining the SA2016 Enclosures to Couplet RM610 Unit #2186
Figure 6-13	Daisy-Chaining the SA2016 Enclosures to RM660 Unit #1.188
Figure 6-14	Daisy-Chaining the SA2016 Enclosures to Couplet RM660 Unit #2189
Figure 6-15	LRC Module Latch Operation190
Figure 6-16	Installing an LRC I/O Module in Bay 4.191
Figure 6-17	Installing a SATA Drive Carrier Module (1).196
Figure 6-18	Installing a SATA Drive Carrier Module (2).197
Figure 6-19	Installing an SATA Drive Carrier Module (3)198
Figure 6-20	Activating the Anti-tamper Lock199
Figure 7-1	Ops Panel LEDs and Switches205
Figure 8-1	Ops Panel212
Figure 8-2	Removing/Inserting an AC Power Supply/Cooling Module (1).222
Figure 8-3	Removing/Inserting an AC Power Supply/Cooling Module (2).223
Figure 8-4	Removing/Inserting an AC Power Supply/Cooling Module (3).224
Figure 8-5	Removing an LRC I/O Module (1)225
Figure 8-6	Removing an LRC I/O Module (2)226
Figure 8-7	LRC I/O Module Latch Operation227

Figure C-1	The SFx016 Enclosure - Front and Rear Views	236
Figure C-2	Enclosure Bay Numbering Convention	237
Figure C-3	Ops Panel Module	238
Figure C-4	Power Supply/Cooling Module	241
Figure C-5	FC-AL LRC I/O Module	242
Figure C-6	Overall Loop Layout	243
Figure C-7	I/O Module Connectors and LEDs	244
Figure C-8	Drive Module	244
Figure C-9	LEDs on Drive Module Front	245
Figure C-10	Anti-Tamper Lock on Drive Modules.	246
Figure C-11	Assembling Rail to Support (Left Hand Assembly).	248
Figure C-12	Fastening Rail/Support Assembly to Square/Through-Hole Rack (Left Hand Side)	249
Figure C-13	Fastening Rail/Support Assembly to Threaded-Hole Rack (Left Hand Side) 250	250
Figure C-14	Latch Operation on Drive Module	250
Figure C-15	Securing Chassis to Square-Through-Hole Rack (Left Hand)	251
Figure C-16	Securing Chassis to Threaded-Hole Rack (Left Hand)	252
Figure C-17	Settings on SF6016 Enclosures	253
Figure C-18	Settings on SF4016 Enclosures	254
Figure C-19	Settings on SF2016 Enclosures	255
Figure C-20	Drive Locations on SF6016 with Enclosure ID Set to 1	256
Figure C-21	Enclosure Label at Rear.	256
Figure C-22	Connecting the SF6016 Enclosures to RM660	257
Figure C-23	Daisy-Chaining the SF6016 Enclosures	258
Figure C-24	Enclosure ID Switch.	259
Figure C-25	Connecting the SF4016 Enclosures to the RM660 (1)	260
Figure C-26	Connecting the SF4016 Enclosures to the RM660 (2)	262
Figure C-27	Daisy-Chaining the SF4016 Enclosures	264
Figure C-28	Connecting the SF2016 Enclosures to the RM660 (1)	266
Figure C-29	Connecting the SF2016 Enclosures to the RM660 (2)	267
Figure C-30	Daisy-Chaining the SF2016 Enclosures	269
Figure C-31	Drive Locations on the FS6016 with Enclosure ID Set to 1	270

Figure C-32	Enclosure Label at Rear270
Figure C-33	Connecting the SF6016 Enclosure to the RM610271
Figure C-34	Daisy-Chaining the SF6016 Enclosures272
Figure C-35	Enclosure ID Switch272
Figure C-36	Connecting the SF4016 Enclosures to the RM610(1)273
Figure C-37	Connecting the SF4016 Enclosures to the RM610274
Figure C-38	Daisy-Chaining the SF4016 Enclosures276
Figure C-39	Connecting the SF2016 Enclosures to the RM610278
Figure C-40	Daisy-Chaining the SF2016 Enclosures280
Figure C-41	Drive Bays on SF4016 Enclosures for SES Communication Paths.	.281
Figure C-42	Drive Bays on SF2016 Enclosures for SES Communication Paths.	.281
Figure C-43	Opening the Drive Module Handle282
Figure C-44	Inserting the Drive Module282
Figure C-45	Activating the Anti-Tamper Lock283
Figure C-46	PSU LEDs309
Figure C-47	Ops Panel LEDs.310
Figure C-48	Disengage Latches on PSU Module318
Figure C-49	PSU Modules in Rear Panel318
Figure C-50	LRC I/O Module Latch Operation319
Figure C-51	LRC I/O Modules Orientation320

Tables

Table 3-1	List of Available Commands	47
Table 3-2	Loop ID to AL_PA Correspondence	53
Table 5-1	Ops Panel Switch Functions152
Table 5-2	LRC I/O Module LEDs153
Table 5-3	Ambient Temperature and Humidity161
Table 6-1	SA2016 Configuration166
Table 6-2	Configuration Options173
Table 6-3	Ops Panel Switch Functions192
Table 6-4	Mode 1 Drive Addressing192
Table 6-5	Drive Slot Arrangement: Enclosure Front View193
Table 7-1	PSU LEDs204
Table 7-2	Ops Panel LED States206
Table 8-1	PSU LEDs211
Table 8-2	Ops Panel LED States213
Table 8-3	LRC I/O Module LEDs213
Table 8-4	LED Functions219
Table C-1	Ops Panel Switch Functions239
Table C-2	SF4016 Enclosure AL_PA and Tier Mapping Chart for RM660296
Table C-3	SF2016 Enclosure AL_PA and Tier Mapping Chart for RM660302
Table C-4	SF4016 Enclosure AL_PA and Tier Mapping Charts for RM610303
Table C-5	SF2016 Enclosure AL_PA and Tier Mapping Charts for RM 610307
Table C-6	Ops Panel LED States310
Table C-7	LED Functions316

About This Guide

This publication provides information about the SGI InfiniteStorage RM610 and RM660.

Chapters 5 through Chapter 8 of this guide provides you with step-by-step instructions on how to install, configure and connect the SA2016 storage enclosures to your host computer system, and how to use and maintain the systems.

This guide is for users that have a working knowledge of the Fibre Channel Arbitrated Loop (FC-AL) environment into which you are installing the SA2016 system. If you do not have these skills, or are not confident with the instructions in this guide, do not proceed with the installation.

Note: Information in the *SA2016 Drive Enclosure for the SGI InfiniteStorage RM610/RM660 User Guide*, 007-4750-001, now resides in Chapters 5, 6, 7, and 8 of this manual.

What This Guide Contains

The manual structure is outlined below:

- Chapter 1, “Introducing the SGI InfiniteStorage RM 610/660”
Provides information on features and functions of the SGI InfiniteStorage RM610 and RM660. Topics covered in this chapter include the features of the RM610/RM660 and the hardware components of the RM610/RM660.
- Chapter 2, “Installing the SGI InfiniteStorage RM610 and RM660”
Contains information regarding installation of the RM controllers. **Procedures described in this chapter of the manual should only be performed by SGI trained personnel.** Topics covered in this section include an installation overview, information regarding hardware installation, and information on configuring the RM controller.:

- Chapter 3, “Using the SGI InfiniteStorage RM610/RM660 Management and Administrative Facilities”
Provides descriptions of the RM controller’s comprehensive management capability as well as instructions on how to set up the Telnet and SNMP functions of the RM controller for remote monitoring and configuration.
- Chapter 4, “Supporting the SGI InfiniteStorage RM610/RM660”
Contains error recovery and system support information. Topics covered in this section include the procedures for how to recover from component failures and drive failures on the RM controllers and descriptions of the RM controller messages.
- Chapter 5, “Drive Enclosure for SGI InfiniteStorage RM610/RM660”
Describes the SA2016 storage enclosures.
- Chapter 6, “Installation of the SA2016 Drive Enclosure”
Describes how to install the SA2016 Enclosure and plug-in modules.
- Chapter 7, “Operation of SA2016 Enclosure Plug-in Modules”
Describes the operation of the SA2016 Enclosure plug-in modules.
- Chapter 8, “Troubleshooting and Problem Solving”
Chapter 8 provides troubleshooting and problem solving information for the SA2016 Enclosure plug-in modules.
- Appendix A, “Technical Specifications for SGI InfiniteStorage RM660 and RM610”
Technical specifications for the SGI InfiniteStorage RM660 and RM610.
- Appendix B, “Safety Guidelines for SGI InfiniteStorage RM610/RM660 Rack Installation”
Provides the safety guidelines that should be followed when installing the RM controllers in a rack.
- Appendix C, “Using the SFx016 Drive Enclosures”
Provides information on installing the SF6016, SF4016, and SF2016 drive enclosures. Features of the enclosures are also described in this appendix.
- Appendix D, “Complete List of Commands at OEM Level”
Covers the Command Line Interface (CLI) commands for OEM (Original Equipment Manufacturer) access. Description and usage examples are given for each command.

Additionally, this manual contains a glossary containing definitions of terms that are used throughout the manual and an index to help you locate information quickly.

International Standards

The SA2016 storage systems complies with the requirements of the following agencies and standards:

- CE to IEC 950/EN60950
- UL
- cUL

Potential for Radio Frequency Interference

USA Federal Communications Commission (FCC)

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

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


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

European Regulations

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

Safety

All plug-in modules are part of the fire enclosure and must only be removed when a replacement can be immediately added. Do not run system without all units in place.



- In order to comply with applicable safety, emission and thermal requirements no covers should be removed and all bays must be fitted with plug-in modules.




LRC Module Caution Label: Do not operate with modules missing



PSU/Cooling Module Caution Label: Do not operate with modules missing

- The enclosure must only be operated from a power supply input voltage range of 100 -120VAC or 200-240 VAC.
- The plug on the power supply cord is used as the main disconnect device. Ensure that the socket outlets are located near the equipment and are easily accessible.
- The equipment is intended to operate with two working PSUs.
- If powered by multiple AC sources, disconnect all supply power for complete isolation



PSU Warning Label: Power Hazards

- The power connection must always be disconnected prior to removal of the PSU/Cooling module from the enclosure.
- A safe electrical earth connection must be provided to the power cord. Check the grounding of the enclosure before applying power.
- Provide a suitable power source with electrical overload protection to meet the requirements laid down in the technical specification.

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



PSU Safety Label: Electric Shock Hazard Inside

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

- A faulty Power Supply/Cooling module must be replaced with a fully operational module within 24 hours.

Rack System Precautions

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

- The rack design should incorporate stabilizing features suitable to prevent the rack from tipping or being pushed over during installation or in normal use.
- When loading a rack with the units, fill the rack from the bottom up and empty from the top down.
- The system must be operated with low pressure rear exhaust installation (back pressure created by rack doors and obstacles not to exceed 5 pascals [0.5mm water gauge]).
- The rack design should take into consideration the maximum operating ambient temperature for the unit, which is 40°C when dual cooling modules are fitted.
- The rack should have a safe electrical distribution system. It must provide overcurrent protection for the unit and must not be overloaded by the total number of units installed in the rack. Consideration of the units nameplate rating should be used when addressing these concerns.

- The electrical distribution system must provide a reliable earth for each unit and the rack.
- Each power supply in each unit has an earth leakage current of 1.8mA. The design of the electrical distribution system must take into consideration the total earth leakage current from all the power supplies in all the units. The rack will require labelling with “HIGH LEAKAGE CURRENT. Earth connection essential before connecting supply”.
- The rack when configured with the units must meet the safety requirements of UL 60950 and IEC 60950.

Fibre Channel Host & Expansion Connectors

Note: If fitted with Optical modules, the modules must be a UL (or other North American NRTL) RECOGNIZED COMPONENT, must be approved by TUV (or other European Product Safety test house) and the laser in the module must comply with Laser Class 1, US 21 CFR (J) and EN 60825-1.

ESD Precautions

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

Data Security

- Power down your host computer and all attached peripheral devices before beginning installation.
- Each enclosure contains up to 16 removable disk drive modules. Disk units are fragile. Handle them with care, and keep them away from strong magnetic fields.
- *All* the supplied plug-in modules and blanking plates must be in place for the air to flow correctly around the enclosure and also to complete the internal circuitry.

- If the enclosure is used with modules or blanking plates missing for more than a few minutes, the enclosure can overheat, causing power failure and data loss. Such use may also invalidate the warranty.
- If you remove any drive module, you may lose data. If you remove a drive module, replace it immediately. If it is faulty, replace it with a drive module of the same type and capacity.
- Ensure that all disk drives are removed from the enclosure before attempting to manhandle or move the rack installation.
- Do not abandon your backup routines. No system is completely foolproof.

Special Tools and Equipment

There are no special tools required but in order to complete the assembly of some configurations you may need the following (not supplied):

- Security keys (one of these should be included with your SA2016 enclosure for use with the drive locks).

Obtaining Publications

You can obtain SGI documentation in the following ways:

- See the SGI Technical Publications Library at <http://docs.sgi.com>. Various formats are available. This library contains the most recent and most comprehensive set of online books, release notes, man pages, and other information.
- If it is installed on your SGI system, you can use InfoSearch, an online tool that provides a more limited set of online books, release notes, and man pages. With an IRIX system, select **Help** from the Toolchest, and then select **InfoSearch**. Or you can type `infosearch` on a command line.
- You can also view release notes by typing either `grelnotes` or `relnotes` on a command line.
- You can also view man pages by typing `man <title>` on a command line.

Conventions

The following conventions are used throughout this publication:

Convention	Meaning
<code>command</code>	This fixed-space font denotes literal items such as commands, files, routines, path names, signals, messages, and programming language structures.
<i>variable</i>	Italic typeface denotes variable entries and words or concepts being defined.
[]	Brackets enclose optional portions of a command or directive line.
...	Ellipses indicate that a preceding element can be repeated.

Reader Comments

If you have comments about the technical accuracy, content, or organization of this document, contact SGI. Be sure to include the title and document number of the manual with your comments. (Online, the document number is located in the front matter of the manual. In printed manuals, the document number is located at the bottom of each page.)

You can contact SGI in any of the following ways:

- Send e-mail to the following address:
techpubs@sgi.com
- Use the Feedback option on the Technical Publications Library Web page:
<http://docs.sgi.com>
- Contact your customer service representative and ask that an incident be filed in the SGI incident tracking system.
- Send mail to the following address:

Technical Publications
SGI
1500 Crittenden Lane, M/S 535
Mountain View, California 94043-1351

SGI values your comments and will respond to them promptly.

Introducing the SGI InfiniteStorage RM 610/660

Introduction

Welcome to the SGI InfiniteStorage RM610/RM660.

The RM 610/RM660 Solution

RM controllers are designed and optimized for the high bandwidth and capacity requirements of departments, rich media and high performance workgroup applications.

RM controllers plug 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 RM610/RM660

RM controllers are 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 a highly parallel Fibre Channel architecture, each RM controller orchestrates a coherent flow of data throughout the SAN, from users to storage, managing data at up to 800 MB/second (200 MB/s per port on RM660). This task is accomplished through virtualized host and storage connections, a DMA-speed shared data access space, advanced network-optimized RAID engine data protection and security—all acting in harmony with sophisticated Fibre Channel storage management intelligence.

With its modular design, RM controllers can be “coupled” together forming data access redundancy while maintaining fully pipelined, parallel bandwidth to the same disk

storage. 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 RM610/RM660

The RM610/RM660 incorporates the following features:

- **Simplifies Deployment of Complex SANs**
RM controllers provide 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 800 MB/s Throughput (RM660)**
The RM660 provides up to 4 individual 200 MB/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.
- **Highly Parallel Fibre Channel Architecture (RM660)**
On the RM660, a highly parallel Fibre Channel 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 RAID engines
 - 10-independent RAID-protected high speed data caches
- **Fibre Channel Connectivity with 400MB/s Throughput (RM610)**
The RM610 provides up to 4 individual 100MBytes/sec Fibre Channel host 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.
- **Highly Paralle Fibre Channel Architecture (RM610)**
On the RM610, a highly parallel Fibre Channel architecture provides system performance and fault-tolerance.
 - 6-independent dual-ported FC drive channels; including separate (asynchronous, “self-optimized”) I/O queues for each channel
 - 6-independent RAID engines

- 6-independent RAID-protected high speed data caches
- **Highly Scalable Performance and Capacity**
The RAID engine 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.
- **Comprehensive, Centralized Management Capability**
RM controllers provide 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 of a LUN 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 (112 when using SFx016 enclosures).
- **Hot-Swappable and Redundant Components**
RM controllers utilize redundant, hot-swappable power supplies and cooling modules that can be replaced while the system is running.


Note: For information on the required operating system release levels and supported platforms for the RM controller, see the release notes that are included on the product CD.

The RM660 System Hardware


This section describes the hardware components of the SGI InfiniteStorage RM660.

The basic RM660 includes:

- A single unit (2.56 GB cache memory)
- Ten SFP (Small Form-factor Pluggable) cables which connect the RM660 to the FC drive enclosures
- Cable(s) for host Fibre Channel connection(s)
- Serial cable for maintenance/diagnostics
- Ethernet cable



▲ Front of RM660 Controller Behind Cover Panel



▲ Rear of RM660

Figure 1-1 The RM660 Front and Rear Views

The RM660 is a 2U, 19-inch rack-mountable chassis that houses four power supply modules, one fan module, and the internal electronics—a motherboard containing all of the drive-side, caching, and command processor functions, and a daughterboard containing the host connectivity electronics.

The system uses ten independent FC drive channels to manage data distribution and storage for up to 125 disk drives per channel (could be limited by drive enclosure type).

Power Supply and Fan Modules

Note: The power supply for the RM660 is interchangeable with the power supply for the RM610.

Each RM660 is equipped with four Power Supply/Cooling modules and one fan module (Figure 1-2). 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 LED mounted on the front of the Power Supply/Cooling module indicates the status of the PSU and the fans (Figure 1-3). It is green when the module is operating normally and turns red when a fault occurs.




Figure 1-2 Power Supply/Cooling Module

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

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 1-3 Fan Module

I/O Connectors and Status LED Indicators

Figure 1-4 shows the position of the I/O connectors at the back of the unit.




Figure 1-4 I/O Ports on Rear Panel of the RM660

The four “HOST” ports (SFP connectors) are used for Fibre Channel host connections. You may connect your host systems directly to these ports or connect the ports to your Fibre Channel switches and hubs. When an FC signal is present on the port, the respective LED turns green.

The “DISK” ports (SFP connectors) are for disk connections. The ten ports are labeled by data channels (ABCDEFGHPS) as illustrated in Figure 1-4. They are also color-coded to facilitate easy installation. When an FC signal is present on the port, the respective LED turns green.

The “CLI” connector provides local system monitoring and configuration capabilities.

The “UPS” connector (DB-9 male) can be connected to a UPS.

The “UPS” connector (DB-9 male) can be connected to a UPS. The “TELNET” connector provides remote monitoring and configuration capabilities. The “ACT” (Activity) LED flashes green when there is Telnet activity. The “LINK” LED is off when there is no Ethernet link. It turns amber when the link speed is 10 MB/s and turns green when the link speed is 100 MB/s.

The “S2A LINK” port is used to connect single RM660 units to form a couplet. The LED will illuminate when a link is present.

The “COM 2” port is a DB-9 serial female connector that is used to connect single RM660 units to form a couplet.

The “SYSTEM” LED is green when the power supplies and fans are operating normally. It turns amber if a power supply or fan fails.

Host and Disk Connectivity

The RM660 includes four 2 Gb/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.

The disk ports utilize ten standard copper SFP connectors.

The RM610 System Hardware

This section describes the hardware components of the SGI InfiniteStorage RM610.

The basic RM610 includes:

- A single unit (1.5GB cache memory)

- Six SFP (Small Form-factor Pluggable) cables which connect the RM610 to the FC drive enclosures
- Cable(s) for host Fibre Channel connection(s)
- Serial cable for maintenance/diagnostics
- Ethernet cable for remote management




Figure 1-5 The RM610 Front & Rear Views

The RM610 is a 1U, 19-inch rack-mountable chassis that houses two power supply/fan modules, one fan module, and the internal electronics—a motherboard containing all of the drive-side, caching, and command processor functions, and a daughterboard containing the host connectivity electronics.

The system uses six independent FC drive channels to manage data distribution and storage for up to 125 disk drives per channel (could be limited by drive enclosure type).

Power Supply and Fan Modules

Note: The power supply for the RM660 is interchangeable with the power supply for the RM610.

Each RM610 is equipped with two Power Supply/Cooling modules and one fan module (Figure 1-6). The PSU (power supply unit) voltage operating ranges are nominally 110V or 230V AC, selected automatically.

The two 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 LED mounted on the front of the Power Supply/Cooling module indicates the status of the PSU and the fans (Figure 1-6). It is green when the module is operating normally and turns red when a fault occurs.




Figure 1-6 Power Supply/Cooling Module

The two modules are installed in the left and right slots at the front of the unit, behind the cover panel (Figure 1-5). Each is held in place by two thumbscrews.

The fan module (Figure 1-7) is installed in the front center slot, behind the cover panel, and is held in place by two thumbscrews.




Figure 1-7 Fan Module

I/O Connectors and Status LED Indicators

Figure 1-8 shows the position of the I/O connectors at the back of the unit.




Figure 1-8 I/O Ports on Rear Panel of RM610

The four “HOST” ports (SFP connectors) are used for Fibre Channel host connections. You may connect your host systems directly to these ports or connect the ports to your Fibre Channel switches and hubs. When an FC signal is present on the port, the respective LED turns green.

The “DISKS” ports (SFP connectors) are for disk connections. The six ports are sequentially numbered and labeled by data channels (ABCDPS) as illustrated in Figure 1-8. They are also color-coded to facilitate easy installation. When an FC signal is present on the port, the respective LED turns green. If no FC signal is present on the port, the respective LED will flash red.

The “CLI” connector (RS-232 DB-9 female) provides local system monitoring and configuration capabilities.

The “UPS” connector (DB-9 male) can be connected to a UPS.

The “TELNET” connector provides remote monitoring and configuration capabilities. The “ACT” (Activity) LED flashes green when there is Telnet activity. The “LINK” LED is off when there is no Ethernet link. It turns amber when the link speed is 10MB/s and turns green when the link speed is 100MB/s.

The “SA LINK” port is used to connect individual RM610s to form a couplet. The LED will illuminate when a link is present.

The “COM 2” port is a DB-9 serial female connector that is used to connect individual RM610s to form a couplet.

The “SYSTEM” LED is green when the power supplies and fans are operating normally. It turns amber if a power supply or fan fails.

Host and Disk Connectivity

The RM610 includes four 1Gb/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.

The disk ports utilize six standard copper SFP connectors.

Installing the SGI InfiniteStorage RM610 and RM660

Installation Overview

This list provides an overview of the installation process. The following sections explain these steps in greater detail.

1. Unpack the system.
2. Install the drive enclosures in the 19" cabinet(s). (See Appendix A, "Technical Specifications for SGI InfiniteStorage RM660 and RM610" for physical dimensions and weight of the RM Controller)
3. Set up and connect the drive enclosures to the RM controller(s).
4. Connect the RM controller to your FC hub/switch, and host computer(s).
5. Connect your RS-232 terminal to the RM controller.
6. Power up the system.
7. Configure the storage array (create and format LUNs) via RS-232(CLI) interface or Telnet(GUI).
8. Define access rights for the clients in your SAN environment. Shared LUNs will need to be managed by SAN management software. Individual dedicated LUN will appear to the client as local storage and does not require any management software.
9. Initialize the LUNs for use with your server/client systems. Partition disk space and create filesystems as needed.

Setting Up the SGI InfiniteStorage RM610/RM660

Note: On IRIS FailSafe and Failover configurations, SGI recommends that you enable dual cache coherency, disable writeback cache, set a hard loop ID on each port, set each controller to have a unique WWN and set up host port zoning as required.

Warning: RM controllers do not have a cache mirroring feature and data cache is not copied from one RM controller to another in couplet configurations, regardless of the dual cache coherency setting. If writeback cache is enabled and an RM controller fails, all “dirty” data in cache will be lost. Thus if data integrity is paramount, writeback cache should always be disabled. Refer to the failover(7M) man page for additional information.

This section provides information on installing the hardware components of your SGI InfiniteStorage RM610/RM660 system.

Unpacking the Rack

If you purchase the 38U rack and drive enclosures together with the RM controller, all the enclosures will have been pre-installed in the rack at the factory.

Tools that you will need are 11mm wrench, 18mm wrench, screwdriver, and scissors.

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).

1. Remove the carton and ALL protective material.
2. The rack is anchored to the pallet with two brackets. Remove the bolts and screws from both brackets (Figure 2-1).




Figure 2-1 Remove Anchoring Brackets

3. Place the two brackets to the side of the pallet marked “DOOR/RAMP” so they make two ramps for the rack to be rolled off from the pallet (Figure 2-2). Make sure the flange on the end of the bracket is inserted in the slot on the side of the pallet.




Figure 2-2 Creating the Ramps for the Pallet

4. Remove the stop block from the “DOOR/RAMP” side only.
5. Given the weight and size of the rack, at least **four** people are required to unload the rack. Carefully roll the rack off of the pallet and slowly move it to the desired location.
6. After the rack is finally positioned, lower the feet to firmly position the rack. Then attach the three trims, supplied with the rack, to the bottom of the rack (Figure 2-3). Attach the folded end trim to the front of the rack. Attach the two straight trims to the two sides of the rack. Use the 12 screws provided to fasten the trims.




Figure 2-3 Attach Trims to Bottom of Rack

7. To open the door: slide up the clear cover, press the button to unlock which will release the handle, turn the handle to unlatch the door (Figure 2-4).




Figure 2-4 Opening the Rack Door

Note: Please follow the safety guidelines for rack installation given in Appendix B.

Unpacking the RM610/RM660

Before you unpack your RM controller, 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 chassis and all components for signs of external damage. If you detect any problems, please call SGI Customer Service.

Your shipping carton contains the following items:

- RM610/RM660 chassis with:
 - power supply/cooling modules
 - one fan module
- two power cords
- RS-232 and Ethernet cables for monitoring and configuration
- Cover panel and rack-mounting hardware
- Rail kit with mounting hardware

Warning: 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 RM610/RM660 Chassis

The RM610/RM660 fits into a 19" rack. Figure 2-5 shows the set of mounting hardware shipped with the system for use in square-hole rack, threaded-hole rack, and through-hole rack.

The cage nuts and speed nuts are supplied as alternative mounting options for different types of rack, and are not used for threaded-hole racks. The cage nuts are used for square-hole racks and the speed nuts are used for "non-threaded" through-hole racks.




Figure 2-5 RM610/RM660 8x00 Mounting Hardware

1. Attach the two short mounting brackets to the front of unit, if not already installed. Using five #6 screws, fasten the bracket to one side of the chassis (Figure 2-6 and Figure 2-7). Be sure to tighten the screws. Similarly, attach the other short bracket to the other side of the chassis.




Figure 2-6 Attaching the Mounting Brackets to Front of the RM660




Figure 2-7 Attaching the Mounting Brackets to the Front of the RM610

2. For square-hole rack, install cage nuts into front of rack, 2 on each side (Figure 2-8). For through-hole rack, install the speed nuts.




Figure 2-8 Square Hole Rack with Cage Nuts Installed

3. Using four #10 screws, attach the two rear mounting brackets to the back of the rack (Figure 2-9). On square-hole and through-hole racks, use the brackets with pems. On threaded-hole rack, use the brackets without pems.

For square-hole racks, attach a mounting plate to the rack frame, as shown in Figure 2-9 (dimples should fit into the square holes).




Figure 2-9 Rack with Rear Mounting Bracket Installed

4. While a partner supports the RM controller, slide the unit into rack catching the unit shoulder screws in rear mounting bracket slots.
5. Secure the front of the RM controller to the rack by inserting four #10 screws (two on each side) through the holes in the mounting brackets and into the rack frame (Figure 2-10 and Figure 2-11).

The front brackets have “tabs” that fit into the square holes. For racks with threaded or through holes, insert a mounting spacer between the mounting bracket and rack frame.

Be sure to tighten the screws.




Figure 2-10 Securing Front of the RM660 in the Rack




Figure 2-11 Securing Front of the RM610 in the Rack

6. Attach the cover panel to the front of chassis using two thumb screws.

Connecting the RM Controllers in Dual Mode

For dual mode configuration only:

1. Connect the "S2A LINK" ports on the two RM controller units using an Ethernet Crossover cable.
2. Connect the "COM2" ports on the two units using a DB-9 male-to-male cable.

Connecting the RM Controllers

The following sections describe how to connect the RM660 and the RM610.

Connecting the RM660

If you are using the SF family of enclosures with the RM660, please refer to "Setting Up the Drive Enclosures" in Appendix C for information on how to set up the enclosures and connect them to the RM660.

1. There are ten disk channels on the RM660. The disk ports are labeled as follows (Figure 2-12):

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)	

Using the ten copper SFP cables provided, connect these disk ports to your ten disk channels. Make sure the latches on the transceivers are engaged.




Figure 2-12 I/O Connectors on Back of the RM660

- Each RM660 supports up to 4 Fibre Channel host connections. You may connect more than four client systems to the RM660 with the use of hubs/switches and you can restrict user access to the LUNs (as described in “Configuring the RM610/RM660”).

The Host ports (SFP) are numbered 1 through 4 as shown in Figure 2-12. Connect your host system(s), hubs or switches to these ports. Make sure the latches on the transceivers are engaged.

Connecting the RM610

If you are using the SF family of enclosures with the RM610, please refer to “Setting Up the Drive Enclosures” in Appendix C for information on how to set up the enclosures and connect them to the RM610

- There are six disk channels on the RM610. The disk ports are labeled as follows (Figure 2-13)
 - 6-A = Channel A
 - 5-B = Channel B
 - 2-C = Channel C
 - 1-D = Channel D
 - 3-P = Channel P (parity)
 - 4-S = Channel S (spare)

Using the six copper SFP cables provided, connect these disk ports to your six disk channels. Make sure the latches on the transceivers are engaged.




Figure 2-13 I/O Connectors on Back of the RM610

- Each RM610 supports up to 4 Fibre Channel host connections. You may connect more than four client systems to the RM610 with the use of hubs/switches and you can restrict user access to the LUNs (as described in “Configuring the RM610/RM660” on page 28).

The Host ports (SFP) are numbered 1 through 4 as shown in Figure 2-13. Connect your host system(s), hubs or switches to these ports. Make sure the latches on the transceivers are engaged.

Selecting AL_PA for Your Drives

Each drive channel on the RM controller is a single Fibre Channel Arbitrated Loop. Hence, all the drives on the same channel must have their unique Arbitrated Loop Physical Addresses—AL_PA.

On JBOD enclosures, unique AL_PA is assigned to each drive usually based on the selected enclosure ID. Make sure the address values are correctly set on all your drive enclosures.

Note: Please refer to Appendix C for further information on setting the SF enclosure ID.

Laying Out your Storage Drives

The following sections describe how to lay out your storage drives on the RM660 and the RM610.

Laying Out your Storage Drives on the RM660

The RM660 is capable of managing up to 125 tiers. Tiers are the basic building blocks of the RM660. One tier contains 10 drives—eight data drives (Channels A through H), one parity drive (Channel P), and one *optional* spare drive (Channel S). Drives that have the same AL_PA across all ten channels are put on the same tier. The tier that contains the highest AL_PA drives is recognized as Tier #1.

Configuration of disks in the enclosures must be in sets of complete tiers (Channels A through P). Allocating one spare drive per tier will give you the best data protection but this is not required. The spare drives on the RM660 are global hot spares.

Figure 2-14 illustrates the channel and tier numbering of the drive slots with respect to their AL_PA values for a system using five SF4016 drive enclosures. All enclosure IDs have been set to 1 and there are eight tiers in the system. See Appendix C for other enclosure AL_PA/tier mapping charts.




Figure 2-14 Channel and Tier Numbering of Drives on SF4016 Enclosures

Laying Out your Storage Drives on the RM610

The RM610 is capable of managing up to 125 tiers. Tiers are the basic building blocks of the RM610. One tier contains 6 drives—four data drives (Channel A through D), one parity drive (Channel P), and one *optional* spare drive (Channel S). Drives that have the same AL_PA across all six channels are put on the same tier. The tier that contains the highest AL_PA drives is recognized as Tier #1.

Configuration of disks in the enclosures must be in sets of complete tiers (Channels A through P). Allocating one spare drive per tier will give you the best data protection but this is not required. The spare drives on the RM610 are global hot spares.

Figure 2-15 illustrates the channel and tier numbering of the drive slots with respect to their AL_PA values for a system using three SF4016 drive enclosures. All enclosure IDs have been set to 1 and there are eight tiers in the system. See “Enclosure AL_PA and Tier Mapping Charts” on page 295 for other enclosure AL_PA/tier mapping charts.




Figure 2-15 Channel and Tier Numbering of Drives on SF4016 Enclosures

Connecting the RS-232 Terminal

For first time set up, you will need to use a RS-232 terminal or terminal emulator (such as Windows Hyperterminal). Then you may set up the remote management functions and configure/monitor the RM controllers remotely via Telnet.

1. Connect your terminal to the CLI port at the back of the RM controller using a standard DB-9 female-to-male Null Modem cable (Figure 2-16).




Figure 2-16 CLI Port on the RM660




Figure 2-17 CLI Port on the RM610

2. Then bring up your terminal window and use the following settings for your serial port:
 - Bits per second: 115,200
 - Data bits: 8
 - Parity: None
 - Stop bits: 1
 - Flow control: None

Powering On the RM Controller

1. Verify that the power switches on the power supplies at the front are off. Using the power cords provided, connect the two AC connectors at the back to your AC power source. For maximum redundancy, connect the two power connectors to two different AC power circuits.
2. Check that all your drive enclosures are powered up and the drives are spun up and ready.
3. Turn on the power supplies on the RM controller.

The RM controller will go through a series of system diagnostics and the bootup sequence is displayed on your terminal. Wait until the sequence is complete and the RM660 or RM610 prompt is displayed.

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

4. You may now configure the system as described in “Configuring the RM610/RM660”.

Configuring the RM610/RM660

This section provides information on configuring your SGI InfiniteStorage RM610 and RM660.

Note that the configuration example provided in this section represents a general guideline and should not be used directly to configure your particular RM controller. The CLI commands used in these examples are fully documented in “Managing the SGI InfiniteStorage RM610/RM660” on page 45 to “Other Utilities” on page 110, though exact commands may change depending on your firmware version.

Setup Planning

The RM610/RM660 offers great flexibility in defining LUN configurations and optimizing performance for specific applications.

Before proceeding with your RM controller configuration, it is necessary to 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 RM660 provides a fully parallelized, distributed, and multi-tasking hardware RAID implementation that uses a 8+1+1 parity scheme in the RAID engine. The RM610 uses a 4+1+1 parity scheme.

The RAID engine in the RM controller combines the virtues of RAID 3, RAID 5 and RAID 0 (Figure 2-18). Like RAID 3, a dedicated parity drive is used per parity group (8+1 on the RM660, 4+1 on the RM610) called a “Tier”. The RAID engine exhibits RAID 3 characteristics such as tremendous large block-transfer—READ and WRITE—capability with NO performance degradation in crippled mode.

Tier Configuration				
Tier	Capacity (Mbytes)	Space Available (Mbytes)	Disk Status	Lun List
1	280012	271820	ABCDEFGHPS	0
2	280012	271820	ABCDEFGHPS	0
3	280012	271820	ABCDEFGHPS	0

Striping across tiers when a LUN is created across multiple tiers ↓

-----> Parity Protection within same tier

Figure 2-18 Striping Across Tiers

Like RAID 5, however, the RAID engine 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 a greatly enhanced ability to handle small I/O (particularly as disk spindles are added) and many, many streams of real-time content.

LUNs can be created on a fraction of a Tier, a full Tier, across a fraction of multiple Tiers, or across multiple full Tiers.

Tiers of drives are added 9 drives (8+1) at a time on the RM660 and 5 drives (4+1) at a time on the RM610, with an arbitrary number of hot spares (up to 125) supported.

On the RM660, a tenth loop is utilized for the spare disks (8+1+1). On the RM610, a sixth loop is utilized for the spare disks (4+1+1). Up to 125 Tiers may be utilized, but the current SF4016 enclosure limits the number of Tiers to 120 (due to the Fibre-Channel ID

numbering scheme in the enclosure itself). When using SF2016 enclosures, this number is limited to 112.

Up to 128 LUNs can be created in total. Each LUN can be subdivided into up to 64 smaller equally-sized LUN segments, giving a total of 8192 LUN segments. The LUN segments of a LUN are managed together and share the same characteristics. LUNs (and LUN segments) can be shared or dedicated to individual user, according to your security level setup with Read or Read/Write privileges granted per user. The users will only have access to “their own” and “allowed to share” LUNs. Shared LUNs will need to be managed by SAN management software. Individual dedicated LUNs will appear to the users as local storage and do not require external management software.

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

In general, 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, please contact SGI Technical Support.

Configuration Options

There are two options to configure the RM controller:

- Command Line Interface (CLI)
- RM Controller Management Tool which is a JAVA-based GUI RAID client-server application manager that provides a user-friendly graphical user interface (GUI)

This user guide provides information for set up using the CLI. If you want to configure the RM controller using the Management Tool, please refer to the *RSM Management Tool User Guide*. Configuration examples are given for both options.

With the RM controller ready, hit <Enter> in the terminal window to get the RM660 or RM610 prompt.

Login as Administrator

Only users with administrator rights are allowed to change the RM controller configuration.

To login, type:
login admin<Enter>

Then enter password:
password<Enter>

“admin” is the default Administrator account name and its password is “password”. (see “Configuration Management” on page 51 for information on how to change the user and administrator passwords)

Setting System Time & Date

The RM controller’s system time and date are factory-configured for U.S. Pacific Standard Time. This needs to be changed if you are located in other time zones so that the time stamps for all events will be correct. In dual mode, changes should always be made on Unit 1. New settings will automatically be applied to both units.

To change the system date to March 1, 2004, for example, type:
date 3 1 2004<Enter>

To change the system time to 2:15:32pm, for example, type:
time 14:15:32<Enter>

Setting Tier Mapping Mode

The RM610/RM660 supports various FC-AL drive enclosures. When the system is first configured, it is necessary to select a tier mapping mode for the attached enclosures. This will tell the RM controller how the tiers are laid out with respect to the drives’ AL_PA. Hence, the system can correctly light the drive’s LED and provide visual indication of a specific drive, tier, and LUN.

To display the current mapping mode, type: `tier map<Enter>`

To change the mapping mode:

- Type: `tier changemap<Enter>`

```

RM660 [1]: tier changemap

Supported mapping modes:
-----
0) Standard Enclosure mapping.
1) SF2012 Enclosure mapping.
2) SF4012 Enclosure mapping.
3) SF2016 Enclosure mapping.
4) SF4016 Enclosure mapping.

Current mapping mode: 4, SF4016 Enclosure.
Enter new mapping mode (0..4), 'e' to escape:

```

Figure 2-21 Tier Mapping Mode Setup Screen

- Then select the appropriate mapping mode and press <Enter> (Figure 2-21). Use Mapping Mode 0 for SF6016 enclosures.
- Type: `restart<Enter>`, for the changes to take effect.

If you have other types of drive enclosures, select Tier Mapping Mode 0. The RM controller might be able to monitor the status of your enclosures, however, SGI is unable to guarantee compatibility with unqualified third party devices.

Verifying Connections for SF6016 Enclosures

If you are using the SF6016 drive enclosures, it is necessary to verify the connections of the enclosures that are currently attached to the system, so that SES can accurately indicate a drive fault through SES commands. Please refer to “Verifying Connections for SF6016 Enclosures” on page 284 and perform all the steps described before proceeding to the next section.

Checking Tier Configuration

Use the “tier” command to display your current tier configuration. Figure 2-22 illustrates the configuration of a system containing 80 drives on 8 tiers.

```
RM660 [1]: tier
```

Tier Configuration					
Tier	Owner	Capacity (Mbytes)	Space Available (Mbytes)	Disk Status	Lun List
1	1	280012	280012	ABCDEFGHPS	
2	1	280012	280012	ABCDEFGHPS	
3	1	280012	280012	ABCDEFGHPS	
4	1	280012	280012	ABCDEFGHPS	
5	1	280012	280012	ABCDEFGHPS	
6	1	280012	280012	ABCDEFGHPS	
7	1	280012	280012	ABCDEFGHPS	
8	1	280012	280012	ABCDEFGHPS	

Automatic disk rebuilding is Enabled
 System rebuild extend: 32 Mbytes
 System rebuild delay: 60
 System Capacity 2240096 Mbytes, 2240096 Mbytes available.

Figure 2-22 Current Tier Configuration Screen

Each letter under the “Disk Status” column represents a healthy drive at that channel. A “space” indicates that the drive is not present (or detected) at that location. A period (.) denotes that the disk was failed by the system. “?” 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.

1. Verify that all your drives can be seen by the RM controller.
 If a drive is missing, make sure that the drive is properly seated and in good condition. Then enter the “disk scan” command to look for the drive. If the same channel is missing on all tiers, check the cable connections for that channel.
2. Check that “Automatic disk rebuilding” is enabled. If not, enter command “tier autorebuild=on” to enable it.

Cache Coherency and Labeling in Dual Mode

1. Use the “DUAL” command to check status of the units are healthy and verify the “Dual” (COM2) and “Ethernet” (S2A LINK) communication paths between the two units are established (Figure 2-23).

```

RM660 [1]: dual
                Dual RM660 Configuration
                RM660 1                      RM660 2
-----
Label           RM660 Unit[1]                RM660 Unit[2]
Status          Healthy                      Healthy

Dual communication:    established.
Ethernet communication: established.
Cache coherency:      Not established.
Cache coherency timeout: 2

```

Figure 2-23 Dual RM Controller Configuration Screen

2. 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, type:

```
DUAL COHERENCY=ON|OFF<Enter>
```

3. You may change the label assigned to each RM controller unit. This allows you to uniquely identify each unit in the RM controller system. Each RM controller can have a label up to 31 characters long.

To change the label, enter command:

```
DUAL LABEL<Enter>
```

Then select which unit you want to rename (Figure 2-24). When prompted, type in the new label for the selected unit. The new name will be displayed.

```
RM660 [1]: dual label
Enter the number of the HSTD you wish to rename.
LABEL=1 for HSTD 1, Test System[1]
LABEL=2 for HSTD 2, Test System[2]
HSTD: 1
Enter a new label for HSTD 1, or DEFAULT to return to the default label.
Up to 31 characters are permitted.
Current HSTDname: Test System[1]
New HSTDname: Ssystem[1]
```

Figure 2-24 Labeling an RM Controller Unit

Configuring the Storage Arrays

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

Now that you have determined your array configuration, you will need to create and format the LUNs. Given below is an example for creating 2 LUNs:

- LUN 0 with 4 LUN segments on Tiers 1 to 8 with capacity of 8192MB each
 - LUN 1 on Tiers 1 and 2 with capacity of 8192MB
1. Enter command “cache” to display the current cache settings. Verify that the amount of cache installed is correct.
 2. Then select a cache segment size for your array. To set the segment size to 128KBytes, for example, type:
cache size=128<Enter>

This setting can also be adjusted on-the-fly for specific application tuning (see “Cache Segment Size” on page 83 for more information). The default setting is 128.
 3. Enter command “lun”. The Logical Unit Status chart should be empty as no LUN is present on the array.

4. To create a new LUN, type:
lun add=0<Enter>
where 0 (zero) is the LUN number. Valid LUN numbers are 0..127.
If only "lun add" is entered, you will be prompted to enter a LUN number.
5. You are then prompted to enter the parameter values for the LUN. You may press "e" at any time to exit the format LUN menu and cancel the command completely. In this example, the parameters to enter will be:
 - Enter a label for the LUN which may contain up to 12 characters (the label may be changed later on using the LUN LABEL command).
 - You are asked if you want to create a LUN group. For LUN 0:
y<Enter>
 - Enter the number of LUNs you want in the LUN group, type:
4<Enter>
 - Enter the capacity (in Mbytes) for a single LUN in the LUN group, type:
8192<Enter>
 - Enter the number of tiers to use.
8<Enter>
 - Then select the tier(s) by entering the Tier number. Enter each one on a new line and press <Enter>. The 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>
This is the recommended block size. A larger block size may give better performance. However, please verify that your OS and file system can support a larger block size before changing the block size from its default value.
6. The message "Operation successful: LUN 0 added to the system" then appears. When you are asked to format the LUN, enter "y" to continue.

7. After you have initiated LUN format, the message “Starting Format of LUN” is displayed. You may monitor the format progress by entering the command “lun” (Figure 2-25).

```

RM660 [1]: lun
                Logical Unit Status
LUN  Label  Owner   Status   Capacity  Block  Tiers  Tier List
      (Mbytes) Size
-----
0    1      1      Format 14%  4x  8192   512    8    1 2 3 4 5 6 7 8
                System Capacity 2240096 Mbytes, 2207328 Mbytes available.
    
```

Figure 2-25 Logical Unit Status - Formatting

8. Upon completion, the message “Finished Format of LUN 0” will be displayed. Enter command “lun” to check the status of the LUN, which should be “Ready” (Figure 2-26).

```

RM660 [1]: lun
                Logical Unit Status
LUN  Label  Owner   Status   Capacity  Block  Tiers  Tier List
      (Mbytes) Size
-----
0    1      1      Ready   4x  8192   512    8    1 2 3 4 5 6 7 8
                System Capacity 2240096 Mbytes, 2207328 Mbytes available.
    
```

Figure 2-26 Logical Unit Status - Ready

9. To create the LUN 1, type:
`lun add=1<Enter>`
10. Then enter the parameters:
 - Enter a label for the LUN 1
 - When asked if you want to create a LUN group, type:
`n<Enter>`
 - For capacity, enter the value in MBytes.
`8192 <Enter>`
 - Enter the number of tiers to use.
`2<Enter>`

- Then select the tier(s) by entering the Tier number. Enter each one on a new line and press <Enter>. 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>

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

Note: Immediate LUN availability is not an option on the RM610 or RM660 products. All LUNs must complete format prior to usage.

Setting Security Levels

After you have formatted all the LUNs you need, you may define the users' right of access to the storage. There are two types of configurations: authorized user and host port zoning.

The Authorized User configuration is highly recommended for use in a SAN environment as your data will be completely secured and no accidental plug-in is allowed to do damage such as data deletion, LUN removal, and LUN format. The authorized users will only have access to "their own" and "allowed to share" data. Administrator 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 will see the same LUN identification scheme regardless of which host port it is connected to.

The Host Port Zoning configuration will give the minimum level of security. The LUN mappings will 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 RM 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, and the HBA. All IRIX servers require a LUN 0 on host port(s) for the server to recognize LUNs.

User Authentication (Recommended for SAN Environment)

Each user connected to the RM controller is identified by its World Wide Name (WWN) and is given a unique user ID number. The RM610/RM660 can store configuration for up to 512 users and the security settings apply to all host ports.

Given below is an example for adding 2 users to a system containing 2 LUNs (numbered 0, 1). Each user will have a dedicated LUN segment and internal LUN 1 will be shared and “read-only”. Both users will see the shared LUN as LUN 0 and they will see their own LUN segment as LUN 1. User 1 will have access to host ports 1 and 4 while User 2 will only have access to host port 2.

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

- Type: zoning<Enter>

```
RM660 [1]: zoning
```

		LUN Zoning	
Port	World Wide Name	External LUN, Internal LUN	
1	2100CB5CBA7F5F1F		
2	2200CB5CBA7F5F1F		
3	2300CB5CBA7F5F1F		
4	2400CB5CBA7F5F1F		

Figure 2-27 LUN Zoning Screen

- Check that the “LUN Zoning” chart is empty (Figure 2-27).
If not, follow the procedure in “Host Port Zoning (Anonymous Access)” on page 43 to remove all mapping to internal LUNs.

To add a user:

1. Type: user audit=on<Enter>
The RM controller will report which users are connected.
2. Type: user add<Enter>. Then enter the ID number of the user to be added (which is 0 in this example, see Figure 2-28).

3. You may assign an alias name for the user. The name may contain up to 12 characters. Type in a name and press <Enter>.
4. Enter “y” to specify host port zoning. Then enter each active port on a new line.
Type:
 1<Enter>
 4<Enter>
 e<Enter>
5. Enter “y” to specify the unique LUN mapping. For this user:
External LUN 0 is mapping to internal LUN: R1
External LUN 1 is mapping to internal LUN: 0.00
External LUN 2 is mapping to internal LUN: e

Internal LUN 1 will be read-only.
6. The LUN zoning chart is displayed for this newly added user.

```

RM660 [1]: user add

  ID  User          Port  S_ID  World Wide Name  Login
-----
  0   Anonymous     1     E8   210000E08B057383  MON JUN 03 13:04:32 2002
 128 Anonymous     2     E8   210000E08B028233  MON JUN 03 13:51:27 2002
 256 Anonymous     3     E8   210000E08B01753A  MON JUN 03 13:51:27 2002

Enter a user ID or 's' to specify a world wide name or 'e' to escape.
0<enter>

Enter an alias name for the user up to 12 characters.
client1<enter>

The host ports can be zoned for this user to indicate which host
ports the user is allowed to log into. The user has access to all
the host ports by default.

Do you want to zone the host ports for this user? (y/n): y<enter>

Enter active ports (1..4) for RM660 [1]
'e' to escape:
1<enter>
4<enter>
e<enter>

A user can have a unique LUN mapping for all of the host ports or
it can use the anonymous LUN mapping from the port zoning command.
Do you want this user to have a unique LUN mapping? (y/n): y<enter>

Enter the new LUN mapping for user client1 210000E08B057383.

For each external LUN specify the internal LUN the user will have access
to. 'R' preceding the LUN indicates the LUN will be read-only for the user.
Entering 'P' indicates a place holder LUN that will have zero capacity.
Enter 'n' for no LUN or 'e' to escape.

Specify the internal LUN (0..127), or 'n' for non or 'e' to escape.
External LUN 0 is mapping to internal LUN: R1<enter>
External LUN 1 is mapping to internal LUN: 0.00<enter>
External LUN 2 is mapping to internal LUN: e<enter>

                                LUN Zoning
User                               External LUN, Internal LUN
-----
000 client1                        R000,001 001,0.00

User client1 210000E08B057383 added.

```

Figure 2-28 Configuring the User Access to LUNs

- Connect user #2 and repeat Steps [2] to [6] above to specify the host port zoning and LUN mappings.

For active host port, enter port 2 only.

For LUN mapping:

External LUN 0 is mapping to internal LUN: R1

External LUN 1 is mapping to internal LUN: 0.01

External LUN 2 is mapping to internal LUN: e

In this scheme, both users have their own customized LUN identification scheme. The internal LUN 1 that is shared by the users will need to be managed by SAN management software. The individual dedicated LUN (or LUN segment) will appear to the user as local storage and does not require external management software.

To display the new security settings, enter command “user” (Figure 2-29).

```
RM660 [1]: user
```

User	World Wide Name	Ports		LUN Zoning External LUN, Internal LUN
		1	2	
000 client1	210000E08B057383	1	4	R000,001 001,0.00
128 client2	210000E08B028233	2		R000,001 001,0.01

User auditing is enabled.

Figure 2-29 Current Security Settings Screen

Host Port Zoning (Anonymous Access)

This type of set up should only be used for non-SAN environment. Users are given the “general admission” to the data. It only provides the minimum level of security.

One zoning configuration is supported for each of the host ports. Any unauthorized user accessing the RM controller storage will be considered “anonymous” and granted the Zoning access for the host port they are connected to.

Given below is an example for adding LUN zoning to host port 1. External LUN 1 will be mapped to internal LUN 0 and it will be read-only for the users.

- Enter command “zoning edit”. The current settings are displayed (Figure 2-30).
- Select a host port (1..4). Enter “1”.

3. Then specify the internal LUN (0..127) to be mapped to the external LUN. In this example, the settings should be:
 External LUN 0 is not mapped. New internal LUN: n
 External LUN 1 is not mapped. New internal LUN: R0
 External LUN 2 is not mapped. New internal LUN: e

```

RM660 [1]: zoning edit

                                LUN Zoning
      Port   World Wide Name      External LUN, Internal LUN
-----
      1     21000001FF000218
      2     22000001FF000218
      3     23000001FF000218
      4     24000001FF000218

Enter the host port (1..4), 'e' to escape.
1<enter>

Enter the new LUN zoning for this host port.

For each external LUN specify the internal LUN the user will have access
to. 'R' preceding the LUN indicates the LUN will be read-only for the user.
Enter 'P' indicates a place holder LUN that will have zero capacity.
Enter 'n' for no LUN or 'e' to escape.

Specify the internal LUN (0..127), or 'n' for none or 'e' to escape.

External LUN 0 is not mapped. New internal LUN: n<enter>
External LUN 1 is not mapped. New internal LUN: R0<enter>
External LUN 2 is not mapped. New internal LUN: e<enter>

                                LUN Zoning
      Port   World Wide Name      External LUN, Internal LUN
-----
      1     21000001FF000218  R001,000
      2     22000001FF000218
      3     23000001FF000218
      4     24000001FF000218

Zoning changed on host port 1.
    
```

Figure 2-30 Zoning Configuration Screen

1. The new settings are displayed.
2. Repeat the above steps to configure other host ports.

The RM610/RM660 installation is complete at this point. Your client systems should now be able to access the LUNs on the RM controller. If you want to remotely monitor the RM controller, please refer to “Remote Management of the RM Controller” on page 114 for information on how to configure the RM controller’s network interface settings.

Using the SGI InfiniteStorage RM610/RM660 Management and Administrative Facilities

Managing the SGI InfiniteStorage RM610/RM660

The RM controller provides a set of management tools that enable administrators to centrally manage storage and network resources, which handle business-critical data. These include Configuration Management, Performance Management, LUN Management, Security Administration, and Firmware Update Management.

Management Interface

All the SAN management information can be accessed both locally and remotely on the RM controller.

Note: Each RM controller may only have one active login (serial or Telnet) at any given time.

Locally via Serial Interface

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

Connect your terminal to the CLI port at the back of the RM controller using a standard DB-9 female-to-male Null Modem cable (Figure 3-1 and Figure 3-2).




Figure 3-1 CLI Port at Back of the RM660




Figure 3-2 CLI Port at back of the RM610

Then bring up your terminal window and use these settings for your serial port:

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

With the RM controller ready, hit <Enter> in your terminal window to get the “RM660” or “RM610” prompt.

It is possible to change the baud rate on RM controller. Refer to “Changing Baud Rate for CLI Interface” on page 113 for information.

Remotely via Telnet

To configure and monitor the RM controller remotely, connect the RM controller to your Ethernet network. Please refer to “Remote Management of the RM Controller” on page 114 for further information on how to set up the RM controller’s network interface.

Available Commands

The “Help” command will display the available commands within the RM610/RM660 Utility (Table 3-1). To get help information on a command, type the command followed by a question mark. For example, `cache ?`.

Table 3-1 List of Available Commands

Commands	Descriptions	Reference Sections
APC_UPS	Displays the status of the APC UPS trap monitor	“APC UPS SNMP Trap Monitor” on page 110
API	Displays/Changes status of API connections	“API Server Connections” on page 110
AV	Displays information about audio/visual settings of the system	“Audio/Visual Settings of the System” on page 85
CACHE	Displays and adjusts the cache settings of the LUNs	“Optimization of I/O Request Patterns” on page 82
COMMENT	Echoes a message to the screen and saves it to the log	“Saving a Comment to the Log” on page 110
CONSOLE	Displays or modifies the console settings of this RM controller	“Changing Baud Rate for CLI Interface” on page 113
DATE	Displays / Changes the current system date	“Setting the System’s Date and Time” on page 72
DEFAULTS	Restores the system to its default configuration	“Restoring the System’s Default Configuration” on page 73
DISK	Displays information about the disks in the system	“Configure and Monitor Status of Storage Assets” on page 55, “Automatic Drive Rebuild” on page 77, “Optimization of I/O Request Patterns” on page 82
DUAL	Displays information about dual system configuration	“Couplet RM Controller Configuration (Cache/Non-Cache Coherent)” on page 79
FAULTS	Displays all current disk, system, and disk enclosure faults	“System and Drive Enclosure Faults” on page 109

Table 3-1 List of Available Commands **(continued)**

Commands	Descriptions	Reference Sections
HELP	Displays help information about system commands	"Available Commands" on page 47
HOST	Displays information about the host fibre channel ports	"Configure and Monitor Status of Host Ports" on page 51, "Optimization of I/O Request Patterns" on page 82, "Resources Allocation" on page 96
LOG	Displays a log of previous system messages	"Message Log" on page 108
LOGIN	Allows the user to log into a (new) CLI or Telnet session (i.e. at a specific security level)	"Administrator and User Logins" on page 49, "Login Names and Passwords" on page 118
LOGOUT	Allows the user to log out of a CLI or Telnet session	"Administrator and User Logins" on page 49
LUN	Displays status of and adds/removes LUNs to/from the system	"Configure and Monitor Status of Storage Assets" on page 55, "LUN Management" on page 73, "Locking LUN in Cache" on page 88
MIRROR	Displays/changes Internal Mirrored Groups (IMG) in the system	"Internal Mirrored Groups (IMG)" on page 111
NETWORK	Displays/changes information about the systems network configuration	"System Network Configuration" on page 66, "Network Interface Set Up" on page 115
PASSWORD	Changes the CLI and Telnet usernames and passwords	"Administrator and User Logins" on page 49, "Login Names and Passwords" on page 118
RESTART	Performs a restart of the system	"Restarting the SGI InfiniteStorage RM610/RM660" on page 71
ROUTE	Displays/updates information about the unit's IP routing table	"System Network Configuration" on page 66, "Network Interface Set Up" on page 115
SAVE	Saves the parameter blocks	"Saving the RM Controller's Configuration" on page 72
SES	Displays all current disk enclosure faults	"Configure and Monitor Status of Storage Assets" on page 55, "Resources Allocation" on page 96
SETTINGS	Displays/changes the CLI/Telnet session control settings	"CLI/Telnet Session Control Settings" on page 113
SHUTDOWN	Performs a shutdown of the system	"Restarting the SGI InfiniteStorage RM610/RM660" on page 71
STATS	Displays the system performance statistics	"System Performance Statistics" on page 90

Table 3-1 List of Available Commands **(continued)**

Commands	Descriptions	Reference Sections
TELNET	Allows the user to display/change whether remote Telnet sessions are currently enabled or disabled for this RM controller	“Remote Login Management” on page 105
TFTP	Performs an update of the firmware on the system using TFTP	“Firmware Update Procedure” on page 104
TIER	Displays information about the tiers in the system	“Configure and Monitor Status of Storage Assets” on page 55, “Tier Mapping for Enclosures” on page 64, “Automatic Drive Rebuild” on page 77, “Resources Allocation” on page 96
TIME	Displays /changes the current system time	“Setting the System’s Date and Time” on page 72
UPTIME	Displays the system's uptime	“Displaying System’s Uptime” on page 109
USER	Displays information and configures the system users	“Monitoring User Logins” on page 100, “User Authentication” on page 102
VERSION	Displays the firmware version of the system	“Displaying Current Firmware Version” on page 103
WHOAMI	Displays the owner of the current CLI or Telnet session	“Administrator and User Logins” on page 49
ZONING	Displays and changes the default LUN zoning for each host port	“Zoning (Anonymous Access)” on page 101

Administrator and User Logins

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

For RS-232 terminal session, the general purpose user does not require login. For 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 the status and configuration information of the system.

If the RM controller determines that the individual does not have the proper privileges, it will return a message:

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

Login

To login, enter command “login”. Then enter the login name and password (Figure 3-3). 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”.

```
RM660 [1]: login
Enter a login name: admin
Enter the password:

Successful CLI session login.
New owner      : admin.
New security level: Administrative.
```

Figure 3-3 Login Screen

Logout

To logout, enter command “logout”. For a terminal session, you will be returned to the general purpose user level. For Telnet, the current session will be disconnected.

Password

The “PASSWORD” command allows the administrator to change the login names and passwords for administrative and general purpose users (Figure 3-4). The associated privileges, however, remain the same regardless of the name or password changes.

```
RM660 [1]: password
Enter new name to replace <admin>:
Enter old password:
Enter new name to replace <user>:
Enter old password:
```

Figure 3-4 Password Configuration Screen

To reset login names and passwords to defaults, enter command:
PASSWORD DEFAULTS

Who Am I

The “WHOAMI” command will display the owner of the current terminal or Telnet session (Figure 3-5).

```
RM660 [1]: whoami
CLI session:
Current owner      : admin.
Current security level: Administrative.
```

Figure 3-5 “Who Am I” Screen

Configuration Management

The RM controller provides uniform configuration management across heterogeneous SAN. Status of host ports and storage assets are continuously being monitored.

Configure and Monitor Status of Host Ports

The status information of the host ports can be obtained at any time.

The “HOST” command will display the current settings and status for each host port (Figure 3-6). It will also display a list of the users currently logged into the system. An unauthorized user will be given the user name “Anonymous”.

The “PORT=X|ALL” parameter specifies the specific host port(s) (1..4) to be affected when used in combination with any of the other parameter: ID, TIMEOUT, UART, WWN. Default is to apply changes to ALL host ports.

```

RM660 [1]: host
                                Host Port Configuration
Host      Hard   Current   Port Speed   Timeout
Port     Loop ID Loop ID   Desired/Actual seconds   World Wide Name   Loop Status
-----
  1       EF     EF       1Gbps/1Gbps   30      2100CB5CBA7F5F1F   Good
  2       EF     EF       1Gbps/-NA-   30      2200CB5CBA7F5F1F   Not connected
  3       EF     EF       1Gbps/1Gbps   30      2300CB5CBA7F5F1F   Good
  4       EF     EF       1Gbps/1Gbps   30      2400CB5CBA7F5F1F   Good

                                Current Logins
      User      Port  Frame  S_ID  World Wide Name   Login
-----
  Client1     1    2048   2    10000000C92135FE  WED JAN 28 17:04:32 2004
  Anonymous   2    2048   E8    200000E08B002D4C  WED JAN 28 17:03:13 2004
    
```

Figure 3-6 Host Ports Status Screen

Host ID

The “HOST ID=<new ID>” command changes the hard loop ID of a host port. This is the Fibre Channel AL_PA value which will be used by the host port. The system will select 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.

Table 3-2 shows the loop ID to AL_PA value correspondence.

Table 3-2 Loop ID to AL_PA Correspondence

Loop ID (dec)	AL_PA (hex)	Loop ID (dec)	AL_PA (dec)						
0	EF	26	BA	51	8F	76	5C	101	33
1	E8	27	B9	52	88	77	5A	102	32
2	E4	28	B6	53	84	78	59	103	31
3	E2	39	B5	54	82	79	56	104	2E
4	E1	30	B4	55	81	80	55	105	2D
5	E0	31	B3	56	80	81	54	106	2C
6	DC	32	B2	57	7C	82	53	107	2B
7	DA	33	B1	68	7A	83	52	108	2A
8	D9	34	AE	69	79	84	51	109	29
9	D6	35	AD	60	76	85	4E	110	27
10	D5	36	AC	61	75	86	4D	111	26
11	D4	37	AB	62	74	87	4C	112	25
12	D3	38	AA	63	73	88	4B	113	23
13	D2	39	A9	64	72	89	4A	114	1F
14	D1	40	A7	65	71	90	49	115	1E
15	CE	41	A6	66	6E	91	47	116	1D
16	CD	42	A5	67	6D	92	46	117	1B
17	CC	43	A3	68	6C	93	45	118	18
18	CB	44	9F	69	6B	94	43	119	17
19	CA	45	9E	70	6A	95	3C	120	10
20	C9	46	9D	71	69	96	3A	121	0F
21	C7	47	9B	72	67	97	39	122	08
22	C6	48	98	73	66	98	36	123	04
23	C5	49	97	74	65	99	35	124	02
24	C3	50	90	75	63	100	34	125	01
25	BC								

Host WWN

The “HOST WWN=<x|0|DEFAULT>” command can be used to over-ride the system ID and specify a different World Wide Name for a host port. This parameter is entered as an 64-bit hex value. Default WWN is 0.

Host Status

The “HOST STATUS” command will display the loop status of each host port and a count of the Fibre Channel errors encountered on each port (Figure 3-7). Use the “HOST STATUSCLEAR” command to reset the error counts.

```

RM660 [1]: host status

                Host Port Status

                Port 1      Port 2      Port 3      Port 4
-----
Running Disparity      7         483         0         0
Word Alignment         7         429         0         0
Code violations        7         482         0         0
Parity Errors          0          0          0         0
Force Deletes          2          0          0         0
Force Inserts          4         306         0         0
Loss of Sync.          4         485         0         0
Link Failures          4         165         0         0
CRC Errors             0          0          0         0
Unexpected EOF         0          0          0         0

Loop Status           Good Not connected Not connected Not connected
    
```

Figure 3-7 Host Port Status Screen

Host LILP Payload

The “HOST LIPINFO” command will display the LILP payload for all host ports (Figure 3-8).

```

RM660 [1]: host lipinfo

                LILP Payload for Host Ports:

+-----+-----+-----+-----+
| Port 1 | Port 2 | Port 3 | Port 4 |
+-----+-----+-----+-----+
| 02EF88FF | 02E8EFFF | 02E8EFFF | 00000000 |
| FFFFFFFF | FFFFFFFF | FFFFFFFF | 00000000 |
+-----+-----+-----+-----+
    
```

Figure 3-8 Host Port LILP Payload Screen

Host Array Parity Checking

The “HOST ARRAYPARITY=ON|OFF>” command enables/disables host array parity checking on all the host ports. This feature is normally used for testing only. Default setting is ON.

Host Port Speed

The “HOST SPEED” command lets you display and change the port speed on the host port(s). You will be prompted for the desired speed as well as for the choice of host port(s).

Configure and Monitor Status of Storage Assets

Disk and Channel Information

The “DISK” command will display the current disk configuration and the status of the ten disk channels on the RM controller (Figure 3-9). This example is for the RM660.

```
RM660 [1]: disk
    Disk Channel Status
Disk Channel A healthy.
Disk Channel B healthy.
Disk Channel C healthy.
Disk Channel D acquiring loop synchronization.
Disk Channel E healthy.
Disk Channel F healthy.
Disk Channel G healthy.
Disk Channel H healthy.
Disk Channel P healthy.
Disk Channel S healthy.

    All disks are healthy.

Disk write caching is Enabled.
Audio/Visual settings Disabled.
Disk command timeout: 15 seconds.
Disk command AV timeout: 0
Fast AV reads: Enabled
Ordered Tag Count: 0

Disk commands outstanding: 0
```

Figure 3-9 RM660 Disk Channel Status Screen

If the channel status is “acquiring loop synchronization”, this may indicate a channel problem. Refer to “Multiple Drive Failures and Channel Failures” on page 124 for recovery information.

To get information about a specific disk, enter command:

```
DISK INFO=<tier><channel>
```

The “DISK LIST” command will display a list of the disks installed in the system and indicate how many were found.

The “DISK SCAN” command will check each disk channel in the system for any new disks and verify that the existing disks are in the correct location. This will also start a rebuild operation on any failed disks which pass the disk diagnostics.

The “DISK STATUS” command will display the loop status of each disk channel and a count of the Fibre Channel errors encountered on each channel (Figure 3-10). To reset the Fibre Channel error counts, issue command: DISK STATUSCLEAR.

RM660 [1]: disk status										
Disk Channel Status										
LUN	A	B	C	D	E	F	G	H	P	S
Link Fail	0	0	0	0	0	0	0	0	0	0
Lost Sync	0	0	0	0	0	0	0	0	0	0
Word Align	0	0	0	0	0	0	0	0	0	0
Disparity	0	0	0	0	0	0	0	0	0	0
Violation	0	0	0	0	0	0	0	0	0	0
Deletes	0	0	0	0	0	0	0	0	0	0
Inserts	0	0	0	0	0	0	0	0	0	0
Unexp EOF	0	0	0	0	0	0	0	0	0	0
Link Fail	0	0	0	0	0	0	0	0	0	0
CRC Error	0	0	0	0	0	0	0	0	0	0
Parity Er	0	0	0	0	0	0	0	0	0	0
DID Error	0	0	0	0	0	0	0	0	0	0
ARB TMOs	0	0	0	0	0	0	0	0	0	0

Figure 3-10 Disk Status Screen

The “DISK LIPINFO” command will show the last LILP payload for all disk channels.

The “DISK DEFECTLIST=<tier><channel> G|P|A” command lets you display the defect list information for a specified disk (Figure 3-11). The second parameter indicates the type of defect list information to be displayed, where “G” indicates the Grown list, “P” indicates the Permanent list, and “A” indicates both Grown and Permanent lists.

```
RM660 [1]: disk defectlist=1a a
Combined Defect List for Drive 1A
```

Item	Physical Sector
1	0002a234
2	0002e930
3	00006c34

Figure 3-11 Disk Defect List Screen

The “DISK RLS=<tier><channel>” command will display Read Link Error Status Block information for the specified drive. If neither the tier nor the channel are specified, the RLS information will be requested from all drives. If only the channel is specified, information will be requested from all drives on that channel. If only the tier is specified, information will be requested from all drives on that tier. The RLS information consists of the following items:

- Link: Link Failure Count
- Sync: Loss of Synchronization Count
- Signal: Loss of Signal Count
- PSPE: Primitive Sequence Protocol Error
- ITW: Invalid Transmission Word
- CRC: Invalid CRC Count
- F7Init: Lip F7 initiated count
- F7Rec: Lip F7 received count
- F8Init: Lip F8 initiated count
- F8Rec: Lip F8 received count

Tier View

Tiers are the basic building blocks of the RM controller. One tier contains ten drives: eight data drives (Channels A through H), one parity drive (Channel P), and an optional spare drive (Channel S). Drives that have the same AL_PA across all ten channels are put on the same tier. The tier that contains the drives with the highest AL_PA value is recognized as Tier #1. 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.

The “TIER” command lets you display the current status and configuration of the tiers in the system (Figure 3-12).

```

RM660 [1]: tier
Tier Configuration
Capacity  Space Available
Tier  Owner  (Mbytes)  (Mbytes)  Disk Status  Lun List
-----
  1    1      280012    271820    ABrDEFGHPS  0 4
  2    1      280012    271820    ABCDEFGHPS  0 4
  3    1      280012    271820    ABCDEFGHPS  1 4
  4    1      280012    271820    ABCDEFGHPS  1 4
  5    1      280012    271820    ABCDEFGHPS  2 4
  6    1      280012    271820    ABCDEFGHPS  2 4
  7    1      280012    271820    ABCDEFGHPS  3 4
  8    1      280012    271820    ABCDEFGHPS  3 4

Automatic disk rebuilding is Enabled
System rebuild extent: 32 Mbytes
System rebuild delay: 60
System Capacity 2240096 Mbytes, 2174560 Mbytes available.
    
```

Figure 3-12 Tier Configuration Screen

The tiers’ total and available capacities are shown under the “Capacity” and “Space Available” columns respectively.

Disk Status shows the status of each disk on the tier. A letter <ABCDEFGHPS> 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. “?” 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. “!” indicates that the disk is in the wrong location.

The rate of rebuild and format operations can be adjusted with the DELAY and EXTENT parameters.

Tier Configuration

To display the configuration information about all of the tiers, enter the “TIER CONFIG” command (Figure 3-13).

```

RM660 [1]: tier config

```

Tier	Owner	Capacity (Mbytes)	Status	Installed	Healthy	Fail	Repl	Spare	H	A	WC
1	1	280012	A.CDEFGHP	ABCDEFHGHP	A CDEFGHP	B	N	1 1	N	N	OFF
2	1	280012	ABCDEFHGHP	ABCDEFHGHP	ABCDEFHGHP	N	N	2 1	N	N	OFF
3	1	280012	ABCDEFHGHP	ABCDEFHGHP	ABCDEFHGHP	N	N	3 3	Y	Y	OFF

```

Other RM660 host busy: 0.
Other RM660 disk busy: 0.

```

Figure 3-13 Displaying the Tier Configuration Information

In the “Fail” column, an “N” indicates that there is no failed disk in that tier. A character in the range <ABCDEFHGHP> indicates that the disk at that channel has failed and is not replaced by a spare disk. Note that if a disk has failed and is replaced by a spare drive, the system will also display an “N”.

In the “Repl” column, an “N” indicates that no disk in the tier is replaced by a spare disk. A character in the range <ABCDEFHGHP> indicates that the disk at that channel is being replaced by a spare disk.

In the “Spare” column, the left number is only valid if a disk is being replaced in the tier (A, B, C, D, E, F, G, H, or P displayed in the “Repl” column) and this number indicates which tier's spare is used. The right number is only valid if a spare disk is healthy (a “Y” displayed in the H column) but not available (an “N” displayed in the A column) and this number indicates where the spare disk is used.

In the “H” (Healthy) column, an “N” indicates that there is no spare disk installed on that tier. A “Y” indicates a spare disk is installed and it is healthy.

In the “A” (Available) column, a “Y” indicates that a spare disk is available on that tier and an “N” indicates that no spare disk is available.

In Figure 3-13, disk B in tier 1 has failed but it has not been replaced by a spare disk as the “Repl” column shows an “N”. In tier 3, a spare disk is healthy and available.

In Figure 3-14, both the “Fail” and “Repl” columns show a “B” for Tier 1 which indicate that disk 1B has failed and is being rebuilt. The left number in the “Spare” column tells you which tier's spare disk is being used, which is 3 in this example. The spare disk in Tier 3 has now become unavailable.

```

RM660 [1]: tier config

          Capacity
Tier Owner (Mbytes)  Status  Installed  Healthy  Fail Repl Spare H A WC
-----
  1   1   280012  ArCDEFGHP  ABCDEFGHPS  A CDEFGHP  B  B  3 1  N N OFF
  2   1   280012  ABCDEFGHP  ABCDEFGHP  ABCDEFGHP  N  N  2 1  N N OFF
  3   1   280012  ABCDEFGHPS  ABCDEFGHPS  ABCDEFGHPS  N  N  3 1  Y N OFF

Other RM660 host busy: 0.
Other RM660 disk busy: 0.
    
```

Figure 3-14 Tier Configuration Information (1)

In Figure 3-15, disk 1B is now replaced by spare disk 3S.

```

RM660 [1]: tier config

          Capacity
Tier Owner (Mbytes)  Status  Installed  Healthy  Fail Repl Spare H A WC
-----
  1   1   280012  ArCDEFGHP  ABCDEFGHPS  A CDEFGHPS  N  B  3 1  N N OFF
  2   1   280012  ABCDEFGHP  ABCDEFGHP  ABCDEFGHP  N  N  2 1  N N OFF
  3   1   280012  ABCDEFGHPS  ABCDEFGHPS  ABCDEFGHP  N  N  3 1  Y N OFF

Other RM660 host busy: 0.
Other RM660 disk busy: 0.
    
```

Figure 3-15 Tier Configuration Information (2)

LUN View

The “LUN” command displays the current status of the LUNs (Figure 3-16). “Ready” indicates that the LUN is in good condition. The percentage of completion is displayed if the LUN is being formatted or rebuilt. A status of “Unavailable” may result from multiple drive failures. “Ready [GHS]” indicates that a spare drive has been successfully swapped for one of the drives in LUN.

```

RM660 [1]: lun
                                Logical Unit Status
LUN  Label  Owner  Status  Capacity  Block  Tiers  Tier List
      (Mbytes)  Size
-----
0    vol1   1     Ready [GHS]  10002   512    1     1
1    vol2   1     Ready        10002   512    1     2
2    vol3   1     Ready        10002   512    1     3
3    vol4   1     Format 14%   10002   512    1     4

System Capacity 2240096 Mbytes, 2200088 Mbytes available.
    
```

Figure 3-16 Logical Unit Status Screen

The “LUN LIST” command displays a list of all valid LUNs in the system. The list shows the capacity, owner, status, and serial number of each LUN (Figure 3-17).

```

RM660 [1]: lun list
                                Logical Unit Status
LUN  Label  Owner  Status  Capacity  Serial
      (Mbytes)  Number
-----
0    vol1   1     Ready [GHS]  10002   00015A1300A7
1    vol2   1     Ready        10002   0001A28101A7
2    vol3   1     Ready        10002   0001A29A03A7
3    vol4   1     Ready        10002   0001A2b10400

System Capacity 2240096 Mbytes, 2200088 Mbytes available.
    
```

Figure 3-17 LUN List Screen

LUN Configuration

The “LUN CONFIG” command displays the configuration information about all the valid LUNs in the system (Figure 3-18). If there are more than 1 segment in the LUN (LUN 1.00 to LUN 1.01 in Figure 3-18), the LUN Offset column shows where each segment is placed on the LUN. If there are more than 1 segment in the LUN, the Tier Start and Tier End columns show where the LUN segments' start and end points are on the tiers.

```

RM660 [1]: lun config
                                Logical Unit Configuration
LUN  Capacity  Block   LUN   Tier   Tier   Tier
(Blocks) Size    Offset Start  End    List
-----
0    1000000    512     0     0     1FFFFF 1 2
1    2000000    512     0     200000 2FFFFF 1 2 3 4
1.00 1000000    512     0     200000 27FFFF
1.01 1000000    512    1000000 280000 2FFFFF
    
```

Figure 3-18 LUN Configuration Screen

LUN Reservations

The “LUN RESERVATIONS” command 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 (Figure 3-19). The command “LUN RELEASE” can be used to release any SCSI reservations on a LUN.

```

RM660 [1]: lun reservations
                                Current SCSI LUN Reservations
LUN  Label  Status  Reservation ID  Port  User Name
-----
0    Ready  No SCSI Reservations
1    Ready  No SCSI Reservations
2    Ready  No SCSI Reservations
    
```

Figure 3-19 LUN Reservations Screen

Adding/Removing Storage Assets

The RM controller supports up to 125 tiers (depending on individual disk enclosure’s AL_PA numbering scheme). As your storage demand grows, new tiers can easily be added without affecting system operations.

Use the “DISK SCAN” command to check 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 will automatically be 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 will display all current drive enclosure failures (Figure 3-20), detected by the SCSI Enclosure Services (SES) as well as provide a means to access SES specific functions such as disk, channel, LUN, or tier visual identification. Drive failures are not shown by the SES command (use the TIER command to view drive status).

```
RM660 1]: ses
EncID:50050CC000033C8: Power Supply 1 :DC Power Failure
```

Figure 3-20 Displaying the Current Disk Enclosure Failures

If your enclosures provide redundant SES communication paths, the error will be reported twice. “EncID” is the World Wide Name 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.

The “SES=ON” command lets you save the SES state to the parameter blocks, and startup the SES monitors.

The “SES=OFF” command lets you save the SES state to the parameter blocks, and shutdown the SES monitors.

Display Information of SES Devices

The “SES SHOWDEVICES” command allows you to display all the SES devices on all channels.

The “SES SHOWALL” command will display all configuration information for all the SES devices on all channels.

The “SES SHOW=<tier><channel>” command lets you display 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..125><ABCDEFGHPS>.

Visual Indication of Drive

The “SES IDDISK=<tier><channel>” command will provide a visual indication of the specified drive, <1..125><ABCDEFGHPS>. The status LED of the drive will blink until

the command “SES ID=OFF” is issued, which will restore the system to its original visual state.

Visual Indication of Tier

The “SES IDTIER=<tier>” command will provide a visual indication of the specified tier, <1..125>. The status LED of the drives will blink until the command “SES ID=OFF” is issued, which will restore the system to its original visual state.

Visual Indication of Channel

The “SES IDCHANNEL=<channel>” command will provide a visual indication of the specified channel, <ABCDEFGHPS>. The status LED of the drives will blink until the command “SES ID=OFF” is issued, which will restore the system to its original visual state.

Visual Indication of LUN

The “SES IDLUN=<LUN number>” command will provide a visual indication of the specified LUN, <0..127>. The status LED of the drives will blink until the command “SES ID=OFF” is issued, which will restore the system to its original visual state.

Verify Connections to SF6016 Enclosures

The “SES VERIFY_6016” command allows you to verify the connections of an SF6016 drive enclosure that is currently attached to the system. The connections must be verified before the SCSI Enclosure Services (SES) can accurately indicate a drive fault through SES commands.

The “SES RESET_6016” command will remove all of the SF6016 drive enclosures from the verified enclosures list for this RM controller unit.

Tier Mapping for Enclosures

The RM controller supports 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 RM controller to properly light the enclosure “fault” LEDs.

The RM controller currently supports the SF and SA families of drive enclosures. If you have other types of drive enclosures, the RM controller might still be able to monitor the status of your enclosures. Component failures such as power supply, fan, and drive will be reported. However, visual indication of the drives, tiers, and LUNs will not be supported.

To display the current mapping mode for the disks in the array, enter the “TIER MAP” command (Figure 3-21).

```
RM660 [1]: tier map
Current mapping mode: 4, 'SF4016 Enclosure 2 x 8'
```

Figure 3-21 Displaying the Current Tier Mapping Mode

To display the current mapping AL_PA values for each tier of disks in the array, enter the “TIER DISPLAYMAP” command (Figure 3-22). A list of 125 tiers and the corresponding AL_PA values are displayed.

```
RM660 [1]: tier displaymap
Tier      ALPA      Current mapping mode: 2, SF4012 Enclosure
-----
1         0xEF
2         0xE8
3         0xE4
4         0xE2
5         0xE1
         .
         .
123      0x26
124      0x18
125      0x17
```

Figure 3-22 Displaying the Current Mapping AL_PA Values for Tiers

To change the current tier mapping for the disks in the array:

- Enter the “TIER CHANGEMAP” command
- Then select the appropriate mapping mode for your SF enclosures and press <Enter> (Figure 3-23). For other enclosures, select mode “0”.

```
RM660 [1]: tier changemap

Supported mapping modes:
-----
0) Standard Enclosure mapping.
1) SF2012 Enclosure mapping.
2) SF4012 Enclosure mapping.
3) SF2016 Enclosure mapping.
4) SF4016 Enclosure mapping.

Current mapping mode: 4, SF4016 Enclosure.
Enter new mapping mode (0..4), 'e' to escape:
```

Figure 3-23 Tier Mapping Mode Setup Screen

- Enter the “RESTART” command for the changes to take effect.

Note: The “changemap” command should only be used when the system is first configured. Changing the mapping mode will alter all the tier information, hence, making LUN information inaccessible.

System Network Configuration

The “NETWORK” command will display the current network interface settings of the RM controller that you are inquiring about (Figure 3-24).

```

RM660 [1]: network
                                     Network Configuration
                                     =====
Gateway:                             172.16.0.254
Netmask:                              255.255.255.0
MAC Address HSTD #1:                 00:01:ff:01:00:ae
IP Address HSTD #1:                   172.16.0.1
IP Address HSTD #2:                   -- Unknown --
Services
-----
Telnet:                               ENABLED
Telnet Port HSTD #1:                  23
Telnet Port HSTD #2:                  23
API Server:                           ENABLED
API Server Port HSTD #1:              8008
API Server Port HSTD #2:              8009
SNMP:                                  ENABLED
SNMP Trap IP Address:                 0.0.0.0
Limited SNMP:                         DISABLED
Syslog:                                ENABLED
Syslog IP Address:                    172.16.0.253
Syslog Port HSTD #1:                  514
Syslog Port HSTD #2:                  514

```

Figure 3-24 Current Network Configuration Screen

The “NETWORK USAGE” command will display the address resolution protocol map. It also displays ICMP (ping), general network, and IP, TCP, and UDP layer statistics.

To change the IP address of the RM controller, enter command:

```
NETWORK IP=<new address>
```

The system must be restarted before the changes will take effect.

To change the netmask of the RM controller, enter command:

```
NETWORK NETMASK=<aaa.bbb.ccc.ddd>
```

To set the current gateway in the network routing table to the supplied Internet address, enter command:

```
NETWORK GATEWAY=<aaa.bbb.ccc.ddd>
```

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=<with no Internet address> will clear out the current gateway.

To ping destination with a single packet, enter command:

```
NETWORK PING=<aaa.bbb.ccc.ddd>
```

Telnet

To enable/disable the Telnet capability on the RM controller, enter command:

```
NETWORK TELNET=ON|OFF
```

The system must be restarted before the changes will take effect.

Note: To affect Telnet session availability only temporarily during the current power-cycle, refer to the “TELNET” command (see “Remote Login Management” on page 105).

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

To change the Telnet port number for the current RM controller, use command:

```
NETWORK TELNETPORT=<port number>
```

The system must be restarted before the changes will take effect. Valid ports are 0 to 32768. Note, however, that 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

To enable/disable the SNMP functionality on the RM controller, use command:

```
NETWORK SNMP=ON|OFF
```

The system must be restarted before the changes will take effect.

To specify whether the SNMP functionality will only report component-level information, or all levels of information, enter command:

```
NETWORK LIMIT_SNMP=ON|OFF
```

Default setting is OFF.

To change the destination IP address for SNMP trap packets, enter command:

```
NETWORK TRAPIP=<aaa.bbb.ccc.ddd>
```

The system must be restarted before the changes will take effect.

To enable/disable the Syslog capability, use command:

```
NETWORK SYSLOG=ON|OFF
```

Note: It is highly recommended that “NETWORK SYSLOG” to 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 “syslogd” program configured to rotate based on log size rather than date.

To change the destination IP address for syslog packets, enter command:

```
NETWORK SYSLOGIP=<aaa.bbb.ccc.ddd>
```

Both RM controllers in the couplet pair will share the same syslog destination IP address but each RM controller can specify a different destination port.

To change the destination port number for syslog packets for the current RM controller, use command:

```
NETWORK SYSLOGPORT=<port number>
```

Both RM controllers in the couplet pair will share the same syslog destination IP address but each RM controller can specify a different destination port. Valid ports are 0 to 32768. Note however, that 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 RM controller sends syslog messages via the local7 (23) facility.

Please refer to “Remote Management of the RM Controller” on page 114 for information on how to set up the Telnet and SNMP functionality on your host computer.

API Server Connections

To enable/disable the API server capability, use command:

```
NETWORK API_SERVER=ON|OFF
```

The SGI 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, refer to the “API” command (see “API Server Connections” on page 110).

To specify the API Server port number for the current RM controller, use command:

```
NETWORK API_PORT=<port number>
```

The system must be restarted before the changes will take effect. Valid ports are 0 to 32768. Note, however, that 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 will display the current routing table of the system (Figure 3-25) and allows the administrator to change it. The routing table describes how the RM controller can communicate with the hosts on other networks.

```

RM660 [1]: route
Gateway 172.16.0.254

Permanent Routing Table:
=====
destination    gateway
-----
0.0.0.0        172.16.0.254
-----

Current Routing Tables:
=====

ROUTE NET TABLE
destination    gateway    flags    Refcnt    Use    Interface
-----
0.0.0.0        172.16.0.254    3        0        46569    fei0
172.16.0.0     172.16.0.1     101       0         3        fei0
192.168.0.0    172.13.0.254   3         0         1        fei0
-----

ROUTE HOST TABLE
destination    gateway    flags    Refcnt    Use    Interface
-----
127.0.0.1     127.0.0.1     5         1         2        lo0
-----

```

Figure 3-25 Current RM Controller Routing Table

To add gateways to the routing table, use command:

```
ROUTE ADD=<aaa.bbb.ccc.ddd> GATEWAY=<aaa.bbb.ccc.ddd>
```

Up to 6 permanent routes can be added to the routing 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 command:

```
ROUTE ADD=90.0.0.0 GATEWAY=91.0.0.3
```

To delete gateways from the routing table, use command:

```
ROUTE DEL=<aaa.bbb.ccc.ddd> GATEWAY=<aaa.bbb.ccc.ddd>
```

To set the current gateway in the network routing table to the specified Internet address, use command:

```
ROUTE GATEWAY=<aaa.bbb.ccc.ddd>
```

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.

Restarting the SGI InfiniteStorage RM610/RM660

System Restart

The “RESTART” command will perform a restart on the RM controller which the command is issued on.

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

When you change the RM controller’s IP address, for example, you must restart the RM controller to initialize the new values. Any changes you make will not take effect until you perform a restart.

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

System Shutdown

The “SHUTDOWN” command will shutdown the RM controller which the command is issued on.

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

All hosts and users actively using the RM controller should be safely shutdown before using this command. The RM controller will halt all I/O requests and save the data to the disks. It will then ask if the disks should be spundown. Disks should be spundown before they are moved. The unit can be safely turned off after using this command.

Note: Be sure to use “Shutdown” whenever you need to power down the RM controller for maintenance. Shutdown flushes any data left in the cache and prepares the RM controller for an orderly shutdown. For couplet RM controller configuration, issue command to both RM controllers.

Setting the System’s Date and Time

Settings are automatically adjusted for leap years. 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.

System Date

To display the current system date, enter command:

```
DATE
```

To change the system date to March 14, 2004, for example, type:

```
DATE 3 14 2004
```

System Time

To display the current system time, enter command:

```
TIME
```

To change the system time to 2:13:10pm, for example, type:

```
TIME 14:13:10
```

Saving the RM Controller’s Configuration

The “SAVE” command can be used to save the system configuration to non-volatile memory (Figure 3-26).



```
RM660 [1]: save
Saving system parameters.. Done.
```

Figure 3-26 Saving System Parameters 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 will display the current status of the system parameters (Figure 3-27).

```
RM660 [1]: save status
```

System Parameters Status			
Number	Revision	Updates	Last Update
0	500	306	WED FEB 04 14:49:42 2004
1	500	185	WED JAN 28 17:23:08 2004

Figure 3-27 Current System Parameters Status Screen

Restoring the System’s Default Configuration

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

Warning: This command will delete all LUN configuration and data unconditionally. Do not issue this command without guidance from SGI’s Technical Support.

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 RM controller 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 125 tiers, depending on individual disk enclosure’s AL_PA numbering scheme (please refer to Glossary for definition of tier). 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 128 LUNs are supported in total. Each LUN can also be subdivided into up to 64 smaller equally-sized LUN segments, giving a total of 8192 LUN segments. The LUN segments of a LUN are managed together and share the same characteristics. In couplet configuration, LUNs are “owned” by the RM controller via which they are created.

You can add and remove LUNs without affecting system operations. The “LUN” command will display the current Logical Unit Status (Figure 3-28).

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

```

RM660 [1]: lun
                                Logical Unit Status
LUN  Label  Owner  Status  Capacity  Block  Tiers  Tier List
      (Mbytes)  Size
-----
0    vol1   1      Ready [GHS]  10002  512    1      1
1    vol2   1      Ready        10002  512    1      2
2    vol3   1      Ready        10002  512    1      3
3    vol4   1      Format 14%   10002  512    1      4
-----
System Capacity 277810 Mbytes, 237802 Mbytes available.
    
```

Figure 3-28 Logical Unit Status Screen

Creating and Formatting a LUN

To add a LUN, enter command:

LUN ADD=x

where “x” is the Logical Unit Number <0..127>. The system will prompt you for all the

necessary information to create the LUN and indicate if the LUN was successfully added to the system (see “Configuring the Storage Arrays” on page 36 for LUN creation example).

The required LUN information includes:

- Capacity (in MBytes) - default is to use all available capacity
- Number of LUN segments in a LUN group
- Number of tiers - default is to use all tiers
- Block size (in Bytes) - default is 512Bytes
- Label - may contain up to 12 characters

A LUN must be formatted before it can be used. The “FORMAT” command can be used to format the LUNs after they are created.

```
LUN FORMAT=x
```

where “x” is the LUN <0..127>. This will perform a destructive initialization on the specified LUN by over-writing all the data on the LUN with zeroes. A LUN segment in a LUN group can be specified by “x.y”, where “x” is the LUN group in the range <0..127> and “y” is the LUN segment of the group in the range <0..63>. If all of the segments in a LUN group are unformatted then the entire LUN group will be formatted even if only one LUN segment was specified.

The rate of format can be adjusted using the DELAY and EXTENT parameters of the TIER command (see “Resources Allocation” on page 96 for more information).

Interrupting a LUN Format Operation

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

- LUN PAUSE - to pause the current format operations.
- LUN RESUME - to release the paused format operations.
- LUN STOP - to abort all the current format operations.

Changing a LUN Label

If you need to change the label of a LUN, enter command:

```
LUN LABEL
```

Select the LUN to change and enter the new label. A LUN label may contain up to 12 characters (Figure 3-29).

```
RM660 [1]: lun label
Enter the LUN to label (0..127), 'e' to escape:
0
Enter a new label for LUN 0, up to 12 characters:
voll

          Logical Unit Status
LUN  Label  Owner  Status  Capacity  Block  Tiers  Tier List
          (Mbytes)  Size
-----
0     voll   1     Ready [GHS]  10002  512    1     1
1     1       2     Ready        10002  512    1     2
2     2       1     Ready        10002  512    1     3
3     3       2     Ready        10002  512    1     4

          System Capacity 277810 Mbytes, 237802 Mbytes available.
```

Figure 3-29 Changing a LUN Label Screen

Moving a LUN (Dual Mode Only)

If you need to change the ownership of a LUN from one RM controller to the partner when the units are in dual mode, enter command:

LUN MOVE=x

where “x” is the Logical Unit number <0..127> (Figure 3-30).

“LUN Move” can only be successfully accomplished if the LUN does not share tiers with other LUNs. Thus, if LUN 0 is on tiers 1,2,3 and LUN 1 is on tiers 4,5,6, both LUNs can be moved. If, however, LUN 0 is on tiers 1,2,3 and LUN 1 is on tiers 3,4,5, neither LUN can be moved.

```
RM660 [1]: lun move=0

LUN 0 is owned by this RM660.

Do you want to move ownership to the OTHER RM660? (y/n):
```

Figure 3-30 Moving a LUN Screen

Deleting a LUN

If you need to delete a LUN from the system, enter command:

```
LUN DEL=x
```

where “x” is the LUN <0..127>. *This will erase all the data on the LUN.* You can only delete a LUN which is owned by the RM controller that you are logged into.

SCSI Reservations

The “LUN RELEASE=x|x.y” command 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..127>. A LUN segment in a LUN group can be specified by “x.y”, where “x” is the LUN group in the range <0..127> and “y” is the segment of the group in the range <0..63>.

The “LUN START” command 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 RM controller’s automatic drive failure recovery procedures ensure that absolute data integrity is assured while operating in degraded mode.

In the event of a drive failure, the RM controller will automatically initiate a drive rebuild using a spare drive if the “autorebuild” function has been enabled. You may use the “TIER” command to display the current setting (Figure 3-31). The rebuild operation can take up to several hours to complete, depending on the size of the disk and rate of rebuild (see “Resources Allocation” on page 96 for information on how to adjust the rate of rebuild).

```

RM660 [1]: tier

                                Tier Configuration

Tier  Owner      Capacity  Space Available  Disk Status  Lun List
-----
  1      280012      280012      280012      ABCDEFGHPS
  2      280012      280012      280012      ABCDEFGHPS
  3      280012      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 3-31 Automatic Disk Rebuilding Parameter

Use the “TIER AUTOREBUILD=ON|OFF” command to enable/disable the automatic disk rebuild function. A disk will only be replaced by a spare disk if it fails and Autorebuild is ON (which is 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 (see “Recovering from Drive Failures” on page 121 for further information on drive failure recovery).

Manual Drive Rebuild

To initiate a rebuild on a specific drive, enter command:
 DISK REBUILD=<tier><channel>

This operation will reconstruct data on the replacement drive and restore a degraded LUN to healthy status.

Drive Rebuild Verify

The “DISK REBUILDVERIFY=ON|OFF” command 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 rebuild correctly. Default is OFF. Note that 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 command:

```
DISK REPLACE=<tier><channel>
```

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

Interrupting a Rebuild Operation

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

- `TIER PAUSE` - to pause the current rebuild operations.
- `TIER RESUME` - to release the paused rebuild operations.
- `TIER STOP` - to abort all the current rebuild operations.

Couplet RM Controller Configuration (Cache/Non-Cache Coherent)

There are two primary couplet RM controller configurations: cache coherent and non-cache coherent.

The “DUAL” command will display information about couplet system configuration (Figure 3-32).

```
RM660 [1]: dual
                Dual RM660 Configuration
                RM660 1                      RM660 2
-----
Label           RM660 Unit[1]                  RM660 Unit[2]
Status          Healthy                       Healthy

Dual communication:    established.
Ethernet communication: established.
Cache coherency:      established.
Cache coherency timeout: 2
```

Figure 3-32 Couplet RM Controller Configuration Screen

Non-Cache Coherent (firmware version 1.06 and above)

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

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

Users, via mapping, can be assigned any combination of LUNs. In a healthy environment, the users will only see LUNs owned by the RM 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. RM controller 1 owns LUNs 0 and 1 while RM controller 2 owns LUN 2. The user is physically connected to RM controller 1, thus, it will only see LUNs 0 and 1. The user will not be able to access LUN 2. If the user was physically connected to RM controller 2, the reverse would be true, only LUN 2 would be accessible. When an RM controller fails, the user will be given access to all mapped LUNs regardless of the physical connection.

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

Cache Coherent (firmware version 2.02 and above)

In this configuration, each RM controller only owns tiers. This is only for disk failures and rebuilds. Spare drives are owned by individual RM controllers, according to tier ownership.

Each RM controller can access all LUNs. Each user, regardless of physical connection, can see all mapped LUNs.

The couplet RM controller communication occurs over the internal UART and external Ethernet. If the RM controllers detect an Ethernet failure, RM controller 2 will be failed. (This means that an external event can cause an RM controller to fail even though the RM controller may be perfectly fine.) Therefore, it is highly recommended the RM controller Ethernet resides on a private Ethernet segment.

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

To enable/disable the cache coherency function, enter command:

```
DUAL COHERENCY=ON|OFF
```

Default is dual coherency disabled which is the non-cache coherent configuration.

The “DUAL TIMEOUT=X” command allows you to set the cache coherency timeout for cache node requests in hundreds of milliseconds. Valid range is 2 to 255 hundred milliseconds. Default is 2 hundred milliseconds.

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

Fail / Restore the Other RM Controller Unit in the Couplet Pair

To fail the other RM controller unit in the system to perform maintenance, for example, enter command:

```
DUAL FAIL
```

The healthy RM controller unit will take ownership of all the LUNs/tiers from the failed RM controller unit.

To restore the other RM controller unit in the system to healthy status after failure recovery, for example, enter command:

```
DUAL HEAL
```

Ownership of LUNs/tiers are transferred back to the RM controller unit.

Labeling the RM Controller Unit(s)

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

To change the label, enter command:

```
DUAL LABEL=1|2
```

Then select which unit you want to rename (Figure 3-33). When prompted, type in the

new label for the selected unit. The new name will be displayed. Entering “DEFAULT” will restore the label of the unit to its default setting.

```
RM660 [1]: dual label

Enter the number of the HSTD you wish to rename.
  LABEL=1 for HSTD 1, Test System[1]
  LABEL=2 for HSTD 2, Test System[2]

HSTD: 1
Enter a new label for HSTD 1, or DEFAULT to return to the default label.
Up to 31 characters are permitted.
Current HSTDname: Test System[1]
New HSTDname: Ssystem[1]
```

Figure 3-33 Labeling an RM Controller Unit

Performance Management

The RM controller offers great flexibility in optimizing performance with extensive monitoring and reporting capability.

Optimization of I/O Request Patterns

The RM controller manages the pre-fetch and cache efficiency by LUN.

Display Current Cache Settings

The “CACHE” command will display the current cache settings for each LUN in the system (Figure 3-34).

```
RM660 [1]: cache
```

Current Cache Settings						
LUN	Write Caching	Maximum Prefetch	MF Bit	Prefetch Ceiling	Read Priority	Write Priority
0	Enabled	x1	On	65535	high	high
1	Enabled	x1	On	65535	high	high
2	Enabled	x1	On	65535	high	high
3	Enabled	x1	On	65535	high	high

writeback limit: 75%

640.0 Mbytes of Cache Installed
(4096 Segments of 128 Kbytes)

Figure 3-34 Current Cache Settings Screen

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 by default. Valid LUN’s are 0 to 127. Default is to apply to all LUNs.

Cache Segment Size

System performance can be optimized by changing the cache segment size to match the size of the host I/O requests. 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 the best 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 KBytes. Valid segment sizes are 64, 128, 256, 512, 1024 and 2048. The default value is 128. 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.

Use the “CACHE WRITEBACK=ON|OFF” command to enable/disable writeback caching for the specified LUN. Default setting is ON.

The “`CACHE WRITELIMIT=x`” command 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 to 100. Default value is 75.

Prefetch Settings

When the system receives a request, it can read more data than it 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.

Use the “`CACHE PREFETCH=x`” command to set 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.

Use the “`CACHE MF=ON|OFF`” command to enable/disable the Multiplication Factor 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.

Use “`CACHE MAX=x`” command to set 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

The “`CACHE DEFAULTS`” command will load the default settings for all of the cache parameters for the specified LUN’s.

Disk Configuration Settings

The writeback cache and disk timeout settings can be configured for the system. The “DISK” command will display the current disk configuration settings (Figure 3-35).

```
RM660 [1]: disk

    Disk Channel Status
Disk Channel A healthy.
Disk Channel B healthy.
Disk Channel C healthy.
Disk Channel D healthy.
Disk Channel E healthy.
Disk Channel F healthy.
Disk Channel G healthy.
Disk Channel H healthy.
Disk Channel P healthy.
Disk Channel S healthy.

    All disks are healthy.

    Disk write caching is Enabled.
    Audio/visual settings Disabled.
    Disk command timeout: 10 seconds.
    Disk command AV timeout: 1
    Fast AV reads: Enabled
    Ordered Tag Count: 0

    Disk commands outstanding: 0
```

Figure 3-35 RM660 Disk Configuration Setting Screen

Use the “DISK TIMEOUT=*x*” command to set the disk timeout for an I/O request in seconds. Valid range is 1 to 512 seconds. Default setting is 21 seconds.

Use the “DISK WRITECACHE=ON|OFF” command to enable/disable writeback caching on the disks. Writeback caching allows the disks to increase the performance of write I/O requests by storing the data in cache and saving the data at a later time. Default is ON.

Audio/Visual Settings of the System

The audio and visual 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 will display information about the audio/visual settings of the system (Figure 3-36).

```

RM660 [1]: av
Current LUN Audio/Visual Settings
LUN Label      Fast AV      Write      Maximum
                  Caching     Prefetch
-----
0              Off          On         x 1
1              Off          On         x 1
2              Off          On         x 1

Disk Audio/Visual settings are Disabled, Using disk defaults
Early Error Recovery:      Enabled
Automatic Read Reallocation: Disabled
Read Retry Count:         1
Write Retry Count:         1
Recovery Time Limit:      65535
Ordered Tag Count:        0

```

Figure 3-36 Current Audio/Visual Settings

Use the “AV FASTAV=ON|OFF” command to enable/disable 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 basis. Use LUN=x command to change the settings for a single LUN. Default setting is OFF.

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

Use the “AV DISKAV=ON|OFF” command to enable/disable the disk audio/video options for streaming data. When enabled, the disks will adjust the disk parameters to minimize the latency for data transfers by disabling non-essential features which may impact performance. Note that 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.

Note: When AV mode is enabled, the RM controller will not retry check conditions on disks. Since this increases the risk of disk failures, this should be used in AV environments only.

The “AV TIMELIMIT=x” command specifies what the recovery time limit is for the drives when DISKAV is enabled. This is the maximum amount of time that a disk can use for the data error recovery in one millisecond increments. The parameter is located in bytes 10 and 11 of the Read-Write Error Recovery mode page on the disks. The recovery time limit of each disk will be set to its default value when DISKAV=OFF. Note that 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. Valid range is 0 to 65535. Default setting is 65535.

The “AV ARRE=ON|OFF” command specifies what the Automatic Read Reallocation Enabled (ARRE) bit will be set to on the drives when DISKAV is enabled. When enabled, the disks will be automatically reallocate defective data blocks during read operations. When disabled, the disks will not reallocate defective data blocks during read operations. Automatic Write Reallocate Enabled (AWRE) is always enabled. The parameter is located in bit 6, byte 2 of the Read-write error recovery mode page on the disks. The ARRE of each disk will be set to its default value when DISKAV=OFF. Note that changing the disk parameters can adversely affect the I/O operation of the system. This parameter should only be adjusted when the systems is idle. Default setting is OFF.

The “AV ERR=ON|OFF” command specifies what the Early Error Recovery (EER) bit will be set to on the drives when DISKAV is enabled. When enabled, the disks disk shall perform the most expedient form of error recovery first. When disabled, the disks will use an error recovery procedure that minimizes the risk of mis-detection or mis-correction. The parameter is located in bit 3, byte 2 of the Read-write error recovery mode page on the disks. The EER of each disk will be set to its default value when DISKAV=OFF. Note that 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 ON.

The “AV READRETRY=X” command specifies what the read retry count will be set to on the drives when DISKAV is enabled. This field indicates the number of times the disks will attempt its recovery algorithm during a read operation. This parameter is located in byte 3 of the Read-write error recovery mode page on the disks. The retry count of each disk will be set to its default value when DISKAV=OFF. Note that 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 1.

The “AV WRITERETRY=X” command specifies what the write retry count will be set to on the drives when DISKAV is enabled. This field indicates the number of times the disks will attempt its recovery algorithm during a write operation. This parameter is located in byte 8 of the Read-write error recovery mode page on the disks. The retry count of each disk will be set to its default value when DISKAV=OFF. Note that 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 1.

The “AV ORDEREDQUEUE=X” command enables the RM controller to use ordered tags when communicating with drives. The value “x” correlates to percentage of ordered tags. 0 indicates no ordered tags, 1 indicates 100% ordered tags, 2 indicates 50% ordered tags, 3 indicates 33% ordered tags, 4 indicates 25% and so on. Valid range is 0 to 255. Default value is 0.

The “AV DISKUPDATE” command tells the system to recheck all of the mode parameters for the disks in the system. This allows the user to update the disk mode parameters after changing several of the AV parameters instead of changing them one at a time.

Locking LUN in Cache

The RM controller provides the ability for the administrator to lock LUNs 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 RM 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.

Initial Cache

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

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

- RM controller reads data from disks
- RM controller locks data in cache
- RM controller 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’ data. Once cache has been allocated to a locked LUN, however, it cannot be used by an unlocked LUN.

Once the size of the locked LUNs exceeds 50% of the total cache, the RM controller will have to 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, LUN 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
-------------------	----------------	----------------	----------------

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

- RM controller determines which data to remove from locked portion of cache, using the LRU algorithm:
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 since it is the least recently used data.
- RM controller reads data from disks
- RM controller locks data in cache
- RM controller 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
-------------------	----------------	----------------	----------------

Locking / Unlocking a LUN

To lock a LUN in the data cache, enter command: LUN LOCK=x where "x" is the Logical Unit number <0..127> (Figure 3-37).

```

RM660 [1]: lun lock=0

          Logical Unit Status
LUN Label Owner   Status   Capacity   Block   Tiers   Tier List
          (Mbytes)   Size
-----
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 3-37 Logical Unit Status - LUN Locked in Cache

To unlock a LUN and release the cache locked by the LUN, enter command “LUN UNLOCK=x”.

System Performance Statistics

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

The “STATS” command will display the performance statistics for the host ports, disk channels, and cache memory (Figure 3-38). It will show the read and write performance of each of the host ports.

```

RM660 [1]: stats

```

System Performance Statistics						
	All Ports	Port 1	Port 2	Port 3	Port 4	
Read MB/s:	4.6	1.0	3.7	0.0	0.0	
Write MB/s:	36.9	24.1	12.8	0.0	0.0	
Total MB/s:	41.5	25.0	16.5	0.0	0.0	
Read IO/s:	932	17	915	0	0	
Write IO/s:	657	386	271	0	0	
Total IO/s:	1590	404	1186	0	0	
Read Hits:	70.4%	92.0%	69.9%	0.0%	0.0%	
Prefetch Hits:	46.8%	68.0%	46.4%	0.0%	0.0%	
Prefetches:	5.7%	16.7%	5.4%	0.0%	0.0%	
Writebacks:	100.0%	100.0%	100.0%	0.0%	0.0%	
Rebuild MB/s:	0.0	0.0	0.0	0.0	0.0	
Verify MB/s:	0.0	0.0	0.0	0.0	0.0	
	Total	Reads	Writes	Pieces	Reads	Writes
Disk IO/s:	576	334	242	1:	4995678	2636344
Disk MB/s:	83.2	42.8	40.4	2:	361919	253470
Disk Pieces:	13335378	6250849	7084529	3:	23056	14424
DBD Pieces:	299526			4:	27414	253470
				5:	26043	10836
Cache Writeback data:	74.4%			6:	18301	10176
Rebuild/Verify data:	0.0%	0.0%		7:	11152	10628
Cache Data locked:	0.0%			8:	4303	451177

Figure 3-38 System Performance Statistics 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 because of prefetching. Prefetches shows the percentage of host read I/O requests to the disks which are due to prefetching.

The bottom of the screen shows the read and write performance of the disks. Disk Pieces shows the total number of disk I/O requests from the host ports. 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. BDB Pieces is the number of host I/O blocking and deblocking requests.

Cache Writeback 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 “STATS CLEAR” command will reset all the statistics back to zero.

The “STATS DELAY” command will display a histogram of the time it takes for the host and disk I/O requests to complete in 100 msec intervals (Figure 3-39).

```
RM660 [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 3-39 Command Delay Statistics Screen

The “STATS HOSTDELAY” command will display a histogram of the time delay between when the last data transfer is set ready and the host command completes (Figure 3-40). The host ready delay information is shown in 100msec intervals.

```
RM660 [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
0.2	0	0	0	1	0	0	0	0	0
0.3	0	0	0	1	2	1	0	0	0
0.4	0	0	1	2	0	2	0	0	0
0.5	0	0	0	0	0	0	0	0	0
0.6	0	0	0	1	0	2	0	0	0
0.7	0	0	0	1	0	2	0	0	0
0.8	0	0	0	0	0	0	0	0	0
0.9	0	0	0	0	0	0	0	0	0
1.0	0	0	0	0	0	2	0	0	0
1.1	0	0	0	0	0	0	0	0	0
1.2	0	0	0	0	0	1	0	0	0
1.3	0	0	0	0	2	1	0	0	0
1.4	0	0	0	0	2	1	0	0	0
1.5	0	0	0	0	0	0	0	0	0
1.6	0	0	0	0	0	0	0	0	0

Figure 3-40 Host Delay Statistics Screen

The “STATS TIERDELAY=<tier>” command will display a histogram of the time it takes for the disk I/O request to complete for all the disks in the specified (Figure 3-41). If no tier is specified, all valid tiers will be displayed. The histogram is displayed in 100msec intervals.

```
RM660 [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, 'e' to escape:

Figure 3-41 Tier Delay Statistics Screen

The “STATS DISK” command will display a histogram of which disks in the system have taken an unusually long time to complete an I/O request (Figure 3-42). The count is incremented for a disk if the 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 the 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 or it may have problems.

```

RM660 [1]: stats disk

                Delayed Disk Command Counts

  0      A      B      C      D      E      F      G      H      P      S
  1      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 3-42 Host Command Offsets Screen

The “STATS DUAL” command displays the statistics for the dual mode messages (Figure 3-43).

```
RM660 [1]: stats dual
```

Dual Message Statistics

Message	Total	Msgs/s
Lock requests	326839	0
Release requests	115	0
Lock acknowledges	301310	0
Lock requests received	301310	0
Release requests received	208	0
Lock acknowledges received	326840	0
Lock releases received	279194	0
Lock releases	294583	0
Total Messages sent	922847	0
Total Messages received	907552	0

Figure 3-43 Dual Message Statistics Screen

The “STATS LENGTH” command will display a histogram of the length of the host I/O requests in 16KByte intervals (Figure 3-44).

```
RM660 [1]: stats length
```

Command Length Statistics

Length Kbytes	Port 1		Port 2		Port 3		Port 4	
	Reads	Writes	Reads	Writes	Reads	Writes	Reads	Writes
> 0	0	0	55499	448D4	554EB	9FDF	0	0
> 16	0	0	0	1	2	1	0	0
> 32	0	0	0	1	2	1	0	0
> 48	0	0	1	2	0	2	0	0
> 64	0	0	228F9	23A45	22B43	2239C	0	0
> 80	0	0	0	1	0	2	0	0
> 96	0	0	0	1	0	2	0	0
> 112	0	0	71372	82626	15C9C	17DA5	0	0
> 144	0	0	0	0	0	0	0	0
> 160	0	0	0	0	0	2	0	0
> 176	0	0	0	0	0	0	0	0
> 192	0	0	0	0	0	1	0	0
> 208	0	0	0	0	2	1	0	0
> 224	0	0	0	0	2	1	0	0
> 240	0	0	0	0	8B7A	902E	0	0
> 256	0	0	0	0	2BB8D	35A2A	0	0

Figure 3-44 Command Length Statistics Screen

The “STATS OFFSET” command will display a histogram of the offset of the host I/O requests into the cache segments (Figure 3-45). 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.

```

RM660 [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 3-45 Host Command Offsets Screen

The “STATS REPEAT=OFF|MBS|IOS” command 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 and their rates can be adjusted to minimize their impact on system performance.

The “TIER” command will display the current rebuild parameter settings for the system (Figure 3-46).

```

RM660 [1]: tier

                        Tier Configuration
                        Capacity  Space Available
                        (Mbytes)  (Mbytes)   Disk Status  Lun List
-----
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 3-46 Displaying the Current Rebuild Parameters

The Rebuild Delay parameter controls the amount of time to wait 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. `DELAY=0` will remove all delays so the rebuild and format operations will go as fast as possible but this could significantly affect the performance of the system. *A delay value smaller than 1 is not recommended.*

The `"TIER DELAY=x"` command can be 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 60 (6 seconds).

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` will allow more data to be rebuilt in a single pass. The recommended setting is to use the default value of 32MBytes and only adjust `DELAY` to match your user load.

The `"TIER EXTENT=x"` command can be used to set the system rebuild/format extent in MBytes. The range is 1 to 128MBytes. Default is 32MBytes.

Background LUN Verify Operations

The `"LUN VERIFY"` command displays the current setting for background verify on all LUNs.

Use the `"LUN VERIFY=x|x.y"` command to turn ON background verify for LUN `"x"`, where `"x"` is the Logical Unit number $\langle 0..127 \rangle$. A LUN segment in a LUN group can be specified by `"x.y"`, where `"x"` is the LUN group in the range $\langle 0..127 \rangle$ and `"y"` is the segment of the group in the range $\langle 0..63 \rangle$.

The `"LUN VERIFY=ON|OFF"` command will prompt you for a list of LUNs on which background verify will be turned either ON or OFF. A `"VERIFY=ON"` command will both turn on the background verify for the specified LUN(s), as well as start up the Verify operation(s). A `"VERIFY=OFF"` command, however, only turns off the Background Verify setting for the specified LUN(s). Therefore, any verifies 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, if at all possible, since it can help increase disk reliability.

The “LUN DELAY=x” command sets the system verify delay value to 'x'. 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). Note that DELAY=0 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 30.

The “LUN EXTENT=x” command 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 Mbytes. Default is 16Mbytes.

SES Device Monitoring Rate

The SES device monitoring rate can be adjusted to minimize its impact on system performance.

The “SES M_WAIT” command displays the current setting in seconds (Figure 3-47).

```
RM660 [1]: ses m_wait
SES timer m_wait = 6 seconds
```

Figure 3-47 SES Device Monitoring Rate

The “SES M_WAIT=x” command sets the SES device monitoring rate for the system in seconds. Valid range is 4 to 90 seconds. The default monitoring rate is 10 seconds.

Note: Improper use of this 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.

The “HOST TIMEOUT=x” command sets the host command timeout for an I/O request in seconds. Valid range is 1 to 512 seconds. Default setting is 30 seconds.

Security Administration

The RM controller’s dual-level data security is unique and powerful. The non-host based security is maintained with scalable features including restricted management access and authentication against authorized listing. No security software is required on the host computers. (Please refer to “Administrator and User Logins” on page 49 for information regarding Telnet and serial port security.)

Each authorized user will have its customized LUN identification scheme which applies to all host ports (Figure 3-48).




Figure 3-48 Mapping Internal LUN’s to External LUN’s

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

The “place holder” LUN feature allows the RM 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, and the HBA.

Monitoring User Logins

The AUDIT function continuously monitors logins to the RM controller and provides alerts in the event of unauthorized login attempts (Figure 3-49).

```
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 3-49 User Login Messages

The “USER AUDIT=ON|OFF” command can be used to enable/disable the user auditing function. When enabled, the system will display a message when a user logs in or out. Default is OFF.

The “USER CONNECTIONS” command will display a list of all the currently connected users and the host port that the user is connected to (Figure 3-50).

```
RM660 [1]: user connections
Current Connections:
User          P          External LUN, Internal LUN
-----
client1       1          000,000
Anonymous     2          000,000
Anonymous     3          000,000
```

Figure 3-50 User Connections Screen

The “USER SHOWALL” command will display the LUN mapping information for all the authorized users (Figure 3-51).

```
RM660 [1]: user showall
```

		LUN Zoning	
User		External LUN, Internal LUN	
000 client1		Using Host Port Zoning	
001 client2		000,010	001,011 002,012

Figure 3-51 Users' LUN Mapping Information Screen

To display the LUN mapping information for a particular user, use the "USER SHOW=x" command where x is the user ID number.

Zoning (Anonymous Access)

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

The "ZONING" command will display the current settings for the host ports. The LUN Zoning chart indicates which internal LUNs the users will have access to (with read-only and read/write privileges) and where the internal LUN will appear to the users. In Figure 3-52, only internal LUN 1 can be accessed and it is read-only for the users. It will appear as LUN 0 to the users.

```
RM660 [1]: zoning
```

		LUN Zoning	
Port	World Wide Name	External LUN, Internal LUN	
1	2100CB5CBA7F5F1F	R000,001	
2	2200CB5CBA7F5F1F	R000,001	
3	2300CB5CBA7F5F1F	R000,001	
4	2400CB5CBA7F5F1F	R000,001	

Figure 3-52 Current Zoning Configuration

The "ZONING EDIT" command 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). The default configuration is to deny access to all the LUNs (See "Host Port Zoning (Anonymous Access)" on page 43 for further information on how to change the settings). The "ZONING DEFAULT" command will restore the zoning of a host port back to its default settings.

```

RM660 [1]: zoning edit

                                LUN Zoning
Port   World Wide Name           External LUN, Internal LUN
-----
1      21000001FF000218
2      22000001FF000218
3      23000001FF000218
4      24000001FF000218

Enter the host port (1..4), 'e' to escape.
1
Enter the new LUN zoning for this host port.

For each external LUN specify the internal LUN the user will have access
to. 'R' preceding the LUN indicates the LUN will be read-only for the user.
Enter 'P' indicates a place holder LUN that will have zero capacity.
Enter 'n' for no LUN or 'e' to escape.

Specify the internal LUN (0..127), or 'n' for none or 'e' to escape.
External LUN 0 is not mapped. New internal LUN:

```

Figure 3-53 Edit Zoning Configuration Screen

User Authentication

This second level of protection provides superior securities for your SAN storage.

The RM controller creates correspondence between users (World Wide Name), 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 “its own” and “allowed to share” data according to its customized LUN identification scheme. Administrator 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 “Zoning (Anonymous Access)” on page 101 above.

The “USER” command will display the current settings for all authorized users (Figure 3-54). Each user is identified by its 64-bit World Wide Name and is given a unique user ID number. The Ports column indicates which host ports, on each RM controller, the user is allowed to log into. 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 will appear to the user.

```

RM660 [1]: user

User                World Wide Name      Ports          LUN Zoning
                   1      2          External LUN, Internal LUN
-----
000 client1         200000E08B00E867    1  4 1 3      R000,000 001,002 003,009
001 client2         210000E08B00C0D8    1234 1234     Using Host Port Zoning.
002 client3         210000E08B01703A    1234 1234     Using Host Port Zoning.

User auditing is enabled.

```

Figure 3-54 Current LUN Mapping Screen

To configure/change the settings, use these commands:

- USER ADD - Adds a new user and defines the right of access.
- USER EDIT - Edits the right of access of an existing user.
- USER DELETE - Deletes an existing user from the system.

See “User Authentication (Recommended for SAN Environment)” on page 40 for further information on how to add a new user.

Firmware Update Management

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

Displaying Current Firmware Version

The “VERSION” command displays version information of the RM controller’s hardware and firmware (Figure 3-55).

```
RM660 [1]: version

Silicon Graphics RM660
Firmware Version: 5.10

Firmware date: Sep 3 2004, 16:39:31
IEEE ULA Number: 00030478
Bootrom Version: 1.08
```

Figure 3-55 Version Information Screen

Firmware Update Procedure

The “TFTP” command enables the administrator to download the new RM controller firmware from a TFTP server to the RM controller. A TFTP server, such as the directMONITOR console, must be running when using this command. This command will “fail” the current RM controller and should not be used during active I/O.

Follow these steps if you need to update the firmware files.

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

3. 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 TIERDELAY
TIER         TIER CONFIG
```

4. Copy the new firmware file to your TFTP server (such as the directMONITOR console).
5. Connect to the RM controller via Telnet or serial (CLI port).
6. Enter command “TFTP”.
7. You will be asked to confirm action (Figure 3-56). Type “y” to continue.

```
RM660 [1]: tftp

WARNING: This requires failing and restarting
the current RM660.

Do you want to continue? (y/n): y

Enter the TFTP server IP Address: 010:123:139:005

Enter the filename: \(\path)\xxx.fsh
```

Figure 3-56 Downloading RM controller Firmware

1. Then enter the TFTP server's IP address.
2. Enter the firmware path and filename.
3. If you have the couplet RM controller configuration, connect and log into the other RM controller. Repeat Steps [6] to [2] above to update the firmware.
4. Issue the "RESTART" command on the RM controller unit(s) to restart.
5. For dual mode only:
After both RM controller are back on-line, use the "DUAL" command ("Couplet RM Controller Configuration (Cache/Non-Cache Coherent)" on page 79) to verify that both RM controller units are healthy. If either RM controller shows failed, login to the healthy RM controller and issue the "DUAL HEAL" command.

Remote Login Management

The "TELNET=ENABLE|DISABLE" command allows the administrator to *temporarily* enable and disable the establishment of a remote Telnet session. Use the "TELNET" command to display the current setting.

Note: The Telnet capability is reset to ON after RM controller restart. To turn off Telnet access permanently, use the "NETWORK" command (see "System Network Configuration" on page 66).

The "TELNET STATS" command allows the administrator to view various statistics maintained on remote Telnet sessions (Figure 3-57). These statistics are kept from the point of power-on.

```

RM660 [1]: telnet stats

Telnet Session Statistics (since Power-On):
=====

Telnet Power-On   :   Time: 21:43:53   Date: 03/27/2002
Current System   :   Time: 19:37:00   Date: 03/29/2002

Sessions have been:  -- ENABLED --
since           :   Time: 21:43:53   Date: 03/29/2002

Number of times Telnet Sessions have been:
    Enabled : 1
    Disabled: 0
    Killed  : 0

Total Telnet attempts           : 2
Total SUCCESSFUL Telnet Sessions : 2
Consecutive UNSuccessful Telnet attempts : 0
Total UNSuccessful Telnet attempts : 0

Breakdown of UNSuccessful Attempts:
-----
Telnet session was already active : 0
Telnet sessions were disabled    : 0
User supplied invalid login information : 0
Could not obtain the remote socket address : 0
Other (miscellaneous)            : 0

```

Figure 3-57 Telnet Statistics

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 (Figure 3-58).

```

RM660 [1]:
New TELNET Session initiated from IP address: 010:123:139:005
[Remote TELNET session ON] Local SUBshell RM660 [1]:

```

Figure 3-58 Telnet Session Initiated

Within the CLI sub-shell, the “TELNET” command allows the administrator to view information regarding the currently active Telnet session (Figure 3-59).

```
[Remote TELNET session ON] Local SUBshell RM660 [1]: telnet

Time: 19:49:48 Date: 03/29/2002

Remote Telnet Session Information:
-----
Owner's Name           : admin
Security Level         : Administrative
Remote Site IP Address : 010:123:139:005
Local SDD IP Address  : 10.123.134.1
Initiated at          : 19:49:06 on 03/29/2002
Duration              : 0 (seconds)
Most Recent Command
  User Entry          : whoami
  Initiated at       : 19:50:52 on 03/29/2002
  Completed at      : 19:50:52 on 03/29/2002
  Idle Time         : 0 (seconds)
```

Figure 3-59 Telnet Session Information

The “TELNET KILL” command lets the administrator terminate the remote Telnet session (Figure 3-60). The KILL parameter may also be specified with 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 if no value is specified. An administrative login is required before the command is processed.

```
[Remote TELNET session ON] Local SUBshell RM660 [1]: telnet kill=1

-- WARNING --

Any CLI command that may be in progress on the remote Telnet site
will need to be completed locally after the remote session has been
terminated.

Enter the administrative (or higher) login name: admin
Enter the appropriate password:

-- Please wait for the remote TELNET session to be terminated. --
.....
Telnet Session termination.
```

Figure 3-60 Terminating a Telnet Session

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 (Figure 3-61). 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.

```
RM660 [1]:  
  
-- The System Administrator will terminate this TELNET Session in 1 minute --  
-- The System Administrator will terminate this TELNET Session in 30 seconds --  
Connection closed by foreign host.
```

Figure 3-61 Telnet Session Being Terminated

Note that if the user is in the middle of running a CLI command at the remote Telnet site when the administrative KILL is effected, the command will continue as it was in progress, on the CLI console. It is up to the administrator to finish out such a command in a graceful fashion.

The System Logs

Message Log

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

To display the log of previous system messages, enter command “LOG”.

To clear the log of all previous messages, enter command “LOG CLEAR”.

The “LOG CHECKCONDITION|CHECKCONDITION=MORE” command will display the Check Condition log. The “MORE” option will display additional information concerning the check condition.

To clear the Check Condition log, enter command “LOG CHECKCLEAR”.

System and Drive Enclosure Faults

The “FAULTS” command will display a list of all current disk, system, and drive enclosure faults or failures (Figure 3-62).

```

RM660 [1]: faults

Disk 1C is failed.
RM660 : Power Supply 4 : FAILURE
SCSI Enclosure Services: Enabled: No faults were detected.

-----
                RM660 1                      RM660 2
-----
Label            RM660 Unit[1]                      RM660 Unit[2]
Status           Healthy                          Failed

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

No APC UPS faults detected via SNMP trap.

```

Figure 3-62 Current System Faults

Displaying System’s Uptime

The “UPTIME” command will display the system’s uptime which is the total time the system has been operational (also known as Power on Hours). 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 (Figure 3-63).

```

RM660 [1]: uptime

System Uptime (YY:DDD:HH:MM) : 00:102:19:38

```

Figure 3-63 Display System Uptime

The “UPTIME RESET” command will reset the system’s uptime to zero.

Saving a Comment to the Log

The “COMMENT <message>” command 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 on the command line.

Other Utilities

APC UPS SNMP Trap Monitor

The “APC_UPS” command will display the status of the APC UPS SNMP trap monitor (Figure 3-64).

```
RM660 [1]: apc_ups
APC UPS SNMP trap monitor is off.
No APC UPS faults detected via SNMP trap.
```

Figure 3-64 Displaying APC UPS SNMP Trap Monitor Status

To enable/disable this monitor, enter command:

```
APC_UPS MONITOR=ON|OFF
```

The “APC_UPS CLEAR_FAULTS” command 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 (Figure 3-65).

```
RM660 [1]: api
API Server connections are currently -- ENABLED--
```

Figure 3-65 Displaying Status of API Connections

To temporarily enable/disable the establishment of connections to the API server, issue command: `API=ENABLE|DISABLE`

When disabled, users at remote locations will be unable to establish a new API connection until an “API ENABLE” command is issued. Note that this command only provides control over API connections during the current power cycle. To “permanently” disable or enable API connections, across power-cycles, use the “NETWORK API_SERVER” command (see “API Server Connections” on page 69). Default setting for this command is enabled.

To display the collected statistics on API connections (Figure 3-66), enter command: `API STATS`

The command “API CLEARSTATS” will reset the collected statistics.

```

RM660 [1]: api stats

          API Server Connections Statistics
          =====
                Time           Date
          -----
System Boot-Up Completion      : 18:12:47   12/16/2003
System's Current                : 20:27:17   01/20/2004

API Server Initiation          : 18:22:46   12/16/2003
API Server Connections have been : -- ENABLED --
                                since : 18:12:46   12/16/2003

API Server services have been :
  Enabled : 1 time.
  Disabled : 0 times.

```

Figure 3-66 API Server Connection Statistics

Internal Mirrored Groups (IMG)

You may create, edit and destroy IMGs from the system, as well as breaking, merging, adding and deleting LUNs within the IMG.

To display the currently created IMGs (Figure 3-67), enter command “MIRROR”.

```

RM660 [1]: mirror

                Internal Mirrored Groups (IMG)
      Primary          Backup LUNS
IMG  LUN  1st  2nd  3rd  4th  5th  6th  7th  Owner Status
-----
-----No IMGs found-----
System Capacity 1044132 Mbytes, 3564 Mbytes available.

```

Figure 3-67 Current Internal Mirrored Groups

The “MIRROR CREATE” command will create a new IMG in the system. The system will prompt you for all the necessary information to create the IMG and indicate if the IMG was successfully added to the system.

The “MIRROR DESTROY” command will permanently delete an IMG from the system. This will delete only the IMG itself, but not the data on the individual LUNs. *Note that all Backup LUNs in the IMG must be broken before a delete can occur.*

The “MIRROR ADD” command can be used to add additional Backup LUNs to the IMG after it has been created. All data on the Backup LUN is destroyed and then it is data-synchronized with the Primary LUN.

The “MIRROR DEL|DELETE” command will delete a broken Backup LUN from the IMG. All data on the Backup LUN is preserved.

The “MIRROR BREAK” command can be used to break Backup LUNs from the IMG where the Backup LUN will no longer be mirroring the Primary LUN. All data on the Backup LUN is preserved. Parentheses surrounding a Backup LUN is an indication that the LUN is broken from the IMG.

The “MIRROR MERGE” command will add the Backup LUN back into the IMG allowing it to mirror the Primary LUN once again. All data from the Primary LUN is copied over onto the Backup LUN so they are again identical.

The “MIRROR PAUSE” command can be used to temporarily pause the merge procedure.

The “MIRROR RESUME” command will release the paused merge operations.

The “MIRROR STOP” command can be used to abort the merge procedure all together.

Changing Baud Rate for CLI Interface

To display the current serial console setting of the RM controller unit, enter command: `CONSOLE` (Figure 3-68)

```
RM660 [1]: console
Serial console baud rate is 115200 baud.
```

Figure 3-68 Displaying the Serial Console Setting

The “`CONSOLE BAUD`” command can be used to change the baud rate of the CONFIG port of the RM controller (Figure 3-69).

```
RM660 [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 3-69 Changing the Baud Rate

CLI/Telnet Session Control Settings

You may change the CLI's (and Telnet's) various session control settings. The “`SETTINGS`” command will display the current setting (Figure 3-70).

```
RM660 [1]: settings

Current Session Control Settings
-----
Lines per page:    0 (No paging - continuous scroll)
```

Figure 3-70 Current Session Control Settings

The “SETTINGS DEFAULTS” command will reset all the CLI and Telnet session control settings to their default values.

The “SETTINGS LINES=<number of lines>” command lets you display or set 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 entirely. 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.

Remote Management of the RM Controller

This section provides information on how to set up the RM controller for remote management and configuration.

The RM 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 RM controller also supports SNMP and allows the system to be remotely monitored.

Network Connection

Connect the Telnet port on the back of the RM controller to your Ethernet network (Figure 3-71). Then set the IP addresses, login names and passwords as described below.




Figure 3-71 Telnet Port on the RM660




Figure 3-72 Telnet Port on the RM610

Currently, the RM 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. The system needs to be restarted before the changes will take effect. However, you can make all the changes and then restart the system at the end.

The “network” command will display the current settings (Figure 3-73).

```

RM660 [1]: network

                                Network Configuration
                                =====
Gateway:                          172.16.0.254
Netmask:                           255.255.255.0
MAC Address HSTD #1:                00:01:ff:01:00:ae
IP Address HSTD #1:                 172.16.0.1
IP Address HSTD #2:                 -- Unknown --

Services
-----
Telnet:                             ENABLED
Telnet Port HSTD #1:                 23
Telnet Port HSTD #2:                 23

API Server:                         ENABLED
API Server Port HSTD #1:             8008
API Server Port HSTD #2:             8009

SNMP:                               ENABLED
SNMP Trap IP Address:               0.0.0.0
Limited SNMP:                       DISABLED

Syslog:                             ENABLED
Syslog IP Address:                  172.16.0.253
Syslog Port HSTD #1:                 514
Syslog Port HSTD #2:                 514

```

Figure 3-73 Current Network Configuration Screen

1. Change the RM controller's IP address for your network environment. Issue command:
`NETWORK IP=<new IP address>`
2. Change the netmask of the RM controller, if needed, using command:
`NETWORK NETMASK=<new netmask>`
3. Enable the Telnet capability, if needed. The command to use is:
`NETWORK TELNET=ON`

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

4. Decide whether the SNMP functionality should be enabled.
`NETWORK SNMP=ON|OFF`
If you are using directMONITOR (the external system console option) with the RM controller, the SNMP function should be enabled.
5. 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>`
6. Decide whether the Syslog capability should be enabled.
`NETWORK SYSLOG=ON|OFF`
If you are using directMONITOR (the external system console) with the RM controller, the syslog function should be enabled.
7. If the SYSLOG function is enabled, enter the destination IP address for the Syslog packets.
`NETWORK SYSLOGIP=<destination IP address>`
Make sure your destination computer supports the SYSLOG feature. For example, on UNIX systems, the SYSLOGD application must be properly installed and running.
8. The default destination port number for Syslog packets is 514, if you need to change it, enter command:
`NETWORK SYSLOGPORT=<port number>`
9. Set up the routing table which describes how the RM controller can communicate with the hosts on other networks. The "route" command will display the current settings (Figure 3-74).

- To set the current gateway in the network routing table to the specified Internet address:
ROUTE GATEWAY=<aaa.bbb.ccc.ddd>
- To delete gateways from the routing table:
ROUTE DEL=<aaa.bbb.ccc.ddd> GATEWAY=<aaa.bbb.ccc.ddd>
- Add new gateway to the table:
ROUTE ADD=<aaa.bbb.ccc.ddd> GATEWAY=<aaa.bbb.ccc.ddd>

```

RM660 8000[1]: route
Gateway NONE

Permanent RM660 Routing Table
destination      gateway
-----
0.0.0.0          191.9.161.222

Current RM660 Routing Table
ROUTE NET TABLE
destination      gateway  flags  Refcnt  Use  Interface
-----
0.0.0.0          191.9.161.222  3      0      0    fei0
160.100.0.0      191.9.161.222  3      0      3    fei0
191.9.0.0         191.9.160.208  101    0      0    fei0

ROUTE HOST TABLE
destination      gateway  flags  Refcnt  Use  Interface
-----
127.0.0.1        127.0.0.1  5      0      0    lo0

```

Figure 3-74 Current RM Controller Routing Table

10. If you have couplet RM controller, connect (or Telnet if this is not the initial set up) and log into the other RM controller. Then repeat Steps [1] to [9] above to set its network parameters.
11. Issue the “RESTART” command on the RM controller unit(s) to restart and make the changes take effect.
12. Upon boot up, verify the network connections using the “ping” command on the RM controller unit(s):
NETWORK PING=<IP address of a system on the network>

It is recommended to ping a host which is on the same subnet as the RM controller, then try another host which is on a different subnet.
13. If firewall is running (which will block traffic and hosts from talking to each other), make sure it is set up to allow the RM controller to pass information on Port 23 (for Telnet), Port 161 (for SNMP), and Port 162 (for SNMP traps).

Login Names and Passwords

The RM controller's two levels of security, administrative and general purpose user access, also applies to remote management. 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 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, via RS-232 or Telnet (see "Remote Login Management" on page 105).

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 command "PASSWORD DEFAULTS", while logged in as "admin", will restore all passwords and user names to the default values.

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 "Remote Login Management" on page 105 for more information).

The Telnet client should have local echoing disabled and use port 23.

Supporting the SGI InfiniteStorage RM610/RM660

Maintaining the SGI InfiniteStorage RM610 and RM660

This section provides information regarding error recovery on the SGI InfiniteStorage RM610 and RM660.

To display the list of all current system and drive enclosure faults, enter command:

```
FAULTS
```

Component Failure Recovery

The RM controller contains redundant and hot-swappable fans and power supplies. A single component failure, therefore, will not shut down the system. In the unlikely event of a component failure, you can replace the failed component while the RM controller is running. The replaced component will automatically be returned to service in the system.

Power Supply Failure

If an RS-232 serial console or a telnet session is being used, 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 red (Figure 4-1 and Figure 4-2).




Figure 4-1 Power Supply/Cooling Modules in the Front of the RM660




Figure 4-2 Power Supply/Cooling Modules in the Front of the RM610

Note: The RM controller should not be operated with any open slots for more than 15 minutes. Make sure you have the replacement power supply available before removing the failed module.

1. If installed, remove the cover panel by removing the two thumbscrews.
2. Locate the failed power supply module which is indicated by a red Status LED. Turn off its power switch.
3. Remove the two thumbscrews. Then slide the module out of the bay.
4. On the new module, check that the power switch is off.
5. Then slide the module into the bay. Make sure it is fully inserted. Install the two thumbscrews to secure it.
6. Turn on the power switch. Check that the Status LED is green, indicating that the module is operating normally.
7. Replace the cover panel if previously removed.

Fan Failure

If an RS-232 serial console or a telnet session is being used, a fan failure error message will be displayed on your console if the fan module fails. Follow these steps to replace the module.

Note: The RM controller should not be operated with any open slots for more than 15 minutes. Make sure you have the replacement fan unit available before removing the failed fan unit.

1. If installed, remove the cover panel by removing the two thumbscrews.
2. Remove the two thumbscrew from the fan module (Figure 4-3). Then slide the module out of the bay.




Figure 4-3 Fan Module

3. Slide the new module into the bay. Make sure it is fully inserted. Install the two thumbscrews to secure it.
4. Replace the cover panel if previously installed.

Recovering from Drive Failures

When a drive failure occurs on the RM controller, 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 in the same tier before 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 RM 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 a 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 loss of data. The logical unit(s) on that tier will continue to operate in degraded mode. If a spare drive is available, the RM controller will automatically rebuild the data on the spare drive if “autorebuild” is enabled. System operation is not affected while recovery is taking place.

If an RS-232 serial console or a telnet session is being used, when a drive failure occurs, the RM controller will display an error message on your console, indicating which drive has failed and which spare drive (if available) is replacing the failed drive. The Fault LED on the failed drive will turn amber. The error is also written to the event log. The RM 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. The “TIER” command will display the status of the disks. The failed drive is denoted by an “r” (Figure 4-4).

If a rebuild is taking place, the percentage of completion is also displayed. Figure 4-4 illustrates a system that has a failed drive on Tier 2, Channel D and it is being rebuilt. When rebuild is complete, a message “Finished Replacing disk 2D” will be displayed.

Tier Configuration					
Tier	Owner	Capacity (Mbytes)	Space Available (Mbytes)	Disk Status	Lun List
1	1	280012	0	ABCDEFGHPS	0 1
2	1	280012	0	ABC r EFHGHS	2 3

Rebuilding disk 2D, 14.6% completed.

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

System Capacity 560024 Mbytes, 0 Mbytes available.

Figure 4-4 Checking the Disk Status

The “DISK INFO” command will tell you which spare drive is replacing the failed drive (Figure 4-5).

<pre> RM660 [1]: disk info=2d Disk: 2D Status: Failed and replaced by spare 2S Capacity: 35002 (Mbytes) Block size 512 Write cache: Enabled AL_PA: E8 Vendor Id: SEAGATE Product ID: ST318203FC Product Rev: 0004 Serial Num: LR225022 Global Id: 2000002037294419 </pre>	<pre> RM660 [1]: disk info=2s Disk: 2S Status: Replacing failed drive 2D Capacity: 35002 (Mbytes) Block size 512 Write cache: Enabled AL_PA: E8 Vendor Id: SEAGATE Product ID: ST318203FC Product Rev: 0004 Serial Num: LR418906 Global Id: 2000002035292423 </pre>
--	--

Figure 4-5 Obtaining the Disk Information

Returning the system to fault-tolerant state:

It is recommended that you replace the failed drive as soon as possible so that the tier can return to its optimal state. If you replace the failed drive while rebuild is taking place, the system will finish rebuilding data on the spare drive first. When 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, you should replace the failed drive as soon as possible. Once the replacement drive is inserted and verified by the system, a rebuild will be automatically started if Autorebuild is ON. If not, you may initiate a rebuild by doing the following:

- Enter command “disk scan”
The system will check each channel and look for newly inserted drive(s).
- Initiate a rebuild using the command
“DISK REBUILD=<tier><channel>”.

Replacing a Failed Disk with Spare Disk Manually

You may manually replace a specified failed disk with a spare disk using the “DISK REPLACE=<tier><channel>” command. 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 “Resources Allocation” on page 96).

Changing the Rate of Rebuild

You may use the commands “TIER DELAY” and “TIER EXTENT” to control the percentage of processor time allocated to rebuild compare to I/O handling. Rebuild is done in steps. DELAY will control the amount of time to wait before rebuilding the next chunk of data. The size of the chunk of data to rebuild each time is controlled by the parameter EXTENT.

To increase the rate of rebuild, reduce the DELAY value (please refer to “Resources Allocation” on page 96 for 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

If you want to abort the rebuild, issue command “TIER STOP” which will stop all current rebuild operations.

Multiple Drive Failures 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 three 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, use the “TIER” command to check the current disk status (Figure 4-6) and see if the drives failed are all on the same channel.

Tier Configuration					
Tier	Owner	Capacity (Mbytes)	Space Available (Mbytes)	Disk Status	Lun List
1		280012	280012	ABC.EFGHPS	
2		280012	280012	ABC.EFGHPS	
3		280012	280012	ABC.EFGHPS	
Automatic disk rebuilding is Enabled					
System rebuild extent: 32 Mbytes					
System rebuild delay: 60					
System Capacity 840036 Mbytes, 840036 Mbytes available.					

Figure 4-6 Checking the Tier Status

Then use the “DISK” command to check the status of the channel. Figure 4-7 below shows that the system is looking for Channel D and there may be a problem with that channel.

```

Disk Channel Status
Disk Channel A healthy.
Disk Channel B healthy.
Disk Channel C healthy.
Disk Channel D acquiring loop synchronization.
Disk Channel E healthy.
Disk Channel F healthy.
Disk Channel G healthy.
Disk Channel H healthy.
Disk Channel P healthy.
Disk Channel S healthy.

All disks are healthy.

Disk write caching is Enabled.
Audio/Visual settings Disabled.
Disk command timeout: 10 seconds.
Disk command AV timeout: 0
Fast AV reads: Enabled
Ordered Tag Count: 0
Disk command outstanding: 0

```

Figure 4-7 Disk Channel Information

Recovery steps:

- You may get guidance from SGI Technical Support to determine the cause of error and steps to recover.
- Isolate the cause of channel failure. (e.g., loose cable connections)
- Rebuild all the drives on that channel individually by issuing command: DISK REBUILD=<tier><channel>
e.g. disk rebuild=1d
disk rebuild=2d
disk rebuild=3d

Several rebuild operations will be processed at the same time. The other rebuild jobs will be queued up and processed in the same order as they were entered.

Component Failure on Enclosures

The RM controller implements the SES (SCSI Enclosure Services) protocol for communications with the drive enclosures. If your enclosures provides SES communications, status information of the enclosures, including power supply, fan, and

presence of drive, will be obtained and evaluated. If a change in status is found, an “SES” message will be displayed on your console (Figure 4-8).

```
ses_fault 12:42:57 Disk missing 1B
ses_E 12:55:58 EncID:50050CC0000033BB: Device 3 SLOT ADDR D5: NOT INSTALLED
ses_A 13:45:23 EncID:50050CC0000033C8: Power Supply 1 :DC Power Failure
```

Figure 4-8 Example of SES Messages

If your enclosures provide redundant SES communication paths, the message will be reported twice. “EncID” is the World Wide Name 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.

To display the current drive enclosure failures, type in command “SES” (Figure 4-9).

```
RM660 [1]: ses
EncID:50050CC0000033C8: Power Supply 1 :DC Power Failure
```

Figure 4-9 Displaying the Current Disk Enclosure Failures

RM Controller Messages and Descriptions

This section provides descriptions of the RM610/RM660 messages.

The following describes the format and content of the system messages displayed by the RM controller. This will help explain the system messages (displayed on the RS-232 terminal, Telnet session, syslog messages, and SNMP traps) to aid administrators in diagnosing system events.

The messages generated by the RM controller will fall into the following categories: Fault messages, Warning messages, Timeout messages, and Informational messages.

Fault Messages

All system fault messages will begin with “Fault:” as shown below.

```
Fault: Disk read piece error!
```

These are the most serious messages displayed by the RM controller and should be dealt with immediately. These messages indicate that a hardware or software error has occurred which is unrecoverable and may cause the RM controller to become unstable. This means that the firmware has detected an error that prevents further operation. The unit may not function correctly after this type of error and a graceful shutdown of the system may not be possible. Steps should be taken to isolate the error that may allow the unit to continue functioning but continued operation is not recommended.

Actions to be taken in the field:

- Run diagnostics on the RM controller to insure it is still operational.

Warning Messages

All system warning messages will begin with “Warning:” as shown below.

```
Warning: SDRAM not detected on disk channel A, slot 0.
```

These messages indicate that the RM controller has encountered a serious error that may require user intervention. These messages may indicate a hardware error (such as a bus parity error), a failure of a hardware component (such as the cache memory on a disk channel), or an internal firmware error (such as a system variable going out of range).

Note: These errors do not include disk failures. Handling of disk failures is part of the normal operation of the RM controller. Warning messages only indicate a problem with the RM controller itself.

The difference between a “Fault” and a “Warning” is that the firmware will take steps to continue operation and insure the integrity of the data. However the firmware may be unable to continue normal operation and data may be lost due to the error. *These errors should be dealt with immediately and continued operation is not recommended.*

Actions to be taken in the field:

- Run diagnostics on the RM controller to insure it is still operational.
- Downgrade the firmware in the RM controller to a previous version known to be stable.

Timeout Messages

All system timeout messages will begin with “Timeout:” as shown in the following example.

```
Timeout: Read Disk 1A, Id:1
```

These messages indicate that the RM controller has timed out waiting for a data transfer such as a read from an initiator. This means that the unit may have trouble communicating with an initiator or a disk in the array. Timeout messages may also be generated under heavy I/O conditions or when an error occurs like a disk failure or a LIP on the Fibre Channel bus. Timeouts should generally not occur during normal operation and should be reported if they keep reappearing.

Informational Messages

All other system messages are used to inform the operator of a change in the system status. These messages do not indicate a problem with the RM controller but may indicate a problem that requires user intervention.

```
Failing Disk 1A  
Replacing Failed Disk 1A with Spare Disk 1S
```

All informational messages that are related to an I/O command from an initiator will be displayed with the initiator information. This will include the initiator’s 64-bit World Wide Name, the host port the initiator is communicating on, the source ID the initiator is using on the Fibre Channel bus, and the exchange ID of the SCSI command (see following example).

```
Read Address error LUN:0 address:80000000 length:1  
WWN:2000000011111111 port:4 S_ID:10 OX_ID:4189
```

```
SDRAM and Cache Error Messages
```

```
Warning: SDRAM not detected on disk channel A, slot 0.  
Warning: SDRAM not detected on disk channel A, slot 1.
```

The above messages indicate that the RM controller cannot locate the SDRAM on the disk channel. The SDRAM should be installed in slot 0 of each disk channel. If the SDRAM is installed in slot 1 on one channel, then all the other channels must have the SDRAM installed in slot 1 as well. This error will prevent the system from running normally. Make sure the SDRAM is installed correctly.

```
Warning: SDRAM configuration mismatch detected on channel A, slot 0.  
Warning: SDRAM configuration mismatch detected on channel A, slot 1.
```

The above messages indicate that there is a mismatch on the SDRAM configurations. The RM controller requires that all the SDRAM installed have the same size, configuration and speed. This error will prevent the system from running normally. Make sure the SDRAM is installed correctly.

```
Warning: Cache diagnostic failure on host port 1.  
Warning: Cache diagnostic failure on disk channel A.
```

The above messages indicate that the RM controller performed a diagnostic test on the SDRAM from the host and disk sides and discovered an error. This error will prevent the system from running normally. Make sure the SDRAM is installed correctly.

```
Warning: Unknown cache diagnostic status on host port 1.  
Warning: Unknown cache diagnostic status on disk channel A.
```

The above messages indicate that the RM controller performed a diagnostic test on the SDRAM from the host and disk sides and discovered an error with the test itself. This error will prevent the system from running normally. This may indicate a problem with the firmware.

```
Timeout: Cache diagnostic host port 1.  
Timeout: Cache diagnostic disk channel A.
```

The RM controller performed a diagnostic test on the SDRAM from the host and disk sides but the test took too long to complete. This error will prevent the system from running normally. This may indicate a problem with the hardware.

```
Warning: SDRAM Cache size is invalid.
```

The RM controller tried to initialize the cache descriptors and found that the size reported for the SDRAM was invalid. This error will prevent the system from running normally. Make sure the SDRAM is installed correctly.

```
Warning: Cache descriptor allocation is full.
```

The RM controller tried to initialize the cache descriptors and ran out of memory. The RM controller will continue to function normally but will be unable to use all of the SDRAM in the system for caching. This error should only occur if an extremely large amount of SDRAM cache is installed in the system and the cache cluster size is set very small. This can be corrected by increasing the cache cluster size of the unit.

```
Warning: Cluster size is invalid. Restoring saved value.
```

The RM controller tried to initialize the cache descriptors and found that the cache cluster size to be used was invalid. The RM controller will load the saved value from the parameter blocks and continue. It should continue to run normally but the cluster size may need to be adjusted. This may indicate a problem with the firmware.

```
Warning: Cluster size is invalid. Restoring default value.
```

The RM controller tried to initialize the cache descriptors and found that the cache cluster size to be used was invalid. The RM controller also found the saved value in the parameter block to be invalid as well. It will load the default value and continue. This unit should continue to run normally but the cluster size may need to be adjusted. This may indicate that the parameter block has been corrupted.

Disk Boot-up Error Messages

```
Warning: Unexpected Disk interrupt channel:A status: msg:  
Warning: Unexpected Host interrupt port:1 status: msg:
```

The RM controller received an interrupt from the host port or disk channel that it was not expecting. This may indicate a problem with the firmware or hardware.

```
Waiting for disks to Boot: 3FF
```

The RM controller sent a request to the disk channels to perform a LIP and to report the results back. This message indicates that the LIP request is taking a long time and the RM controller is still waiting. The LIP results will be delayed if a disk is in the process of spinning up. It is normal for the disks to take a while to boot when the unit is first powered on. The hex number at the end of the message is a bitmap that indicates which disk channels the unit is still waiting for. This message is informational only.

```
Warning: Disk boot incomplete: 3FF
```

The RM controller sent a request to the disk channels to perform a LIP and to report the results back. This message indicates that the LIP request took too long and at least one disk channel failed to report in. The hex number at the end of the message is a bitmap

that indicates which disk channels failed to report in. This error will prevent the system from running normally. This may indicate a problem with the hardware or the disks in the array.

```
Warning: Disk channel A boot failure.
```

The above message indicates that a disk channel failed to boot properly. This error will prevent the system from running normally. This may indicate a problem with the hardware.

```
Warning: Disk 1A requested AL_PA EF, received 01
```

The above message indicates that the disk did not receive the hard assigned AL_PA it had requested on the Fibre Channel bus. This indicates a problem with the configuration of the disks in the array. Check to make sure the disks are all assigned unique addresses, power cycle the disks and reboot the RM controller.

```
Disk 1A returned invalid inquiry data.
```

The above message indicates that the RM controller received an error when asking for the SCSI inquiry data from the disk. The disk may have returned an error for the inquiry command, the inquiry data may be incorrect or the inquiry data may indicate that the device is not a disk. Make sure the disk is installed correctly.

```
Disk 1A returned invalid capacity data.
```

The above message indicates that the RM controller received an error when asking for the SCSI read capacity data from the disk. The disk may have returned an error for the read capacity command. Make sure the disk is installed correctly.

```
Disk 1A returned an invalid blocksize of 400 bytes.
```

The above message indicates that the RM controller received an error when asking for the SCSI read capacity data from the disk. The disk returned a blocksize that is not supported by the unit. Check the disk to make sure it is correct.

```
Disk 1A returned an invalid capacity of 100 blocks.
```

The above message indicates that the RM controller received an error when asking for the SCSI read capacity data from the disk. The disk returned a capacity that is too small to be supported by the unit. Check the disk to make sure it is correct.

```
Could not find Disk 1A
```

The above message indicates that the RM controller could not locate a disk that is part of a valid LUN in the system. The data in the LUN will not be available until the RM

controller can locate the disk or the disk is marked as failed. Make sure the disk is installed correctly.

```
Warning: Could not login to disk 1A. Device unknown.
```

The above message indicates that the RM controller found a device on the Fibre Channel bus but could not login with it. Make sure the disk is installed correctly and cycle power on the disk.

```
Disk 1A Capacity of %ld is less than the minimum required of %ld.
```

The above message indicates that the capacity of the disk is not large enough to support the existing LUN's that use the disk. Check the disk to make sure it is correct.

```
Warning: Disk 1A reports it should be disk 2A.
```

The above message indicates that a disk that is part of a valid LUN in the system is not in the correct location. It may be on the wrong channel or may have a different AL_PA. The data in the LUN's will not be available until the disk is restored to the correct location or the disk is marked as failed. Make sure the disk is installed correctly.

```
Warning: Disk 1A does not match the installed disk.
```

The above message indicates that the RM controller does not recognize a disk that is part of a valid LUN in the system. It may be in the wrong location or it may have been replaced without informing the RM controller. The data in the LUN's will not be available until the original disk is restored or the disk is marked as failed. Make sure the disk is installed correctly.

```
Warning: Internal parameter block 0 is invalid. Defaults loaded.
```

The above message indicates that the internal parameter block was invalid so the RM controller loaded a default copy. This is a normal response when a RM controller is powered on for the first time. It will attempt to get a valid copy of the parameter block from the disks. Saving the parameter blocks from the terminal with the SAVE command will fix this error. If this error occurs after the system has been configured properly then this may indicate that the parameter block has been corrupted, the hardware is faulty, or there is a problem with the firmware.

```
Found older parameter block 0 on disk 1A.  
Old Date: TUE MAY 02 11:18:58 2000  
New Date: WED MAY 03 11:18:58 2000
```

The above message indicates that all the disk copies of this parameter block were older than the internal copy of the parameter block. This usually occurs when the system is

shut down before the RM controller can save an updated parameter block to the disks. In this case the RM controller will correct the problem the next time the parameters are saved. Saving the parameter blocks from the terminal with the SAVE command will fix this error. This error may also occur if a new RM controller with a valid configuration is attached to disks with a different configuration. In this case, disconnect the new RM controller from the disks and use the DEFAULTS command to reset its settings and then reconnect it to the disks. The new RM controller will then load the correct parameter blocks from the disks.

```
Loaded parameter block 0 from disk 1A.
Old Date: TUE MAY 02 11:18:58 2000
New Date: WED MAY 03 11:18:58 2000
```

The above message indicates that the internal copy of the parameter block did not match any of the disk copies of the parameter block. The RM controller will load the latest copy of the parameter blocks from the disks. In this case the RM controller will correct the problem the next time the parameters are saved. Saving the parameter blocks from the terminal with the SAVE command will also fix this error. This error may also occur if a new RM controller is attached to disks with a different configuration.

```
Fault: Parameter block error: %s is %ld should be %ld
```

The above message indicates that the RM controller detected a firmware incompatibility error. This indicates a problem with the firmware. The RM controller may not function correctly after this type of error. *Do not use the RM controller in this state or the system may become permanently corrupted.* Downgrade the firmware in the RM controller to a previous version known to be stable.

```
Warning: Error saving parameters, Invalid block 0
```

The above message indicates that the RM controller has detected an internal firmware error. The RM controller may not function correctly after this type of error and data may have been lost. This indicates a problem with the firmware. Downgrade the firmware in the RM controller to a previous version known to be stable.

```
Disk bad current cache mode page data: 1A
Disk bad changeable cache mode page data: 1A
Disk bad new cache mode page data: 1A
Disk bad current error recovery mode page data: 1A
Disk bad changeable error recovery mode page data: 1A
Disk bad new error recovery mode page data: 1A
Disk bad current disconnect mode page data: 1A
Disk bad disconnect mode page data: 1A
Disk bad new disconnect mode page data: 1A
```

The above messages indicate that the RM controller received an error when asking for the SCSI mode page data from the disk. The disk may have returned an error for the mode sense command, the data returned may be incorrect, or the disk may not support the mode page requested. The mode sense data is not required for the normal operation of the RM controller. Make sure the disk is installed correctly.

Disk Error Messages

```
Failing Disk 1A
```

The above message indicates that the RM controller failed the disk. The disk may have encountered an error, returned bad status to a SCSI command, failed to respond, or took too long to process a command. Check the disk and replace it if it is no longer functional.

```
Replacing Failed Disk 1A with Spare Disk 1S
```

The above message indicates that the RM controller failed the disk and replaced it with an available spare. The disk may have encountered an error, returned bad status to a SCSI command, failed to respond, or took too long to process a command. Check the disk and replace it if it is no longer functional.

```
No Spares available.
```

The above message indicates that the RM controller failed a disk in the array but could not find a spare disk to replace it with. This message is informational only. The tier containing the failed disk will begin operating in degraded mode. Check the disk and replace it if it is no longer functional.

```
Spare too small %d, %x  
LUN 1 won't fit on spare 1S, %x
```

The above message indicates that the RM controller failed a disk in the array and replaced it with a spare disk but discovered that the spare disk is not large enough to support the LUN's. The hex value at the end of the message indicates the capacity of the disk. This message is informational only. The tier containing the failed disk will begin operating in degraded mode. Check the disk and replace it if it is no longer functional.

```
Restoring Disk 1A
```

The above message indicates that the RM controller is restoring a failed disk by rebuilding the parity information from the other disks in the array. This message is informational only. The RM controller will indicate when the restore operation has completed.

Healing Disk 1A

The above message indicates that the RM controller is healing a failed disk without rebuilding the parity information from the other disks in the array. The operator must request this action. This message is informational only. The RM controller will indicate when operation has completed.

Failing Disk channel A

The above message indicates that the RM controller has failed an entire disk channel due to a hardware error. The RM controller will continue to operate and the data will be available unless other disks or channels have failed in the system. Reboot the RM controller and run diagnostics.

Healing Disk channel A

The above message indicates that the RM controller has restored a disk channel to healthy status after it was failed due to a hardware error. The operator must request this action. This message is informational only.

```
Timeout: Read Disk 1A, Id:EF
Timeout: Write Disk 1A, Id:EF
Timeout: Command 2A Disk 1A, Id:EF
Timeout: Manual command 2A Disk 1A
```

The above messages indicate that a SCSI command sent to the disk has taken too long to complete. The timeout messages may be generated under heavy I/O conditions when an error occurs, such as a disk failure, or a LIP on the Fibre Channel bus. This could also indicate that the RM controller might have trouble communicating with the disk. Timeouts should generally not occur during normal operation and should be reported if they keep reappearing. Check the disk and replace it if it is no longer functional.

SCSI Check Condition Disk 1A

The above message indicates that the RM controller received a check condition status on a SCSI command it sent to a disk. The RM controller will display any valid sense data returned with the status. Check the disk and replace it if it is no longer functional.

SCSI Reservation Conflict Disk 1A

The above message indicates that the RM controller received a reservation conflict status on a SCSI command it sent to a disk. The disks should not be reserved in the normal operation of the RM controller. The RM controller will display any valid sense data returned with the status.

SCSI Command Terminated Disk 1A

The above message indicates that the RM controller received a command terminated status on a SCSI command it sent to a disk. The RM controller will display any valid sense data returned with the status. Check the disk and replace it if it is no longer functional.

Disk missing 1A

The above message indicates that the RM controller sent a SCSI command to a disk that it could no longer locate on the Fibre Channel bus. Check the disk and replace it if it is no longer functional.

Disk not responding 1A

The above message indicates that the RM controller sent a SCSI command to a disk that is no longer responding to commands on the Fibre Channel bus. Check the disk and replace it if it is no longer functional.

Disk ALPA changed 1A

The above message indicates that the AL_PA which the disk was using on the Fibre Channel bus has unexpectedly changed. This indicates a problem with the configuration of the disks in the array. Check to make sure the disks are all assigned unique addresses, cycle power on the disks and reboot the RM controller.

Disk Login failed 1A

The above message indicates that the RM controller found a device on the Fibre Channel bus but could not login with it. Make sure the disk is installed correctly and cycle power on the disk.

Disk Fibre channel loop failure 1A

The above message indicates that the RM controller found a device on the Fibre Channel bus but could not login with it. Make sure the disk is installed correctly and cycle power on the disk.

Disk Command Timeout 1A

The above message indicates that the SCSI command sent to the disk was aborted because the command took too long to complete. Check the disk and replace it if it is no longer functional.

Disk Command Data Under-run 1A

The above message indicates that a read or write SCSI command completed with a data under-run status. This indicates a problem with the firmware or the configuration of the disks in the array.

```
Disk Command aborted 1A ID EF
```

The above message indicates that the disk aborted a SCSI command. Check the disk and replace it if it is no longer functional.

```
Warning: SDRAM parity error detected 1A ID EF, p:%x
```

The above message indicates that the RM controller detected a parity error when reading the data on a disk channel. This indicates a problem with the hardware. Make sure the SDRAM is installed correctly and run the cache diagnostics on the RM controller.

```
Warning: SDRAM parity error recovered 1A ID EF, p:%x
```

The above message indicates that the RM controller detected a parity error when reading the data on a disk channel but was able to recover after a retry. This indicates a problem with the hardware. Make sure the SDRAM is installed correctly and run the cache diagnostics on the RM controller.

```
Warning: Galeforce parity error detected 1A ID EF, p:%x
```

The above message indicates that the RM controller detected a parity error inside a disk Fibre Channel chip when reading the data on a disk channel. This indicates a problem with the hardware. Run the cache diagnostics on the RM controller.

```
Warning: Galeforce parity error recovered 1A ID EF, p:%x
```

The above message indicates that the RM controller detected a parity error inside a disk Fibre Channel chip when reading the data on a disk channel but was able to recover after a retry. This indicates a problem with the hardware. Run the cache diagnostics on the RM controller.

```
Warning: Bus parity error detected 1A ID EF
```

The above message indicates that the RM controller detected a parity error inside a disk channel when reading the data on a disk channel. This indicates a problem with the hardware. Run the cache diagnostics on the RM controller.

```
Warning: Bus parity error recovered 1A ID EF
```

The above message indicates that the RM controller detected a parity error inside a disk channel when reading the data on a disk channel but was able to recover after a retry.

This indicates a problem with the hardware. Run the cache diagnostics on the RM controller.

```
SCSI Error Status: %x Disk 1A ID EF
```

The above message indicates that the RM controller received a SCSI error status on a command it sent to a disk. The RM controller will display any valid sense data returned with the status. Check the disk and replace it if it is no longer functional.

```
Data recovered disk:1A  
Data rebuild recovered disk:1A
```

The above message indicates that the RM controller successfully rebuilt a section of missing data on a disk. This is usually due to an error encountered while reading from a disk drive such as a media error or a recovered error. This message is informational only. This could indicate a problem with the disk.

```
Warning: Retry failure detected! disk:1A
```

The above message indicates that the RM controller was not able to rebuild a section of missing data on a disk after several retries. This is usually due to an error encountered while reading from a disk drive, such as a media error or a recovered error, which does not go away. This indicates a problem with the disk.

```
Warning: Data loss detected! disk 1A  
Warning: Data rebuild loss detected! disk:1A
```

The above message indicates that the RM controller was saving the data to the disks when it encountered an error on more than one disk and the retries were unsuccessful. This indicates a problem with the disks.

```
Fault: Invalid manual command semaphore 2A disk 1A  
Fault: Disk Invalid operation code:  
Fault: Disk read piece error!  
Fault: Disk normal piece error!  
Fault: Disk write invalid error:  
Fault: DMT DISKQ != NULL  
Warning: Unknown manual command 2A disk 1A  
Warning: Invalid Manual command buffer %02x disk 1A  
Warning: Disk Fail error t:1, c:A  
Warning: Disk channel fail error %d  
Warning: Disk channel heal error %d  
Warning: Disk manual command invalid c:%c %x:%x  
Warning: Disk manual command length error c:%c L:%x RC:%x  
Warning: Disk command invalid:1A rsp:%d cmd:%d, status:%x  
Warning: Disk command previously finished 1A rsp:%d cmd:%d,
```

```
status:%x
Warning: Disk timeout check overflow!
Warning: DMT disk reuse error!
Warning: WBK count reset:
```

The above messages indicate that the RM controller has detected an internal firmware error. The RM controller may not function correctly after this type of error and data may have been lost. This indicates a problem with the firmware. Downgrade the firmware in the RM controller to a previous version known to be stable.

Host Error Messages

```
Warning: Host Array parity error.
```

The above message indicates that the RM controller detected an array parity error when a host Fibre Channel port was reading the data from the SDRAM. This indicates that the parity data may not be consistent or the data may be corrupted. This may have been caused by an error with the disks.

```
Warning: Host Bus parity error: %03x
Warning: Host %d, Bus and array parity error: %03x
```

The above messages indicate that the RM controller detected a bus parity error when a host Fibre Channel port was reading the data from the SDRAM. The hex number at the end of the message is a bitmap that indicates which disk channels the error was detected on. This indicates a problem with the hardware. Run the cache diagnostics on the RM controller.

```
Warning: Host Internal error.
```

The above message indicates that the RM controller detected an internal error in a host Fibre Channel port when transferring data. This indicates a problem with the hardware or the firmware. Run the cache diagnostics on the RM controller.

```
Timeout: Host SCSI command:%02X LUN:1 Lane:%d
```

The above message indicates that a SCSI command from an initiator has taken too long to complete. Timeout messages may be generated under heavy I/O conditions or when an error occurs, such as a disk failure. This could also indicate that the RM controller might have trouble communicating with the initiator. Timeouts should generally not occur during normal operation and should be reported if they keep reappearing. Check the connection to the initiator on the host port.

```
Command %02X aborted LUN:1 T:%x
```

The above message indicates that a SCSI command from an initiator was aborted. The abort may have come from the initiator, a target reset, or from a timeout in the RM controller. This could also indicate that the RM controller might have trouble communicating with the initiator. Aborts should generally not occur often during normal operation and should be reported if they keep reappearing. Check the connection to the initiator on the host port.

```
Target Reset from host 1
```

The above message indicates that a SCSI target reset message was received on a host port. The target reset came from an initiator connected to the RM controller. This could also indicate that the RM controller might have trouble communicating with the initiator. Target resets should generally not occur often during normal operation and should be reported if they keep reappearing. Check the connection to the initiator on the host port.

```
Read Address error LUN:1 address:%x length:%x
Write Address error LUN:1 address:%x length:%x
Write Flush Address error LUN:1 address:%x length:%x
```

The above messages indicate that an initiator sent a SCSI command to the RM controller for an address beyond the capacity of the LUN. This means that the initiator is sending an illegal address request to the RM controller. The RM controller will report the error to the initiator without transferring any data. This often occurs if a LUN is deleted and added back with different capacity without informing the initiator. This message is intended to be informational only.

```
Write Under Run LUN:1 length:%x DL:%x
Read Under Run LUN:1 length:%x DL:%x
```

The above messages indicate that an initiator sent a SCSI read or write command to the RM controller with a length in the CDB which is less than the length in the DL field of the command frame. This means that the initiator is not building the command frames correctly. The RM controller will transfer the data using the length specified in the CDB and report the error to the initiator. This often occurs if a LUN is deleted and added back with a different blocksize without informing the initiator. This message is intended to be informational only.

```
Read Over Run LUN:1 length:%x DL:%x
Write Over Run LUN:1 length:%x DL:%x
```

The above messages indicate that an initiator sent a SCSI read or write command to the RM controller with a length in the CDB which is greater than the length in the DL field of the command frame. This indicates that the initiator is not building the command frames correctly. The RM controller will transfer the data using the length specified in the

DL field of the command frame and report the error to the initiator. This often occurs if a LUN is deleted and added back with a different blocksize without informing the initiator. This message is intended to be informational only.

```
Fault: Host invalid operation code: %x
Fault: Load lock error n:%x m:%x
Warning: Host node finished read error L:%d N:%x RC:%x RD:%x P:%x
Warning: Host node finished write error L:%d N:%x WC:%x WD:%x
```

The above messages indicate that the RM controller has detected an internal firmware error. The RM controller may not function correctly after this type of error and data may have been lost. This indicates a problem with the firmware. Downgrade the firmware in the RM controller to a previous version known to be stable.

Rebuild Error Messages

```
Diagnostics Failed disk 1A
```

The above message indicates that the RM controller was unable to rebuild a failed disk because the disk did not pass the diagnostic tests. Check the disk and replace it if it is no longer functional.

```
Timeout: Host 1 Rebuild LUN:0 Address:%x
```

The above message indicates that a rebuild request in the RM controller has taken too long to complete. Timeout messages may be generated under heavy I/O conditions or when an error occurs, such as a disk failure. This could also indicate that the RM controller is having hardware problems. Timeouts should generally not occur during normal operation and should be reported if they keep reappearing.

```
Warning: Spare tier 1 configuration error: %d, %d
```

The above message indicates that the RM controller detected a discrepancy in the configuration information for the spare disk. The RM controller may remove the spare from use if it cannot determine the spare disk's proper configuration. This may indicate a problem with the configuration of the system or a firmware error.

```
Rebuild not needed on LUN: 0
```

The above message indicates that the RM controller tried to rebuild data on a LUN that no longer exists. This usually occurs when a LUN is deleted while the system is rebuilding a disk. This message is intended to be informational only.

```
Rebuild aborted on LUN: 0
Format aborted on LUN: 0
```

The above messages indicate that the RM controller aborted a rebuild or format operation on a LUN. This usually occurs when the user cancels a format or rebuild operation. These messages are intended to be informational only.

```
Warning: Host 1 rebuild bus error: %x
```

The above message indicates that the RM controller detected a bus parity error when a host Fibre Channel port was reading the data from the SDRAM. The hex number at the end of the message is a bitmap that indicates which disk channel the error was detected on. This indicates a problem with the hardware. Run the cache diagnostics on the RM controller.

```
Warning: Disk rebuild error %d,%d
Warning: Rebuild Disk invalid: %d,%d %d
```

The above messages indicate that the RM controller has detected an internal firmware error. The RM controller may not function correctly after this type of error and data may have been lost. This indicates a problem with the firmware. Downgrade the firmware in the RM controller to a previous version known to be stable.

SES (SCSI Enclosure Services) Messages

```
Unsupported SES device configuration detected on channel A.
```

The above message indicates that the RM controller does not recognize the drive enclosure that is connected to Channel x (i.e. the enclosure is not supported). The RM controller, therefore, will not light the Fault LED of a failed drive in that enclosure, but probably will still be able to monitor the common drive enclosure elements such as power supply and cooling fans.

Full SES monitoring and visual disk fault indication are supported on the SGI drive enclosures.

```
ses_C 15:06:43 EncID:50050CC0000033C8: Power Supply 1: DC Power
Failure.
```

The above message indicates that there is a change in status on the drive enclosure where "EncID" (Enclosure ID) is the disk enclosure's WWN. Each element (such as power supply, cooling fans, and temperature sensors) is checked for status changes. If an

element's status has changed since it was previously checked (OK -> failure, failure -> OK), a message is printed to the console.

Other Messages

```
Fault: taskLock %x
Fault: taskUnlock %x
Fault: semGive %x
Fault: semTake %x
Fault: semDelay %x
Fault: taskSpawn %x,%x
Fault: task ended
```

These messages mean that the firmware has detected a problem with the operating system of the RM controller that prevents further operation. The RM controller may not function correctly after this type of error and a graceful shutdown of the system may not be possible. *Continued operation is not recommended.* This may indicate a problem with the firmware. Run diagnostics on the RM controller to insure it is still operational.

Drive Enclosure for SGI InfiniteStorage RM610/RM660

Introduction to the SA2016 System

Figure 5-1 shows the SA2016 system.



Figure 5-1 The SA2016 System

The SA2016 System design concept is based on a subsystem together with a set of plug-in modules. The SA2016 subsystem as supplied comprises:

- Chassis and Backplane with integral Operators Panel. (See Figure 5-7 on page 155)
- Up to 16 Serial ATA (SATA) Drive Carrier modules (See Figure 5-9 on page 157)
 - Serial ATA (SATA) drives with appropriate switching card.
- Dummy drive carrier modules.
- Two AC, 450W Power Supply/Cooling plug-in modules (see Figure 5-1)
- One Serial ATA Control (LRC) Input/Output modules, dependent on configuration required: 1.5Gb internal operating speed with 1 Gb or 2Gb external operating speed. (See Figure 5-6).
- Blank (LRC) I/O module

Note: The SGI **Loop Resiliency Circuit (LRC) I/O Module** is called the **SCM Module** by DDN.

Enclosure Chassis

The chassis consists of a sheet metal enclosure assembly containing a Backplane PCB and module runner system. This chassis assembly also includes an integral Operators (Ops) Panel, mounted at the rear.

The chassis assembly contains 16 drive bays at the front, each of which accommodates a plug-in drive carrier module. The 16 drive bays are arranged in 4 rows of 4 drives. At the rear, the chassis assembly contains the integral ops panel and four module bays to house two Power Supply/Cooling modules, one LRC I/O module and a Blank (LRC) I/O module.

The Backplane PCB provides logic level signal and low voltage power distribution paths. [Figure 1-2](#) and [Figure 1-3](#) show front and rear views of an RS-1602 chassis respectively.

The chassis is fitted with 19 inch Rack mounting features which enables it to be fitted to standard 19 inch racks and uses 3EIA units of rack space.

- A Bay is defined as the space required to house a single 1.0" high 3.5 inch disk drive in its carrier module. e.g. a 1 x 4 bay module would take the space of 1 drive width by 4 drive bays high (in rack mount configuration).
- A 4 x 4 Chassis fitted with 19 inch Rack mounting features enables it to be fitted to standard 19 inch racks. It uses 3EIA units of rack space




Figure 5-2 Enclosure Chassis (Front)




Figure 5-3 Enclosure Chassis (Rear)

The Plug-in Modules

An SA2016 System Enclosure requires one or more of the following modules for normal operation:

Power Supply/Cooling Module

Two auto ranging AC 450W Power Supply/Cooling modules (see Figure 5-4) are supplied mounted in the rear of the enclosure as part of the subsystem core product.




Figure 5-4 Power Supply Cooling Module

PSU voltage operating ranges are nominally 115V or 230V AC, selected automatically.

Four LEDs mounted on the front panel of the Power Supply/Cooling module (see Figure 2-1) indicate the status of the PSU and the fans.

Multiple Power Supply/Cooling Modules

The SA2016 System must always be operated with two Power Supply/Cooling modules fitted. The two Power Supply/Cooling modules operate together so that if one fails the other maintains the power supply and cooling while you replace the faulty unit.

Module replacement should only take a few minutes to perform but must be completed within 10 minutes from removal of the failed module.

Operators Panel

Supplied as an integral part of the Enclosure core product, a typical Operators (Ops) Panel is shown in Figure 5-5.

The Ops Panel provides the enclosure with a micro controller which is used to monitor and control all elements of the Enclosure.




Figure 5-5 Ops Panel

Ops Panel Indicators and Switches

The Ops Panel includes Light Emitting Diodes (LEDs) which show the status for all modules, an Audible Alarm which indicates when a fault state is present, a push-button Alarm Mute Switch and a thumb wheel Enclosure ID Switch.

The Ops Panel switch functions are shown in Table 5-1.

Note: Table 5-1 shows default settings for SA2016 System LRC usage at 1Gb/s

Table 5-1 Ops Panel Switch Functions

Switch Number *See Sw 11	Function	Recommended Setting		Definition
1	Not Used	Off		
2	Not Used	Off		
3	Not Used	Off		
4	Not Used	Off		
5	Not Used	Off		
6	Not Used	Off		
7 & 8	Drive Loop Speed Select	Sw 7	Sw 8	
S2A Controller Settings		Off	Off	Force 1Gb/s
		On	Off	Force 2Gbs (not supported)
9 & 10	Drive Addressing Mode Selection	Sw 9	Sw 10	
S2A Controller Settings		Off	On	Mode 1
		On	On	Mode 0
		On	Off	Mode 2
		Off	Off	Mode 3 (<i>not supported</i>)*
11	SOFT SELECT	On		Selects functions using the hardware switches
12	Not Used	Off		

Note: * Please contact your supplier for further information.

Important: Switch settings are only read at Power On.

LRC Input/Output Module

The SA2016 System storage subsystem includes an enclosure with rear facing bays which houses one or two SATA Control Interface Adaptor (LRC I/O) modules (see Figure 5-3 on page 147), dependent on configuration required.

The plug-in LRC I/O modules have been designed for integration into a SA2016 System storage subsystem, providing external FCAL cable interfacing with up to 16 SATA disk drives.

Processors housed on the LRC I/O modules provide enclosure management and interface to devices on the Backplane, PSU, LRC and Ops Panel, to monitor internal functions.

The module incorporates the following LED indicators:

Important: If only one LRC module is fitted, the LRC module must be installed in Rear Bay 3 location (see Figure 5-3 on page 147) and a Blank I/O module fitted in the unused bay.

Table 5-2 LRC I/O Module LEDs

LED	Definition	Color	Normal Status	Fault Status
FC Host Port 0 Signal Good	Incoming FC signal is GOOD No connection or incorrect connection Invalid SFP connection	Green	On	Off Flashing
FC Host Port 1 Signal Good	Incoming FC signal is GOOD No connection or incorrect connection Invalid SFP connection	Green	On	Off Flashing
Router Status	Storage Router Device Ready Storage Router Device not ready or defective	Green	On	Off
ESI/LRC Module Fault	Fault present (also <i>On when booting</i>) Successful controller initialization	Amber	Off	On




Figure 5-6 LRC I/O Module




Figure 5-7 LRC Front Panel (as viewed from rear of Enclosure)




Figure 5-8 SA2016 System Enclosure with LRC I/O Modules Installed

Drive Carrier Module

The Drive Carrier Module comprises a hard disk mounted in a carrier. Each drive bay can house a single Low Profile 1.0 inch high, 3.5 inch form factor disk drive in its carrier. The carrier has mounting locations for ATA or FC-AL drives.

Each disk drive is enclosed in a die-cast aluminum carrier which provides excellent thermal conduction, radio frequency, and electro-magnetic induction protection and affords the drive maximum physical protection.

The front cap also supports an ergonomic handle which provides the following functions:

- Camming of carrier into and out of drive bays.
- Positive 'spring loading' of the drive/backplane connector.
- An anti-tamper lock operated by a torx socket type key.




Figure 5-9 Drive Carrier Module

SATA Dual Path Transition Card

For SATA drives, a SATA dual path transition card is used. This is attached to the rear of each drive to provide a SCA-2 interface to the drive carrier using the same pins as Fibre Channel. In addition, the card provides a SATA path switch (1.5Gb/s) suitable for dual redundant I/O module applications.

Drive Status Indicators

Each drive carrier incorporates two indicators, an upper (Green) and lower (Amber). In normal operation the green indicator will be ON and will flicker as the drive operates. The amber indicator is OFF during normal operation and ON if there is a drive fault present.

Anti-tamper Locks

Anti-tamper locks are fitted in the drive carrier handles (see Figure 5-10) and are accessed through the small cutout in the latch section of the handle. These are provided to disable the normal 'pinch' latch action of the carrier handle and so prevent accidental or unauthorized removal of drives.




Figure 5-10 Anti-tamper Lock

Dummy Carrier Modules

Dummy carrier modules are provided for fitting in all unused drive bays. They are designed as integral drive module front caps with handles and must be fitted to all unused drive bays to maintain a balanced airflow.

Blank Modules

Blank LRC I/O modules must be fitted in the vacant LRC bay (slot 3) at the rear of the enclosure (see Figure 5-3 on page 147) to maintain airflow and ensure correct operation.

Warning: Operation of the Enclosure with *ANY* modules missing will disrupt the airflow and the drives will not receive sufficient cooling. It is *ESSENTIAL* that all apertures are filled before operating the unit. Dummy Carriers and/or Blank modules are available for this purpose.

Visible and Audible Alarms

The functional modules have associated status LEDs. The Ops Panel shows a consolidated status for all modules.

LEDs show constant green for good or positive indication. Constant Amber LEDs indicate there is a fault present within that module.

The Ops Panel also incorporates an Audible Alarm to indicate when a fault state is present and also an Alarm Mute push-button.

Warning: The Ops Panel is an integral part of the enclosure chassis assembly and is not field replaceable.

SA2016 System Technical Specification

Dimensions

Rack Enclosure	inches	millimeters
Height	5.13	130
Width across mounting flange	19	483
Width across body of enclosure	17.6	447
Depth from flange to rear of enclosure body	21	532
Depth from flange to maximum extremity of enclosure (rear hold down)	21.7	551
Depth from flange to furthest extremity at front of unit	0.67	17

Weight

Maximum Configuration	Rack mount:	37kg (81lb)
Empty Enclosure (Rack)		9kg (19.8lb)
PSU/Cooling Module		4kg (8.8lb)
LRC I/O Module		0.9kg (1.98lb)

AC Power (450W PSU)

Voltage Range	100-120 / 200-240 VAC Rated
Voltage Range Selection	Automatic
Frequency	50/60 Hz
Inrush Current	50A @ 260VAC
Power Factor	>0.98
Harmonics	Meets EN61000-3-2

PSU 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

Environment

Table 5-3 Ambient Temperature and Humidity

	Temperature Range	Relative Humidity	Max. Wet Bulb
Operational	5°C to 40°C	20% to 80% non-condensing	23°C
Non-Operational	0°C to +50°C	8% to 80% non-condensing	27°C
Storage	1°C to +60°C	8% to 80% non-condensing	29°C
Shipping	-40°C to +60°C	5% to 100% non-precipitating	29°C

Airflow

System must be operated with low pressure rear exhaust installation

(Back pressure created by rack doors and obstacles not to exceed 5 pascals [0.5mm Water gauge])

Altitude, Operational

0 to 2133 m (0 to 7,000ft)
(10,000ft at maximum ambient of 35°C)

Altitude, Non-Operational

-305 to 12,192m (-1000 to 40,000ft)

Shock, Operational

Vertical axis 5g peak 1/2 sine, 10ms

Shock, Non-Operational

30g 10ms 1/2 sine

Vibration, Operational

0.21grms 5-500 Hz Random

Vibration, Non-Operational

1.04grms 2-200 Hz Random

Vibration, Relocation

0.3g 2-200 Hz sine

Acoustics	Sound Pressure Operating - Less than 58 dB LpA average measured at the bystander positions. (The 4 bystander positions are 1m horizontal and 1.5m off the floor positioned front, back, left and right. The unit under test will be measured on the floor) Measured at 20°C
Orientation & Mounting	19" Rack mount (3EIA Units)
• Rack Rails	To fit 800mm depth Racks compliant with IEC 297
• Rack Characteristics	Back pressure not exceeding 5 pascals (0.5mm water gauge)
Safety & Approvals	CE, UL, cUL
• EMC	EN55022 (CISPR - A), FCC A

Interfaces

Drive support	See drive carrier specification
Attachment	1 FCAL loop of 16 drives Passive Backplane with 1 or 2 Loop Resiliency Circuit (LRC) I/O Modules. Host Port: FC-AL SFP - SFP cables Expansion Port: SFP to SFP cables
	Maximum external cable length: see manufacturer's specification

LRC I/O Module Specification

Speed	1.5Gb/s internal to each drive, 1Gb/s or 2Gb/s external <ul style="list-style-type: none">• Creates connections to a single loop of 16 drives• 1 External FC loop with two SFP connections
Mounting	Rear, single bays 3 and/or 4 (<i>see <Reference>Figure 5-3</i>)
Connectors	<ul style="list-style-type: none">• 2 x Expansion Ports: SFP connector• 2 x Host Ports: SFP connector (2 per port)

LED Functions

FC Host Port 0 Signal Good	<ul style="list-style-type: none"> – Green: On: Incoming FC signal is GOOD – Off: no connection or incorrect connection – Flashing: Invalid SFP connection
FC Host Port 1 Signal Good	<ul style="list-style-type: none"> – Green: On: Incoming FC signal is GOOD – Off: no connection or incorrect connection – Flashing: Invalid SFP connection
Router Status	<ul style="list-style-type: none"> – Green: On: Storage Router Device Ready – Off: Storage Router Device not ready or defective
ESI/LRC Module Fault	<ul style="list-style-type: none"> – Amber: On: Fault present (also <i>On when booting</i>) – Off: successful controller initialization

Power Dissipation 20W max.

Drive Carrier Module Specification

Please contact your supplier for details of approved drives.

Important: Operating the SA2016 System subsystem with non-approved drives may invalidate the warranty.

Module Dimensions	Height 29.1mm Width 106.55mm Depth 44 mm
Weight	0.98kg (1.0" 36Gb drive)
Transition card	mounting locations for ATA - SATA drives with transition card attached
Operating Temperature	5° C to 40° C (when installed in an RS-1602 system enclosure with dual Power Supply/Cooling Modules)
Power Dissipation	17 Watts maximum

Software Enclosure Services (SES) Support

The enclosure has a sophisticated self-monitoring and reporting function which conforms to ANSI SES specifications. This reports on such topics as:

- Enclosure temperature
- Fan speed

- Drive condition
- Operator panel status

Installation of the SA2016 Drive Enclosure

In this chapter, you are shown how to install your SA2016 Enclosure and plug-in modules into an industry standard 19 inch rack cabinet.

Caution: When connecting up the SA2016 subsystem, use only the power cords supplied or cords which match the specification quoted in “Power Cord” on page 160.

Note: The SGI **Loop Resiliency Circuit (LRC) I/O Module** is called the **SCM Module** by DDN.

Planning Your Installation

Before you begin installation you should become familiar with the configuration requirements of your SA2016 system, detailed in Table 6-1. The correct positions of each of the optional plug-in modules are shown in “Planning Your Installation” on page 167. Please refer to sections “LRC I/O Module Configurations” on page 172 and “LRC I/O

Module Installation” on page 190 for details of LRC I/O module configurations and installation.

Table 6-1 SA2016 Configuration

Module	Location
Drive Bays	ALL drive bays must be fitted with either a drive carrier module or a dummy carrier, no bays should be left completely empty. Drive carrier modules 0 & 15 provide SES Management Services.
Power Supply/Cooling Modules	Two Power Supply/Cooling modules must be fitted. Full power and cooling redundancy is provided while a faulty module is replaced. Install the Power Supply/Cooling modules in rear Bays 1 and 5. Note: Rear bays are numbered from 1 to 5 commencing from the right side.
LRC I/O Module	One loop resiliency circuit (LRC) I/O module should be installed in rear Bay 3. A blank module is fitted rear Bay 4.
Blank LRC I/O Modules	Install in rear Bay 4.
Ops Panel	(integral part of chassis assembly). Installed in rear Bay 2.

Note: The SGI **Loop Resiliency Circuit (LRC) I/O Module** is called the **SCM Module** by DDN.

Caution: Dummy Carriers and Blank Modules **MUST** be fitted to ALL unused bays. There is inadequate drive cooling if any are left open.




Figure 6-1 Module locations

Enclosure Bay Numbering Convention

The enclosure bay numbering convention is shown in Figure 6-1. A **bay** is defined as the space required to house a single 1.0 inch high 3.5 inch disk drive in its carrier module. For example, a 1 x 4 bay module would take the space of 1 drive width by 4 drive bays high (in the rack mount configuration).

The SA2016 subsystem is housed in a 4 x 4 enclosure, that is, 4 bays wide by 4 bays high.

- The front bays are numbered 1 to 4 from left to right, as viewed from the front. Bays are numbered from 1 (top row) to 4 (bottom row). Drive Carrier Module locations are identified from a matrix of the top and side numbers.
- The rear bays are numbered 1 to 5 from right to left, as viewed from the rear.

Enclosure Installation Procedures

Caution: The SA2016 Enclosure with all its component parts installed is too heavy for a single person to easily install into a Rack cabinet. The following procedures describe the installation of the SA2016 enclosure and highlights any critical co-requisite requirements and good handling practices which you should follow so as to ensure that a successful installation is achieved in the easiest manner.



Warning: Ensure that you have fitted and checked a suitable anti-static wrist or ankle strap and observe all conventional ESD precautions when handling SA2016 modules and components. Avoid contact with Backplane components and module connectors, etc.

Pre-Requisites

The SA2016 enclosure is designed for installation into an industry standard 19 inch cabinet capable of holding the unit.

- Minimum depth: 500 mm from front flange to rear metalwork (excludes rear cabling).
- Weight: up to 37kg dependent upon configuration per enclosure.
- Clearances: a minimum gap of 25mm (1inch) clearance between the rack cover and front of drawer; and 50mm (2 inches) rear clearance between rear of drawer and rear of rack is recommended in order to maintain the correct air flow around the enclosure.
- Maximum back pressure: The rack should present a maximum back pressure of 5 pascals (0.5mm water gauge).

Rack Mounting Rail Kit

A set of mounting rails is available for use in 19 inch rack cabinets. These rails have been designed and tested to handle the maximum enclosure weight and to ensure that

multiple SA2016 enclosures may be installed without loss of space within the rack. Use of other mounting hardware may cause some loss of rack space.

Parts Check List

- Rack Mounting Rail Kit

Chassis Installation

Parts Check List

- Chassis (complete with Backplane and Ops Panel installed but excluding all plug-in modules).
- Rack mount front flange mounting screws (4 off).

Procedure for Chassis Installation

Procedure 6-1 Procedure for Chassis Installation

1. Check for damage.
2. Slide the chassis assembly onto the rack rails until the front flanges engage on the rack.
3. Ensure the chassis is centrally located. If in doubt about correct orientation, the drive bays (at front) should have their black drive connectors toward the bottom of each bay.
4. Screw the 4 front rack mount screws through the flanges and tighten.
5. Fit and tighten the rear hold down screws ensuring the enclosure is in tight contact to both the side and top of the chassis to avoid any movement of the chassis in the rack.

Power Supply/Cooling Module Installation

Two power supply/cooling modules to be installed in the rear of the enclosure in positions 1 and 5.



Warning: Do not remove covers from the power supply unit. Danger of electric shock inside. Return the PSU to Customer Service for repair.

Parts Check List

- 2 x AC, 450W Power Supply/Cooling Modules

Procedure for Power Supply/Cooling Installation

Procedure 6-2 Power Supply/Cooling Module Procedure

Important: PSU0 (RH rear bay) must be fitted “upside-down” as shown in Figure 6-1 on page 167.

1. Check for damage, especially to the rear connector on the supply.

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 PSU handle in the open position (Figure 6-2), slide the module into the enclosure (Figure 6-3).

Important: Install the Power Supply/Cooling module (PSU 0) in the bay on the right side (Rear Bay 1) of the enclosure in an “upside down” orientation.

3. Cam the module home by manually closing the PSU handle (Figure 6-4). You should hear a click as the handle latches engage.
4. Connect the power supply cord to the power source and switch the power supply ON.




Figure 6-2 AC Power Supply/Cooling Module - Handle in Open Position




Figure 6-3 Installing an AC Power Supply Cooling Module (1)




Figure 6-4 Installing an AC Power Supply Cooling Module (2)

LRC I/O Module Configurations

Important: Please refer to “Drive Carrier Configuration” on page 194, for information on SATA drive configurations.

Internal Loop Structures

The SA2016 enclosure is configured with one internal loop of 16 drives.

FC-AL Interface

The LRC (I/O) interface module provides an external FC-AL interface via SFP connection. The LRC module facilitates the emulation of fibre channel drives by ATA/SATA drives by converting FC protocols to ATA/SATA protocols. This is done

through protocol conversion on the LRC module, which directly connects to each drive using the serial ATA protocol and then presents to the host via one FC loop hubbed to two or four SFP connections, depending on configuration.

Note: There are no external terminators required with Fibre Channel architecture and any drive may be hot plugged during operation.

Connecting Multiple Enclosures

Multiple enclosures are connected to the RAID Head/Host by using SFP to SFP (Small Form Factor) cables. The configuration options supported are shown in Table 6-2

Table 6-2 Configuration Options

Configuration	Figure
SA2016 Enclosures Connection to RM610 Unit #1	Figure 6-5 on page 176
SA2016 Enclosures Connection RM610 Unit #2	Figure 6-6 on page 177
SA2016 Enclosures Connected to RM660 Unit #1 (1)	Figure 6-7 on page 180
SA2016 Enclosures Connected to RM660 Unit #1 (2)	Figure 6-8 on page 181
SA2016 Enclosures Connected to Couplet RM660 Unit #2 (1)	Figure 6-9 on page 182
SA2016 Enclosures Connected to Couplet RM660 Unit #2 (2)	Figure 6-10 on page 183
Daisy-Chaining the SA2016 Enclosures to RM610 Unit#1	Figure 6-11 on page 185
Daisy-Chaining the SA2016 Enclosures to Couplet RM610 Unit #2	Figure 6-12 on page 186
Daisy-Chaining the SA2016 Enclosures to RM660 Unit #1	Figure 6-13 on page 188
Daisy-Chaining the SA2016 Enclosures to Couplet RM660 Unit #2	Figure 6-14 on page 189

Connection SA2016 Enclosures on the RM610

The basic configuration of the ten SA2016 enclosures. Each enclosure will connect to one channel on the RM610. You may daisy-chain up to seven SA2016 enclosures to each channel.

There is one LRC I/O module in each SA2016 enclosure. Each module connects to one of the two internal drive loops. In couplet RM610 configurations, connections to the other internal drive loop will provide redundant data paths.

The following steps explain how to connect 10 enclosures to the RM660. Enclosure can hold up to 16 disk drives, so the configuration can house 16 full tiers (160 drives)

1. For Channel A: Connect a SFP cable between the "Disk A" port on the back of the RM610 and the upper right "In-Arrow-Direction" SFF connector on Enclosure #1 (See Figure 6-5 on page 176).
2. For Channel B: Connect a SFP cable between the "DISK B" port on the back of the RM610 and the upper right "In-Arrow-Direction" SFF connector on Enclosure #2 (See Figure 6-5).
3. Similarly, connect the other enclosures to Channels CDPS on the RM610 (See Figure 6-5).

RM610 Channels	TO	SA2016 Enclosures

A	TO	1
B	TO	2
C	TO	3
D	TO	4
P	TO	5
S	TO	6

4. Using the ID Range switch on the Ops panel modules, select ID "1" for all ten enclosures.
5. For couplet configurations the RM610 disk channel connections for Unit #2 should be attached between the "Disk Channel" port on the back of the RM610 and the lower right "In-Arrow-Direction" SFF connector of the each enclosure (See Figure 6-6 on page 177).




Figure 6-5 SA2016 Enclosures Connection to RM610 Unit #1




Figure 6-6 SA2016 Enclosures Connection RM610 Unit #2

Connection SA2016 Enclosures on the RM660

The basic configuration of the ten SA2016 enclosures. Each enclosure will connect to one channel on the RM660. You may daisy-chain up to seven SA2016 enclosures to each channel.

There is one LRC I/O module in each SA2016 enclosure. Each module connects to one of the two internal drive loops. In couplet RM660 configurations, connections to the other internal drive loop will provide redundant data paths.

The following steps explain how to connect 10 enclosures to the RM660. Enclosure can hold up to 16 disk drives, so the configuration can house 16 full tiers (160 drives)

1. For Channel A: Connect a SFP cable between the "Disk A" port on the back of the RM660 and the upper right "In-Arrow-Direction" SFF connector on Enclosure #1 (See Figure 6-7 on page 180).
2. For Channel B: Connect a SFP cable between the "DISK B" port on the back of the RM660 and the upper right "In-Arrow-Direction" SFF connector on Enclosure #2 (See Figure 6-7).
3. Similarly, connect the other enclosures to Channels CDEFGHPS on the RM660 (See Figure 6-7 and Figure 6-8 on page 181).

RM660 Channels	TO	SA2016 Enclosures
-----------------------	-----------	--------------------------

A	TO	1
B	TO	2
C	TO	3
D	TO	4
E	TO	5
F	TO	6
G	TO	7
H	TO	8
P	TO	9
S	TO	10

4. Using the ID Range switch on the Ops panel modules, select ID "1" for all ten enclosures.
5. For couplet configurations the RM660 disk channel connections for Unit #2 should be attached between the "Disk Channel" port on the back of the RM660 and the lower right "In-Arrow-Direction" SFF connector of the each enclosure (See Figure 6-9 on page 182 and Figure 6-10 on page 183).




Figure 6-7 SA2016 Enclosures Connected to RM660 Unit #1 (1)




Figure 6-8 SA2016 Enclosures Connected to RM660 Unit #1 (2)




Figure 6-9 SA2016 Enclosures Connected to Couplet RM660 Unit #2 (1)




Figure 6-10 SA2016 Enclosures Connected to Couplet RM660 Unit #2 (2)

Daisy-Chaining the SA2016 Enclosures on the RM610

The following steps explain how to daisy-chain the SA2016 enclosures. You may daisy-chain up to seven enclosures to each channel, giving a total of 112 tiers.

1. For Channel A: Connect a SFP cable between the "Disk A" port on the back of the RM660 and the upper right "In-Arrow-Direction" SFF connector on Enclosure #1 (See Figure 6-11 on page 185).
2. For Channel: Connect a SFP cable between the "DISK B" port on the back of the RM660 and the upper right "In-Arrow-Direction" SFF connector on Enclosure #2 (See Figure 6-11).
3. To daisy-chain the second set of enclosures (See Figure 6-11):
 - For Channel A: Connect a daisy-chain cable between the upper right "Out-Arrow-Direction" SFF connector on Enclosure #1 and the upper right "In-Arrow-Direction" SFF connector on Enclosure #7.
 - For Channel A: Connect a daisy-chain cable between the upper right "Out-Arrow-Direction" SFF connector on Enclosure #2 and the upper right "In-Arrow-Direction" SFF connector on Enclosure #8.
 - Similarly, connect the other expansion enclosures to Channels CDPS.
4. Repeat Step [2] above to connect additional sets expansion enclosures.
5. Using the ID Range switch on the Ops panels, select an ID number for each enclosure. Use ID "1" on the first set of enclosures that are directly connected to the RM610s. Then use ID "2" for the second set of enclosures in the chain, ID "3" for the third set of enclosures, and so on.
6. For couplet configurations the RM610 disk channel connections for Unit #2 should be attached between the "Disk Channel" port on the back of the RM610 and the lower right "In-Arrow-Direction" SFF connector of the each enclosure, with daisy-chain connections attached from the lower right Out-Arrow-Direction SFF connector to the lower right "In-Arrow-Direction" SFF connector of the next enclosure (See Figure 6-12 on page 186).




Figure 6-11 Daisy-Chaining the SA2016 Enclosures to RM610 Unit#1




Figure 6-12 Daisy-Chaining the SA2016 Enclosures to Couplet RM610 Unit #2

Daisy-Chaining the SA2016 Enclosures on the RM660

The following steps explain how to daisy-chain the SA2016 enclosures (see Figure 6-13 and Figure 6-14). You may daisy-chain up to seven enclosures to each channel, giving a total of 112 tiers.

1. 1. For Channel A: Connect a SFP cable between the “Disk A” port on the back of the RM660 and the upper right “In-Arrow-Direction” SFF connector on Enclosure #1 (See Figure 6-13 on page 188).
2. For Channel B: Connect a SFP cable between the “DISK B” port on the back of the RM660 and the upper right “In-Arrow-Direction” SFF connector on Enclosure #2 (See Figure 6-13).
3. To daisy-chain the second set of enclosures (See Figure 6-13):
 - For Channel A: Connect a daisy-chain cable between the upper right “Out-Arrow-Direction” SFF connector on Enclosure #1 and the upper right “In-Arrow-Direction” SFF connector on Enclosure #11.
 - For Channel A: Connect a daisy-chain cable between the upper right “Out-Arrow-Direction” SFF connector on Enclosure #2 and the upper right “In-Arrow-Direction” SFF connector on Enclosure #12.
 - Similarly, connect the other expansion enclosures to Channels CDEFGHPS.
4. Repeat Step [2] above to connect additional sets expansion enclosures.
5. Using the ID Range switch on the Ops panels, select an ID number for each enclosure. Use ID “1” on the first set of enclosures that are directly connected to the RM610s. Then use ID '2" for the second set of enclosures in the chain, ID “3” for the third set of enclosures, and so on.
6. For couplet configurations the RM660 disk channel connections for Unit #2 should be attached between the “Disk Channel” port on the back of the RM610 and the lower right “In-Arrow-Direction” SFF connector of the each enclosure, with daisy-chain connections attached from the lower right Out-Arrow-Direction SFF connector to the lower right “In-Arrow-Direction” SFF connector of the next enclosure (See Figure 6-14 on page 189).




Figure 6-13 Daisy-Chaining the SA2016 Enclosures to RM660 Unit #1




Figure 6-14 Daisy-Chaining the SA2016 Enclosures to Couplet RM660 Unit #2

LRC I/O Module Installation

Important: The LRC I/O module must be installed in Rear Bay 3 location (see Figure 6-1 on page 167) and a Blank I/O module fitted in the unused bay.

Parts Check List

- 1 LRC I/O Modules
- 1 Blank (I/O) module *(if required)*

Procedure for LRC I/O Module Installation

Check for damage especially to the interface connector. Do not install if any pins are bent.




Figure 6-15 LRC Module Latch Operation

1. Install the module in rear Bay 3 of the enclosure (Figure 6-1).
2. With the latch in the open position (see Figure 6-15), slide the LRC module into the enclosure until the latch engages automatically.
3. Cam the module home by manually closing the latches (Figure 6-16). You should hear a click as the latch engages.




Figure 6-16 Installing an LRC I/O Module in Bay 4

Drive Enclosure Device Addressing

Each enclosure has 16 drive bays. The SEL_ID of each drive is determined by the device slot (0-15) in which it is installed and the address range setting, which is set by means of the Enclosure ID switch on the Ops Panel (shown in Figure 6-5 on page 176) at the rear of the enclosure. The switch settings are shown in Table 6-3.

Note: Table 6-3 shows default settings for SA2016 LRC usage at 1Gb/s

Table 6-3 Ops Panel Switch Functions

Switch Number	Function	Recommended Setting		Definition
1 thru 6	Not Used	Off		
7 & 8	Drive Loop Speed Select	Sw 7	Sw 8	
S2A Controllers		Off	Off	Force 1Gb/s
9 & 10	Drive Addressing Mode Selection	Sw 9	Sw 10	
S2A Controllers		Off	On	Mode 1
11	SOFT SELECT	On		Selects functions using the hardware switches
12	Not Used	Off		

Note: ON = switch to the left, OFF = switch to the right.

Table 6-4 Mode 1 Drive Addressing

Thumb Wheel Switch	Device Slot SEL_ID Mode 1															
0	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
1	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
2	32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47
3	48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63

Table 6-4 Mode 1 Drive Addressing

Thumb Wheel Switch	Device Slot SEL_ID Mode 1
4	64 65 66 67 68 69 70 71 72 73 74 75 76 77 78 79
5	80 81 82 83 84 85 86 87 88 89 90 91 92 93 94 95
6	96 97 98 99 100 101 102 103 104 105 106 107 108 109 110 111
7	96 97 98 99 100 101 102 103 104 105 106 107 108 109 110 111
8	96 97 98 99 100 101 102 103 104 105 106 107 108 109 110 111
9	96 97 98 99 100 101 102 103 104 105 106 107 108 109 110 111
10	96 97 98 99 100 101 102 103 104 105 106 107 108 109 110 111
11	96 97 98 99 100 101 102 103 104 105 106 107 108 109 110 111
12	96 97 98 99 100 101 102 103 104 105 106 107 108 109 110 111
13	96 97 98 99 100 101 102 103 104 105 106 107 108 109 110 111
14	96 97 98 99 100 101 102 103 104 105 106 107 108 109 110 111
15	96 97 98 99 100 101 102 103 104 105 106 107 108 109 110 111

Table 6-5 Drive Slot Arrangement: Enclosure Front View

Row/Column	1/#	2/#	3/#	4/#
#/1	Drive 0*	Drive 1	Drive 2	Drive 3
#/2	Drive 4	Drive 5	Drive 6	Drive 7
#/3	Drive 8	Drive 9	Drive 10	Drive 11
#/4	Drive 12	Drive 13	Drive 14	Drive 15*

Note: 1. Drives are numbered row/column. 2. With only one active PSU the enclosure takes approximately 96 seconds to start all drives from Power On. 3 * Denotes SES drives which should always be fitted.

Drive Carrier Configuration

Planning and Configuring Your Installation

System Configuration

Important: Before you begin installation you should become familiar with the configuration requirements of your SA2016 system. Please refer to Section <Reference> for information on your overall system configurations.

There must be a drive present in Bay 1/1 (drive 0) or 4/4 (drive 15) to enable SES Communications to operate. Installing drives in both of these bays provides redundant SES communication paths.

When planning your system configuration, please remember that:

- All SA2016 enclosure drive bays must be filled with either a drive carrier or front dummy fascia. No bays should be left completely empty.

Drive Configuration

Important: After you have installed the drive carrier modules in your SA2016 enclosure, please refer to “LRC I/O Module Configurations” on page 172 for configuration information relevant to the I/O module you are installing.

Drive Carrier Installation

Parts Check List

- Drive Carrier module, or
- Dummy Carrier module

Procedure for Drive Carrier Installation

1. Ensure that the anti-tamper lock is disengaged.
2. Release the carrier handle by pressing the latch in the handle towards the right.
3. Insert the carrier into the enclosure (see Figure 6-17).

Important: For a Rack Mounted System: Ensure that the carrier is oriented so that the drive is uppermost and the handle opens from the left. Slide the carrier, gently, all the way into the enclosure until it is stopped by the camming lever on the right of the carrier (see Figure 6-18).

4. Cam the carrier home - the camming foot on the base of the carrier engages into a slot in the enclosure. Continue to push firmly until the handle fully engages. You should hear a click as the latch engages and holds the handle closed (see Figure 6-19).

Note: Ensure that the Handle always opens from the left.




Figure 6-17 Installing a SATA Drive Carrier Module (1)




Figure 6-18 Installing a SATA Drive Carrier Module (2)




Figure 6-19 Installing an SATA Drive Carrier Module (3)

Note: Removal is the reverse of this procedure (press on the latch to release the handle).

Dummy Carrier Modules

Any unused drive bays must be fitted with a dummy carrier module.

Engaging the Anti-tamper Locks

The anti-tamper locks are fitted in the drive carrier handles and are accessed through the small cutout in the latch section of the handle.

Drives are supplied with the locks set in the locked position.

Activating the Locks

1. Carefully insert the lock key provided into the cutout in the handle.
2. Locate the key into its socket.
3. Rotate the key in a clockwise direction until the indicator is visible in the aperture beside the key.
4. Remove the key.




Figure 6-20 Activating the Anti-tamper Lock

Deactivating the Locks

De-activation is the reverse of this procedure. To deactivate, rotate the key in an anti-clockwise direction until the indicator is no longer visible in the aperture beside the key.

Note: A drive carrier cannot be installed if its anti-tamper lock is activated outside the enclosure.

Power Cord Connection

Parts Check List

- Power cord to requisite local standards

Procedure

1. Attach the power cord to the Power Supply/Cooling Modules.
2. Attach the power cord to the in-line IEC connector in this cord.
3. Switch on each Power Supply/Cooling Module. The PSU Good and AC Fail LEDs on the PSU indicate whether AC mains power is present.

Caution: The power connections must always be disconnected prior to removal of the Power Supply/Cooling module from the enclosure.

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.

Operation of SA2016 Enclosure Plug-in Modules

Before powering up the enclosure please ensure that all the modules are firmly seated in their correct bays.

Power On

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.

Note: Please refer to Figure 7-1 on page 205, for details of the Ops Panel LEDs and related fault conditions.

Procedure 7-1 Power on Drive Enclosure

Follow the procedure below to Power On the enclosure.

1. Apply AC Mains power to the enclosure. Turn the Power Supply modules to ON.
2. On the Ops Panel, the Audible Alarm beeps once, all LEDs flash for 7 seconds then the Alarm double beeps.
3. All LEDs on the Ops Panel should be lit (Green) when the enclosure power is activated (and the disk drive motors should start).

Note: All LEDs on the Ops Panel should be lit Green at power up to indicate that the system is functioning correctly. If any show Amber then a problem exists. Follow the appropriate procedure in Chapter 8, "Troubleshooting and Problem Solving".

Important: If mains power is lost for any reason, on restoration of power the enclosure will re-start automatically.

Power Supply/Cooling Module LEDs

The Power Supply/Cooling module incorporates 4 LEDs, located below the On/Off switch and shown in Table 7-1.

- 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 7-1 PSU LEDs

AC PSU	
 Power AC Fan Power Good Fail Fault Fault	
PSU Good	Green
AC input Fail	Amber
Fan Fault	Amber
DC Output Fail	Amber

Ops Panel LEDs

The Ops Panel LEDs fault and status conditions are defined in Table 3-2 and shown in Figure 7-1.




Figure 7-1 Ops Panel LEDs and Switches

Please refer to Chapter 8, “Troubleshooting and Problem Solving” for details of any fault indication.

Table 7-2 Ops Panel LED States

LED	Definition	Color	Normal Status	Fault Status
Invalid Address	Indicates that an invalid Enclosure ID has been selected or that the selection has changed after Power On	Amber	Off	Flashing
Power On	Enclosure Powered On	Green	On	Off
System Fault	System/SCM Fault	Amber	Off	On
PSU/Cooling Fault	PSU Cooling fault or enclosure over-temperature.	Amber	Off	On
2Gb Link Speed	Indicates link speed	Green	Off	On
Hub Mode	Not Used	Green	Off	Off

Starting the Drives

Unless otherwise selected during installation, all drives in the enclosure should automatically start their motors. If this has not occurred one of the following conditions may exist:

- There may be a power problem (an alarm and power fault indication would normally be active).
- If there is only one Power Supply/Cooling Module present, the drive motors will spin up in a delayed sequence.

Disk Drives LEDs

Each drive carrier incorporates two indicators, an upper (GREEN) and lower (AMBER). In normal operation the Green LED will be ON and will flicker as the drive operates. The amber indicator is OFF during normal operation and ON if there is a drive fault present.

Power Down an Enclosure

Procedure 7-2 Power Down an Enclosure

To power the Enclosure down, perform step 1 or step 2, as follows;

1. Switch off the Power Supply/Cooling modules installed in the Enclosure.
2. Remove AC Mains at the power source

Troubleshooting and Problem Solving

The SA2016 Enclosure includes a processor and associated monitoring and control logic to enable it to diagnose problems within the enclosure's power, cooling and drive systems.

The Enclosure Services Processor is housed along with the Ops Panel in the rear of the enclosure.

The sensors for power and cooling conditions are housed within the Power Supply/Cooling modules. There is independent monitoring for each unit.

If a fault is indicated on the Ops Panel, please refer to Table 7-2 on page 206

Emulation Limitations

Each LRC presents each SATA disk drive as a single ported FC device. However unlike real FC-AL device certain data such as emulated WWN is not available on Power On, so the emulated devices will not present on the FC loop and participate in FC loop initialization until the drive has spun up. Please note that your RAID Head/Host system needs to be aware of this extra spin up time during Power On situations, whether they be planned ON or NOT ON (e.g. as the result of an unexpected power interrupt).

Note: The SGI **Loop Resiliency Circuit (LRC) I/O Module** is called the **SCM Module** by DDN.

Initial Start-up Problems

Faulty Cords

First check that you have wired up the subsystem correctly. Then, if:

- cords are missing or damaged
- plugs are incorrect
- cords are too short

Call your supplier for a replacement.

Alarm Sounds On Power Up

Please refer to “Audible Alarm” on page 214.

Green “Signal Good” LED on LRC Not Lit

Check that the cables have not been reversed during installation.

Computer Doesn’t Recognize the SA2016 Subsystem

1. Check that the FC-AL interface cables from the SA2016 enclosure to the host computer, or RAID controller, are fitted correctly.
2. Check the Enclosure ID switch settings on your SA2016 subsystem and on your system host.
3. Check that the LEDs on all installed drive carrier modules are illuminated Green. Note that the drive LEDs will not be lit during drive spinup.
4. Check that all drive carrier modules have been correctly installed.
5. Check that there is a valid FC_AL signal present at the I/O connector (see “Internal Loop Structures” on page 172.) If there is no signal present check that the cable has not been reversed during installation.
6. Check the LRC I/O module setup as follows:
 - Check that the LRC I/O module has been correctly installed and all external links and cables are securely fitted.
 - Check that the maximum cable length has not been exceeded.

LEDs

Green LEDs are always used for good or positive indication, flashing Green/Amber if non-critical conditions exist. Amber LEDs indicate there is a critical fault present within the module.

Power Supply/Cooling Module

The Power Supply Cooling LEDs are shown in Table 8-1.

- 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 8-1 PSU LEDs

AC PSU	LED	Status
 <p>Power AC Fan Power Good Fail Fault Fault</p>	<ul style="list-style-type: none"> • PSU Good • AC input Fail • Fan Fault • DC Output Fail 	<p>Green</p> <p>Amber</p> <p>Amber</p> <p>Amber</p>

Ops Panel

The Ops Panel displays the aggregated status of all the modules. The Ops Panel LEDs are shown in Figure 8-1 and defined in Table 8-2. For details on how to remove and replace a module see “Replacing a Module” on page 220.




Figure 8-1 Ops Panel

Note: The Ops Panel is supplied as an integral part of the Enclosure core product and is not user replaceable.

The Ops Panel LED states are shown in Table 8-2.

Ops Panel LEDs

Table 8-2 Ops Panel LED States

LED	Definition	Color	Normal Status	Fault Status
Invalid Address	Indicates that an invalid Enclosure ID has been selected or that the selection has changed after Power On	Amber	Off	Flashing
Power On	Enclosure Powered On	Green	On	Off
System Fault	System/SCM Fault	Amber	Off	On
PSU/Cooling Fault	PSU Cooling fault or enclosure over-temperature.	Amber	Off	On
2Gb Link Speed	Indicates link speed	Green	Off	On
Hub Mode	Not Used	Green	Off	Off

LRC I/O Module LEDs

The LRC I/O module LEDs are shown in Table 8-3.

Table 8-3 LRC I/O Module LEDs

LED	Definition	Color	Normal Status	Fault Status
FC Host Port 0 Signal Good	Incoming FC signal is GOOD No connection or incorrect connection Invalid SFP connection	Green	On	Off Flashing

Table 8-3 LRC I/O Module LEDs

FC Host Port 1 Signal Good	Incoming FC signal is GOOD	Green	On	Off Flashing
	No connection or incorrect connection Invalid SFP connection			
Router Status	Storage Router Device Ready	Green	On	Off
	Storage Router Device not ready or defective			
ESI/LRC Module Fault	Fault present (also <i>On when booting</i>)	Amber	On	Off
	Successful controller initialization			

Audible Alarm

The Ops Panel also includes an Audible Alarm which indicates when a fault state is present. However, when a Drive Fault and/or LUN Fault condition exists the audible alarm will not sound. The following conditions will activate the Audible Alarm:

- Fan Fault
- Voltage out of range
- Thermal overrun
- System fault

Audible Alarm Mute

When the Audible Alarm sounds, it may be muted by pressing the Alarm Mute push-button. Automatic muting will take place after two minutes if the mute switch is not manually operated. The Alarm Mute push-button is located above the indicators on the Ops Panel (see Figure 8-1 on page 212).

When the alarm is muted it will continue to sound with short intermittent bleeps to indicate that a problem still exists. It will be silenced when all problems are cleared. (See also Thermal Shutdown states in “Thermal Shutdown” on page 218).

LED Test Mode

The Alarm Mute push-button can also be used to test the LEDs on the Ops Panel. When the Mute push-button is held, all LEDs will be illuminated if there are no faults present.

Troubleshooting

The following sections describe common problems, with possible solutions, which can occur with your SA2016 system

System Faults

Symptom	Cause	Action
<ol style="list-style-type: none"> 1 The SYSTEM LED will illuminate AMBER on the LRC 2 Audible Alarm sound 	<p>The ESI processor has detected an internal fault (e.g. failure of an internal communications path)</p>	<ol style="list-style-type: none"> 1 Check for other AMBER LED indications on the Power Supply/Cooling modules. If there is a PSU error present there may be a communications problem with that Power Supply/Cooling module. Remove and then re-fit the module, if the problem persists then change the module. 2 Check for other AMBER LED indications on the drive carriers. If none are evident then there may either be an ESI processor problem or a Backplane problem. 3 Ops Panel module faulty. Please contact your supplier.

Note: See also “Thermal Shutdown” on page 218.

Power Supply/Cooling Faults

Symptom	Cause	Action
1 Ops Panel FAULT LED AMBER	1 Any power fault.	1 Check Power On/Off Switch on rear of Power Supply/Cooling module is switched ON.(not accessible on later models)
2 An AMBER LED on one or more Power Supply/Cooling Modules.	2 A fan failure.	2 Check AC Mains Connections to Power Supply/Cooling module is live.
3 Audible Alarm Sounding.	3 A thermal condition which could cause PSU overheating.	3 Disconnect the Power Supply/Cooling module from mains power and remove the module from the system. Re-install: if problem persists, replace Power Supply/Cooling Module. 4 Reduce the ambient temperature.

Thermal Control

The SA2016 Enclosure uses extensive thermal monitoring and takes a number of actions to ensure component temperatures are kept low and also to minimize acoustic noise. Air flow is from front to rear of the enclosure.

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. 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. Note: This threshold changes according to the number of drives and power supplies 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

Symptom	Cause	Action
<ol style="list-style-type: none">1 Ops Panel FAULT LED AMBER.2 An AMBER LED on one or more Power Supply/Cooling Modules.3 Audible Alarm Sounding.4 Air temperature exiting PSU above 55°C.	If the internal temperature measured in the airflow through the enclosure exceeds a pre-set threshold a thermal alarm will sound.	<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: For thermal warnings please refer to the SES specification.

An Enclosure will shut down when a critical temperature threshold is exceeded in order to prevent permanent damage to the disk drives.

Drive Carrier Module Faults

Disk drive status is monitored by a Green LED and an Amber LED mounted on the front of each Drive Carrier Module, providing the following indications:

Table 8-4 LED Functions

State	Green	Amber
No drive fitted	Off	Off
Drive Power ON	On	Off
Drive Activity	On/Blink off	Off
Drive Fault	On	On

Drive activity - LED may be off for a length of time during power up.

Dummy Carrier Modules

Dummy Carrier modules must be fitted to all unused drive bays to maintain a balanced air flow.

Dealing with Hardware Faults

Ensure that you have obtained a replacement module of the same type *before* removing any faulty module.

Warning: If the SA2016 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 warranty.

- Replace a faulty drive with a drive of the same type and equal or greater capacity.
- All drive bays must be fitted with a Drive Carrier module or a dummy carrier module in order to maintain a balanced air flow.

- All the supplied plug-in power supply units, 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.

In addition, each enclosure contains two Power Supply /Cooling modules, either of which can maintain power and cooling to the subsystem while the other is replaced.

Replacing a Module

Warning: 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, "Installation of the SA2016 Drive Enclosure" for information on the initial installation of the plug-in modules in the SA2016 enclosure.

Warning: Observe all conventional ESD precautions when handling SA2016 modules and components. Avoid contact with Backplane components and module connectors, etc.

Power Supply/Cooling Modules

Warning: Do not remove covers from the Power Supply/Cooling (PSU) module. Danger of electric shock inside. Return the PSU to your supplier for repair.

Removing a Power Supply/Cooling Module

Warning: Do not remove the faulty Power Supply/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 Power Supply/Cooling module.

As there should always be two power supply units installed, you can continue working while replacing the faulty module.

Note that the power cord may also need to be removed from the PDU in order to have enough clearance in order to remove the PSU.

1. Make sure you identify the faulty Power Supply/Cooling module correctly, from the two modules installed.
2. Switch off and disconnect the power supply cord.
3. Squeeze the two latches on the PSU handle together (Figure 8-2 on page 222) and open the handle to cam the PSU out of the enclosure (Figure 8-3 on page 223).
4. Grip the handle and withdraw the PSU (Figure 8-4 on page 224).

Inserting the Power Supply/Cooling Module

1. Check for damage, especially to the rear connector on the PSU.

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 PSU handle (Figure 8-3 on page 223) in the open position, slide the module into the enclosure.

Important: install the Power Supply/Cooling module in the right hand bay (Rear Bay 1) of the enclosure in an “upside down” orientation.

3. Cam the module home by manually closing the PSU handle (see Figure 8-4 on page 224). A click should be heard as the handle latches engage (see Figure 8-2 on page 222).
4. Connect the power supply cord to the power source and switch the power supply ON.

Note: The alarm will sound until the new Power Supply/Cooling module is operating correctly.



Figure 8-2 Removing/Inserting an AC Power Supply/Cooling Module (1)




Figure 8-3 Removing/Inserting an AC Power Supply/Cooling Module (2)




Figure 8-4 Removing/Inserting an AC Power Supply/Cooling Module (3)

Ops Panel

The Ops Panel is an integral part of the enclosure chassis assembly and is not field replaceable.

LRC I/O Module

Please refer to “LRC I/O Module Installation” on page 190 for full information on installing the LRC I/O module.

Removing the the LRC 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.

Electrically this unit can be hot plugged but the FC connection will be disconnected.

1. Using two hands, grasp each latch between the thumb and forefinger of each hand. Squeeze thumb and forefinger together to release the latch. Pull the latches forward to cam the module out of the enclosure (Figure 8-7).
2. Grip the latch handles and withdraw the module (Figure 8-5-).




Figure 8-5 Removing an LRC I/O Module (1)




Figure 8-6 Removing an LRC I/O Module (2)

Inserting the the LRC I/O Module

Important: If only one LRC module is fitted, it must be installed in Module A location (Rear Bay 3) [see Figure 6-1 on page 167] and an I/O blank module fitted in the unused bay.

1. With the latch in the open position, slide the LRC I/O module into the enclosure until the latch engages automatically.
2. Cam the module home by manually closing the latches (see Figure 8-7).
3. A click should be heard as the latch engages.




Figure 8-7 LRC I/O Module Latch Operation

Drive Carrier Module

Please see “Drive Carrier Installation” on page 195, for information on the initial installation of the plug-in modules in the SA2016 enclosure.

Warning: Observe all conventional ESD precautions when handling SA2016 modules and components. Avoid contact with backplane components and module connectors, etc.

Removal and Replacement

Caution: Drive spin down

Damage can occur to a drive if it is removed while still spinning. If possible use the operating system to spindown 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:

1. Release the carrier handle, by pressing the latch in the handle towards the right
-

Note: The anti-tamper lock must be off.

2. Gently withdraw the Drive Carrier Module approximately 1 inch (25mm) and wait 30 seconds.
3. Withdraw the module from the drive bay and fit a replacement module in accordance with the instructions in “Drive Carrier Installation” on page 195.

Spare Parts and Ancillary Items

The following replaceable parts are available for the SA2016 subsystem:

- Chassis (including Backplane)
- AC Power Supply/Cooling Module
- Drive Carrier Module
- External FC-AL Interface Cables
- LRC Y cable
- SFP connectors
- 19 inch rack mounting rail kit
- Dummy Carrier Module
- (Country specific) power cords
- Keys for Drive Carrier modules.
- All documentation

Technical Specifications for SGI InfiniteStorage RM660 and RM610

This section provides the technical specifications for the SGI InfiniteStorage RM660 and the SGI InfiniteStorage RM610.

SGI InfiniteStorage RM660 Technical Specification

The technical specifications of the RM660 are given below.

Configuration, Performance & Capacity

Host interface	Fibre Channel (copper or fiber optic)
Drive interface	FC-AL
Management interface	RS-232 and Ethernet (Telnet)
Cache memory	2.5GB
Independent drive channels	10
No. of host ports	4
No. of host connections	Up to 512
No. of drives supported	Up to 1250 (1000 drives for data)
No. of LUN groups supported	Up to 128 (each can be subdivided into 64 equally-sized LUN segments)
Hot spare capability	Yes (up to 125 spare modules)
Fibre Channel host transfer rate	Up to 200MB/s (×4)
Fibre Channel drive transfer rate	Up to 1Gb/s (×10);

Configuration, Performance & Capacity

Mirrored cache capability	Yes with couplet configuration
Full duplex 10/100BaseT onboard	Yes

Reliability

SES (SCSI Enclosure Services) protocol support	Yes
Temperature monitoring	Yes
Redundant hot-swappable power supplies	2 modules
Redundant hot-swappable cooling fans	1 module

Physical, Power & Environmental

RM610 chassis (H×W×D, Weight)	1.75" × 17.0" × 24.4", 22.5 lbs
Electrical/AC	100-120V / 200-240V @ 47-63Hz
Power consumption (maximum current)	3.8A @ 110VAC, 1.8A @ 230VAC Couplet: 6.5A @ 110VAC, 3.1A @ 230VAC
Power consumption (average current)	2.8A @ 110VAC, 1.35A @ 230VAC Couplet: 4.6A @ 110VAC, 2.2A @ 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
Certification	FCC, UL, cUL (CSA,TUV), CE

Specifications subject to change without notice.

SGI InfiniteStorage RM610 Technical Specifications

The technical specifications of the RM660 are given below.

Configuration, Performance & Capacity

Host interface	Fibre Channel (copper or fiber optic)
Drive interface	FC-AL
Management interface	RS-232 and Ethernet (Telnet)
Cache memory	1.5GB
Independent drive channels	6
No. of host connections	Up to 512
No. of drives supported	Up to 750(500 drives for data)
No. of LUN groups supported	Up to 128 (each can be subdivided into 64 equally-sized LUN segments)
Hot spare capability	Yes (up to 125 spare modules)
Fibre Channel host transfer rate	Up to 100MB/s (x4)
Fibre Channel drive transfer rate	Up to 1Gb/s (x6);
Mirrored cache capability	Yes with couplet configuration
Full duplex 10/100BaseT onboard	Yes

Reliability

SES (SCSI Enclosure Services) protocol support	Yes
Temperature monitoring	Yes
Redundant hot-swappable power supplies	4
Redundant hot-swappable cooling fans	2

Physical, Power & Environmental	
RM660 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

Specifications subject to change without notice.

Safety Guidelines for SGI InfiniteStorage RM610/RM660 Rack Installation

Given below are the safety guidelines that should be followed when installing the SGI InfiniteStorage RM610 or RM660 in a rack:

Elevated Operating Ambient Temperature

If 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, "Technical Specifications for SGI InfiniteStorage RM660 and RM610".

Reduced Air Flow

Installation of the system in a rack should be such that the amount of air flow required for safe operation of the system is not compromised.

Mechanical Loading

Mounting of the system in the rack should be such that a hazardous condition does not occur due to uneven mechanical loading.

Circuit Overloading

Consideration should be given to the connection of the system to the supply circuit and the effect that overloading of circuits might have on overcurrent protection and supply wiring.

Reliable Earthing

Reliable earthing of rack-mounted systems should be maintained. Particular attention should be given to supply connections other than direct connections to the branch circuit (e.g. use of power distribution units).

Using the SFx016 Drive Enclosures

This appendix provides information on using the SF6016, SF4016, and SF2016 drive enclosures with the SGI InfiniteStorage RM660 and RM610.

The Drive Enclosure

The SFx016 drive enclosure is designed to be used with the RM660 and RM610. Each enclosure can hold up to 16 removable Fibre Channel drive modules in the front and contain interface modules and redundant power supply/cooling modules in the back (Figure C-1).




Figure C-1 The SFx016 Enclosure - Front and Rear Views

Enclosure Bay Numbering Convention

The enclosure bay numbering convention is shown in Figure C-2. The front bays are numbered 1 to 4 from left to right, viewed from the front, and 1 to 4, from top to bottom. Drive module locations are identified from a matrix of the top and side numbers. For example, Drive 14 is installed in Bay 3/4. The rear bays are numbered 1 to 5 from right to left, viewed from the rear.

Note: For proper operation of SES (SCSI Enclosure Services), drive module must always be installed in bays 1/1 (drive 0) and 4/4 (drive 15).

Viewing from Front

Column				Row
1	2	3	4	
Drive 0	Drive 1	Drive 2	Drive 3	1
Drive 4	Drive 5	Drive 6	Drive 7	2
Drive 8	Drive 9	Drive 10	Drive 11	3
Drive 12	Drive 13	Drive 14	Drive 15	4

Viewing from Back

Column				
5	4	3	2	1
PSU/ Cooling 1	I/O Module B	I/O Module A	Ops Panel	PSU/ Cooling 2

Figure C-2 Enclosure Bay Numbering Convention**Operators Panel Module**

The Operators (Ops) Panel provides the enclosure with a micro controller which is used to monitor and control the elements, such as power, cooling, temperature, and device status, within the enclosure (Figure C-3).

The Ops panel module is used in the drive enclosure to provide an enclosure management interface through SCSI Enclosure Services (SES) protocol. The SES function is communicated via drives in bays 1/1 and 4/4. Hence, drive bays 1/1 and 4/4 must be occupied.




Figure C-3 Ops Panel Module

The Ops module includes LEDs which indicate the status for all modules, an Audible Alarm which activates when a fault occurs, an Alarm Mute push-button, and an enclosure ID range thumb wheel switch.

The Ops Panel configuration switch functions are shown in Table C-1. **Switch settings are only read at Power On.**

Table C-1 Ops Panel Switch Functions

Switch No. (*see Sw11)	Function	Function When Off		Function When On
1*	Loop Select: Single (1x16) or Dual (2x8)	LRC operates as 2 loops of 8 drives—SF4016 (see also Addressing Mode 2 below)		LRC operates as single loop of 16 drives— SF2016/SF6016
2	Loop Terminate Mode	If no signal is present on External FC port, the loop will be left “open”		If no signal is present on External FC port, the loop will be “healed” internally
3*	Hub Mode Select (not applicable)	RAID host FC ports will be independently connected		RAID host FC ports will be linked together internally
4	Not used			
5 & 6	RAID Host Hub Speed Select switches (not applicable)	Sw 5	Sw 6	Function
		Off	Off	Force 1 Gb/s
		On	Off	Force 2 Gb/s
		Off	On	Reserved
		On	On	Auto loop speed detect (based on LRC port signals)
7 & 8	Drive Loop Speed Select	Sw 7	Sw 8	Function
		Off	Off	Force 1 Gb/s
		On	Off	Force 2 Gb/s
		Off	On	Speed selected by EEPROM bit
		On	On	Auto loop speed detect (based on LRC port signals)

Table C-1 Ops Panel Switch Functions (continued)

Switch No. (*see Sw11)	Function	Function When Off		Function When On
9 & 10	Drive Addressing Mode Selection	Sw 9	Sw 10	Function
		Off	Off	Mode 3—SF2016
		On	Off	Mode 2: Selecting this mode will force dual loop selection—SF4016 (Sw 1 above)
		Off	On	Mode 1—SF6016
		On	On	Mode 0
11	Soft Select	Select functions (marked * in this table) using the values stored in enclosure EEPROM		Select functions (marked * in this table) using the hardware switches
12	Not used			

Power Supply/Cooling Module

Each drive enclosure is equipped with two Power Supply/Cooling modules (Figure C-4). PSU (power supply unit) voltage operating ranges are nominally 115V or 230V AC, selected automatically.

The two modules provide a 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.

Four LEDs mounted on the front panel of the Power Supply/Cooling Module indicate the status of the PSU and the fans.

Note: Module replacement should only take a few minutes to perform but must be completed within 10 minutes from removal of the failed module.




Figure C-4 Power Supply/Cooling Module

I/O Modules

The SFx016 enclosure contains two Loop Resiliency Circuit (LRC) I/O modules (Figure C-5). These plug-in modules are used to interface with the RM660 and expansion enclosures. The FC-AL backplane incorporates two independent loops formed by Port Bypass Circuits within the LRC I/O modules.




Figure C-5 FC-AL LRC I/O Module

Processors housed in the I/O modules provide enclosure management and interface to devices on the backplane, PSU, LRC, and Ops panel to monitor internal functions. These processors operate in a master-slave configuration to allow failover.

Depending on your configuration, the I/O modules are factory-configured to support:

- one loop of 16 drives operating at 1Gb (SF2016 and SF6016)
- 2 loops of 8 drives operating at 1Gb (SF4016)

The internal dual loop structure is shown in Figure C-6.




Figure C-6 Overall Loop Layout

Note: Figure C-6 shows the logical routing of the FC-AL interface through the drive bays, this figure does not represent the physical wiring of the enclosure.

Each loop is routed through independent LRC I/O modules. Either of these may be removed while the other is operating, thus providing fully redundant FC-AL operation.

Each module includes four Small Form-factor Pluggable (SFP) connector modules with copper HSSDC-2 connectors. Auto-bypass at the output ports is provided. The module also contains one Fault LED (amber) and four FC-AL signal indicators (Figure C-7).




Figure C-7 I/O Module Connectors and LEDs

Drive Modules (FC-AL)

Each disk drive is encapsulated inside an aluminum canister which provides thermal conduction, radio frequency and electromagnetic induction protection and affords the drive maximum physical protection (Figure C-8).




Figure C-8 Drive Module

The module contains an ergonomic handle which provides the following functions:

- Camming of carrier into and out of drive bays
- Positive “spring loading” of the drive/backplane connector
- A tamper resistant lock operated by a T10 Security Torx type bit

Drive LEDs

Each drive module contains two LED indicators, a Status (Green) LED and a Fault (Amber) LED (Figure C-9). In normal operation the Status LED will be ON and will flash as the drive operates. The Fault LED is OFF during normal operation and ON if there is a drive fault present.




Figure C-9 LEDs on Drive Module Front

Anti-Tamper Locks

Anti-tamper locks are fitted in the drive module handles (Figure C-10).




Figure C-10 Anti-Tamper Lock on Drive Modules

The lock can be accessed through the small cutout in the latch section of the handle, using a T10 Security Torx type bit. These locks are provided to disable the normal “pinch” latch action of the module handle and so prevent accidental or unauthorized removal of drives.

Note: Do NOT overtighten the anti-tamper lock!

Dummy Drive Modules and Blank Modules

Dummy drive modules are needed for fitting in all unused drive bays in the front of the enclosure. They are designed as integral drive module front caps with handles and must be installed in all unused drive bays to maintain a balanced air flow.

Blank modules must also be used to cover all vacant bays at the back in order to maintain an efficient air flow pattern within the enclosure.

Warning: Operation of the enclosure with ANY modules missing will disrupt the airflow pattern and the drives will not receive sufficient cooling. It is ESSENTIAL that all bays are filled before operating the unit. Dummy modules and/or blank modules are available for this purpose.

Visible and Audible Alarms

The functional modules have associated status LED indicators. LEDs show constant green for good or positive indication. Constant amber LEDs indicate there is a fault present within that module. The Ops Panel shows a consolidated status for all modules.

Warning: The Ops Panel is an integral part of the SFx016 enclosure and should only be replaced by trained personnel.

The Ops Panel also incorporates an Audible Alarm which beeps when the following occurs:

- fan slow
- voltage out of range
- over temperature
- thermal overrun
- UPS two minute warning
- system fault

You can use the Alarm Mute push-button on the module to turn down the alarm. The mute function provides the reduction of the audible alarm, but leaves an intermittent beep (at approximately 10 second intervals) to show that the failure is still present. The alarm will be turned off completely when all problems are cleared. Automatic muting will take place after two minutes if the mute switch is not manually operated.

Setting Up the Drive Enclosures

The procedures below explain how to install the enclosures in a rack, how to connect the SFP cables between the enclosures and RM module, and how to install the drive modules.

Installing the Drive Enclosures in Rack

1. Assemble the two rail/support assemblies.

Using three #10 screws, washers, and keps nuts, assemble the rail inside the support (Figure C-11). Only finger-tighten the screws.




Figure C-11 Assembling Rail to Support (Left Hand Assembly)

2. Adjust the length of the rail/support assembly to fit the depth of the rack. Then tighten the three screws and nuts on the assembly. Fasten the assembly to the rack frame, as described below. Make sure you leave enough space above the rail/support assembly to accommodate the enclosure. Then tighten all screws.

On square-hole and through-hole racks, use #10 screws and washers (13-00024-001) for the rear. Use #10 screws for the front (see Figure C-12).




Figure C-12 Fastening Rail/Support Assembly to Square/Through-Hole Rack (Left Hand Side)

On threaded-hole racks, use #10 screws and washers for the rear. Use #10 screws, #10 washers, and washers for the front (see Figure C-13).




Figure C-13 Fastening Rail/Support Assembly to Threaded-Hole Rack (Left Hand Side)

3. Similarly, assemble and install the rail/support assembly in the other side of the rack, making sure that the two rails are level (horizontally) in the rack.
4. To reduce the weight of the unit, it is best to remove the power supplies and drive modules from the enclosure.

To remove a power supply, squeeze the two latches on the handle together and open the handle to cam the module out of the enclosure.

To remove a drive module, press the latch in the handle towards the right to release the handle (Figure C-14).




Figure C-14 Latch Operation on Drive Module

Then swing open the handle and pull the module out of the bay. If the anti-tamper lock is activated, see Step [5] on page 283 for information on how to de-activate the lock. Dummy drive modules may be left in the slots.

- Carefully place the chassis on the two rail/support assemblies and push it all the way into the rack.

If in doubt about correct orientation, the drive bays (at front) should have their black drive connectors toward the bottom of each bay.

- Using two #10 screws, secure rear of chassis to rail/support assemblies (see Figure C-15).

Secure the front of the chassis to the rack frame:

For square-hole and through-hole racks, attach cage nuts (12-00022-002) or speed nuts (12-00062-010) to frame as shown in Figure C-15. Then use #10 screws, #10 washers, and washers to secure the chassis.




Figure C-15 Securing Chassis to Square-Through-Hole Rack (Left Hand)

For threaded-hole racks, use #10 screws, #10 washers, and washers to secure the chassis, as shown in Figure C-16.




Figure C-16 Securing Chassis to Threaded-Hole Rack (Left Hand)

7. Replace the two power supplies in the back.
Insert one module into the left most bay and the other module into the right most bay. Swing up the handle to cam the module home. Make sure that the latches are engaged.
8. Replace all the drive modules.
9. Repeat Steps [1] to [8] above to install the other enclosures in the rack.

Note: Please follow the safety guidelines for rack installation given in Appendix B.

Verify Configuration Switch Settings on Ops Panel Modules

On all enclosures, verify that the settings for the configuration switches on the Ops panel are the same as listed below.




Figure C-17 Settings on SF6016 Enclosures

Switch Setting			
ON	OFF	Option Description	Setting
1		Loop Select: Single (1x16)	LRC operates as 1 loop of 16 drives
	2	Loop Terminate Mode	If no signal is present on External FC port, the loop will be left "open"
	3	Hub Mode Select (RAID only)	Not applicable
	4	Not used	
	5	RAID Host Hub Speed Select	Not applicable
	6		Not applicable
	7	Drive Loop Speed Select	Sw 7 off & 8 off = 1Gb/s
	8		
	9	Drive Addressing Mode Select	9 off & 10 on = Mode 3
	10		
	11	Soft Select	Use values set by hardware switches 1 and 3
	12	Not used	




Figure C-18 Settings on SF4016 Enclosures

Switch Setting

ON	OFF	Option Description	Setting
	1	Loop Select: Dual (2x8)	LRC operates as 2 loops of 8 drives (see also Addressing Mode 2 below)
2		Loop Terminate Mode	If no signal is present on External FC port, the loop will be "healed" internally
	3	Hub Mode Select (RAID only)	Not applicable
	4	Not used	
	5	RAID Host Hub Speed Select	Not applicable
	6		Not applicable
	7	Drive Loop Speed Select	Sw 7 & 8 off = 1Gb/s
	8		
9		Drive Addressing Mode Select	9 on & 10 off = Mode 2: forces dual loop selection (Sw 1 above)
	10		
11		Soft Select	Use values set by hardware switches 1 and 3
	12	Not used	




Figure C-19 Settings on SF2016 Enclosures

**Switch
Setting**

ON	OFF	Option Description	Setting
1		Loop Select: Single (1x16)	LRC operates as 1 loop of 16 drives
2		Loop Terminate Mode	If no signal is present on External FC port, the loop will be "healed" internally
	3	Hub Mode Select (RAID only)	Not applicable
	4	Not used	
	5	RAID Host Hub Speed Select	Not applicable
	6		Not applicable
	7	Drive Loop Speed Select	Sw 7 & 8 off = 1Gb/s
	8		
	9	Drive Addressing Mode Select	9 on & 10 off = Mode 3
	10		
11		Soft Select	Use values set by hardware switches 1 and 3
	12	Not used	

Cable Connections and Enclosure ID Settings on the RM660

Connecting SF6016 Enclosures on the RM660

The basic configuration consists of two SF6016 enclosures. Each enclosure will connect to five channels on the RM660. Drive bay 4/4 is not used (Figure C-20). The enclosures are labelled by channels at the back (Figure C-21).

Viewing from Front

1A	2A	3A	1E
1B	2B	3B	2E
1C	2C	3C	3E
1D	2D	3D	

1F	2F	3F	1S
1G	2G	3G	2S
1H	2H	3H	3S
1P	2P	3P	

Figure C-20 Drive Locations on SF6016 with Enclosure ID Set to 1




Figure C-21 Enclosure Label at Rear

There are two I/O modules installed in each SF6016 enclosure. Each module connects to one of the two internal drive loops. In dual RM660 configuration, connections to the two I/O modules will provide redundant data paths.

Figure C-22 illustrates how to connect the first set of enclosures to the RM660. Each enclosure can hold up to 15 drives, so this configuration can house 3 full tiers (30 drives).

In dual mode, RM660 Unit 1 connects to I/O module A (right) of the enclosure and RM660 Unit 2 connects to I/O module B.




Figure C-22 Connecting the SF6016 Enclosures to RM660

Figure C-23 illustrates how to daisy-chain the enclosures.




Figure C-23 Daisy-Chaining the SF6016 Enclosures

Using the ID Range switch on the Ops panels (Figure C-24), select an ID number for each enclosure. Use ID “1” on the first set of enclosures that are directly connected to the RM660s. Then use ID “2” for the second set of enclosures in the chain, ID “3” for the third set of enclosures, and so on.




Figure C-24 Enclosure ID Switch

Connecting SF4016 Enclosures on the RM660

The basic configuration consists of five SF4016 enclosures. Each enclosure will connect to two channels on the RM660. You may daisy-chain up to fifteen SF4016 enclosures to each channel.

There are two I/O modules installed in each SF4016 enclosure. Each module connects to one of the two internal drive loops. In dual RM660 configuration, connections to the two I/O modules will provide redundant data paths.

The following steps explain how to connect 5 enclosures to the RM660. Each enclosure can hold up to 16 drives, so this configuration can house 8 full tiers (80 drives).

1. For Channel A: Connect a SFP cable between the “DISK A” port on the back of the RM660 and the upper right “1” connector on Enclosure #1 (Figure C-25).
2. For Channel B: Connect a SFP cable between the “DISK B” port on the back of the RM660 and the lower left “1” connector on Enclosure #1 (Figure C-25).




Figure C-25 Connecting the SF4016 Enclosures to the RM660 (1)

3. For Channel C: Connect a SFP cable between the “DISK C” port on the back of the RM660 and the upper “1” connector on Enclosure #2’s right I/O module (Figure C-26).
4. For Channel D: Connect a SFP cable between the “DISK D” port on the back of the RM660 and the lower “1” connector on Enclosure #2’s left I/O module (Figure C-26).
5. Similarly, connect the other enclosures to Channels EFGHPS on the RM660 (Figure C-26).

RM660 Channels	TO	SF4016 Enclosures	
A	TO	Enclosure #1	upper “1” connector on right I/O module
B	TO	Enclosure #1	lower “1” connector on left I/O module
C	TO	Enclosure #2	upper “1” connector on right I/O module
D	TO	Enclosure #2	lower “1” connector on left I/O module

RM660 Channels	TO	SF4016 Enclosures	
E	TO	Enclosure #3	upper "1" connector on right I/O module
F	TO	Enclosure #3	lower "1" connector on left I/O module
G	TO	Enclosure #4	upper "1" connector on right I/O module
H	TO	Enclosure #4	lower "1" connector on left I/O module
P	TO	Enclosure #5	upper "1" connector on right I/O module
S	TO	Enclosure #5	lower "1" connector on left I/O module

- Using the ID Range switch on the Ops panel modules, select ID "1" for all five enclosures.

Daisy-Chaining the SF4016 Enclosures on the RM660

The following steps explain how to daisy-chain the SF4016 enclosures. You may daisy-chain up to fifteen enclosures to each channel, giving a total of 120 tiers.

1. Follow the steps given page 259 to connect the first set of enclosures (#1 to #5) to the RM660.
2. To daisy-chain the second set of enclosures (Figure C-27):
 - For Channel A: Connect a daisy-chain cable between the upper right “2” connector on Enclosure #1 and the upper right “1” connector on Enclosure #6.
 - For Channel B: Connect a daisy-chain cable between the lower left “2” connector on Enclosure #1 and the lower left “1” connector on Enclosure #6.
 - Similarly, connect the other expansion enclosures to Channels CDEFGHPS.
3. Repeat Step [2] above to connect additional sets of expansion enclosures.
4. Using the ID Range switch on the Ops panels, select an ID number for each enclosure. Use ID “1” on the first set of enclosures that are directly connected to the RM660s. Then use ID “2” for the second set of enclosures in the chain, ID “3” for the third set of enclosures, and so on.




Figure C-27 Daisy-Chaining the SF4016 Enclosures

Connecting SF2016 Enclosures on the RM660

The basic configuration consists of ten SF2016 enclosures. Each enclosure will connect to one channel on the RM660. You may daisy-chain up to seven SF2016 enclosures to each channel.

There are two I/O modules installed in each SF2016 enclosure. Each module connects to one of the two internal drive loops. In couplet RM660 configuration, connections to the two I/O modules will provide redundant data paths.

The following steps explain how to connect 10 enclosures to the RM660. Each enclosure can hold up to 16 drives, so this configuration can house 16 full tiers (160 drives).

1. For Channel A: Connect a SFP cable between the “DISK A” port on the back of the RM660 and the upper right “1” connector on Enclosure #1 (Figure C-28).
2. For Channel B: Connect a SFP cable between the “DISK B” port on the back of the RM660 and the upper right “1” connector on Enclosure # (Figure C-28).
3. Similarly, connect the other enclosures to Channels CDEFGHPS on the RM660 (Figure C-28 and Figure C-29).

RM660 Channels	TO	SF2016 Enclosures
A	TO	1
B	TO	2
C	TO	3
D	TO	4
E	TO	5
F	TO	6
G	TO	7
H	TO	8
P	TO	9
S	TO	10

4. Using the ID Range switch on the Ops panel modules, select ID “1” for all ten enclosures.




Figure C-28 Connecting the SF2016 Enclosures to the RM660 (1)




Figure C-29 Connecting the SF2016 Enclosures to the RM660 (2)

Daisy-Chaining the SF2016 Enclosures on the RM660

The following steps explain how to daisy-chain the SF2016 enclosures. You may daisy-chain up to seven enclosures to each channel, giving a total of 112 tiers.

1. Follow the steps given page 265 to connect the first set of enclosures (#1 to #10) to the RM660.
2. To daisy-chain the second set of enclosures (Figure C-30):
 - For Channel A: Connect a daisy-chain cable between the upper right “2” connector on Enclosure #1 and the upper right “1” connector on Enclosure #11.
 - For Channel B: Connect a daisy-chain cable between the upper right “2” connector on Enclosure #2 and the upper right “1” connector on Enclosure #12.
 - Similarly, connect the other expansion enclosures to Channels CDEFGHPS.
3. Repeat Step [2] above to connect additional sets of expansion enclosures.
4. Using the ID Range switch on the Ops panels, select an ID number for each enclosure. Use ID “1” on the first set of enclosures that are directly connected to the RM660s. Then use ID “2” for the second set of enclosures in the chain, ID “3” for the third set of enclosures, and so on.




Figure C-30 Daisy-Chaining the SF2016 Enclosures

Cable Connections and Enclosure ID Settings on the RM610

Connecting SF6016 Enclosures on the RM610

The basic configuration consists of one SF6016 enclosure. Each enclosure will contain one spare drive (Figure C-31). The enclosures are labelled by channels at the back (Figure C-32).

Viewing from Front	1A	2A	3A	1P
	1B	2B	3B	2P
	1C	2C	3C	3P
	1D	2D	3D	S

Figure C-31 Drive Locations on the FS6016 with Enclosure ID Set to 1




Figure C-32 Enclosure Label at Rear

There are two I/O modules installed in each SF6016 enclosure. Each module connects to one of the two internal drive loops. In dual RM610 configuration, connections to the two I/O modules will provide redundant data paths.

Figure C-33 illustrates how to connect the first enclosure to the RM610. Each enclosure can hold up to 16 drives, so this configuration can house 3 tiers with one spare drive.

In dual mode, RM610 Unit 1 connects to I/O module A (right) of the enclosure and RM610 Unit 2 connects to I/O module B.




Figure C-33 Connecting the SF6016 Enclosure to the RM610

Figure C-34 illustrates how to daisy-chain the enclosures.




Figure C-34 Daisy-Chaining the SF6016 Enclosures

Using the ID Range switch on the Ops panels (Figure C-35), select an ID number for each enclosure. Use ID “1” on the first enclosure that is directly connected to the RM610s. Then use ID “2” for the second enclosure in the chain, ID “3” for the third enclosure, and so on.




Figure C-35 Enclosure ID Switch

Connecting SF4016 Enclosures on the RM610

The basic configuration consists of three SF4016 enclosures. Each enclosure will connect to two channels on the RM610. You may daisy-chain up to fifteen SF4016 enclosures to each channel.

There are two I/O modules installed in each SF4016 enclosure. Each module connects to one of the two internal drive loops. In couplet RM610 configuration, connections to the two I/O modules will provide redundant data paths.

The following steps explain how to connect 3 enclosures to the RM610. Each enclosure can hold up to 16 drives, so this configuration can house 8 full tiers (48 drives).

1. For Channel A: Connect a SFP cable between the “6-A” port on the back of the RM610 and the upper right “1” connector on Enclosure #1 (Figure C-36).
2. For Channel B: Connect a SFP cable between the “5-B” port on the back of the RM610 and the lower left “1” connector on Enclosure #1 (Figure C-41).




Figure C-36 Connecting the SF4016 Enclosures to the RM610(1)

3. Similarly, connect the other enclosures to Channels CDPS on the RM610 (Figure C-37).

RM610 Channels	TO	SF4016 Enclosures	
A	TO	Enclosure #1	upper right "1" connector
B	TO	Enclosure #1	lower left "1" connector
C	TO	Enclosure #2	upper right "1" connector
D	TO	Enclosure #2	lower left "1" connector
P	TO	Enclosure #3	upper right "1" connector
S	TO	Enclosure #3	lower left "1" connector

- Using the ID Range switch on the Ops panel modules, select ID "1" for all three enclosures.




Figure C-37 Connecting the SF4016 Enclosures to the RM610

Daisy-Chaining the SF4016 Enclosures on the RM610

The following steps explain how to daisy-chain the SF4016 enclosures. You may daisy-chain up to fifteen enclosures to each channel, giving a total of 120 tiers.

1. Follow the steps given on page 273 to connect the first set of enclosures (#1 to #3) to the RM610.
2. To daisy-chain the second set of enclosures (Figure C-38):
 - For Channel A: Connect a daisy-chain cable between the upper right “2” connector on Enclosure #1 and the upper right “1” connector on Enclosure #4.
 - For Channel B: Connect a daisy-chain cable between the lower left “2” connector on Enclosure #1 and the lower left “1” connector on Enclosure #4.
 - Similarly, connect the other expansion enclosures to Channels CDPS.
3. Repeat Step [2] above to connect additional sets of expansion enclosures.
4. Using the ID Range switch on the Ops panels, select an ID number for each enclosure. Use ID “1” on the first set of enclosures that are directly connected to the RM610s. Then use ID “2” for the second set of enclosures in the chain, ID “3” for the third set of enclosures, and so on.




Figure C-38 Daisy-Chaining the SF4016 Enclosures

Connecting SF2016 Enclosures on the RM610

The basic configuration consists of six SF2016 enclosures. Each enclosure will connect to one channel on the RM610. You may daisy-chain up to seven SF2016 enclosures to each channel.

There are two I/O modules installed in each SF2016 enclosure. Each module connects to one of the two internal drive loops. In couplet RM610 configuration, connections to the two I/O modules will provide redundant data paths.

The following steps explain how to connect 6 enclosures to the RM610. Each enclosure can hold up to 16 drives, so this configuration can house 16 full tiers (96 drives).

1. For Channel A: Connect a SFP cable between the “6-A” port on the back of the RM610 and the upper right “1” connector on Enclosure #1 (Figure C-39).
2. For Channel B: Connect a SFP cable between the “5-B” port on the back of the RM610 and the upper right “1” connector on Enclosure #2. (Figure C-39)
3. Similarly, connect the other enclosures to Channels CDPS on the RM610. (Figure C-39)

RM610 Channels	TO	SF2016 Enclosures
A	TO	1
B	TO	2
C	TO	3
D	TO	4
P	TO	5
S	TO	6

4. Using the ID Range switch on the Ops panel modules, select ID “1” for all six enclosures.




Figure C-39 Connecting the SF2016 Enclosures to the RM610

Daisy-Chaining the SF2016 Enclosures on the RM610

The following steps explain how to daisy-chain the SF2016 enclosures. You may daisy-chain up to seven enclosures to each channel, giving a total of 112 tiers.

1. Follow the steps given page 265 to connect the first set of enclosures (#1 to #6) to the RM610.
2. To daisy-chain the second set of enclosures (Figure C-30):
 - For Channel A: Connect a daisy-chain cable between the upper right “2” connector on Enclosure #1 and the upper right “1” connector on Enclosure #7.
 - For Channel B: Connect a daisy-chain cable between the upper right “2” connector on Enclosure #2 and the upper right “1” connector on Enclosure #8.
 - Similarly, connect the other expansion enclosures to Channels CDPS.
3. Repeat Step [2] above to connect additional sets of expansion enclosures.
4. Using the ID Range switch on the Ops panels, select an ID number for each enclosure. Use ID “1” on the first set of enclosures that are directly connected to the RM610s. Then use ID “2” for the second set of enclosures in the chain, ID “3” for the third set of enclosures, and so on.




Figure C-40 Daisy-Chaining the SF2016 Enclosures

Installing Drive Modules

Each SFx016 enclosure holds up to 16 disks. Configuration of disks in the enclosures must be in sets of complete tier. One tier contains eight data drives (Channel A through H), one parity drive (Channel P), and one *optional* spare drive (Channel S). Allocating one spare drive per tier will give you the best data protection but this is not required.

The SES (SCSI Enclosure Services) requires disks to be present in bays 1/1 and 4/4 of each enclosure. When planning your system configuration, please remember that all drive bays must be filled with either a drive module or dummy drive module, no bays should be left completely empty. The drives are not pre-formatted and can be inserted in any order.

When using SF4016 enclosures, the minimum RM660 configuration requires 10 drives. Once the SES positions are filled (indicated by shaded background in Figure C-41), disks can be added in sets of 9.




Figure C-41 Drive Bays on SF4016 Enclosures for SES Communication Paths

When using SF2016 enclosures, the minimum RM660 configuration requires 20 drives. Once the SES positions are filled (indicated by shaded background in Figure C-42), disks can be added in set of 9.




Figure C-42 Drive Bays on SF2016 Enclosures for SES Communication Paths

1. \Release the module handle by pressing the latch in the handle towards the right (Figure C-43).

If the anti-tamper lock is activated, see Step [5] below for information on how to de-activate the lock.




Figure C-43 Opening the Drive Module Handle

2. Insert the module into the bay (Figure C-44).




Figure C-44 Inserting the Drive Module

3. Slide the module, gently, all the way into the bay until it is stopped by the camming lever on the right of the module.

4. Cam the module home - the camming foot on the base of the module will engage into a slot on the enclosure. Continue to push firmly until the handle fully engages. A click should be heard as the latch engages and holds the handle closed.
5. You may activate the anti-tamper locks on the drive module handles.

Carefully insert the lock key provided into the cutout in the handle and align it with the socket. Rotate the key in a clockwise direction until the lock indicator is visible in the aperture beside the socket (Figure C-45). **Do not over-tighten the anti-tamper lock.**

To de-activate, rotate the key in a counter-clockwise direction until the lock indicator is no longer visible in the aperture beside the socket.




Figure C-45 Activating the Anti-Tamper Lock

Powering on the Drive Enclosures

1. Attach the power cords to the Power Supply/Cooling Modules.
For maximum redundancy, connect the two power cords on each SFx016 enclosure to two different AC power circuits.
2. Turn on the switch on all Power Supply/Cooling modules. Verify that the PSU LEDs are green.
3. The enclosures will go through the power up sequence. Check that the Power LEDs on the Ops panels are green.
4. The drives will automatically be spun up in sequence. Wait until all drives are spun up and ready which is indicated by a green Status (upper) LED on the drives.
5. Refer to “Configuring the RM610/RM660” on page 28 for information on how to configure the system.

Verifying Connections for SF6016 Enclosures

It is necessary to verify the connections of the SF6016 drive enclosures that are currently attached to the system. The connections must be verified before the SES can accurately indicate a drive fault through SES commands. Make sure you have selected Tier Mapping Mode 0 as described in “Setting Tier Mapping Mode” on page 32.

Note: This command requires that the RM controllers be offline from any host I/O.

Note: The following procedure shows example output from the RM660.

1. Login to Unit 1 and enter command: `ses on`<Enter>

```
RM660 [1]: ses on
Starting the SES monitors...
Done.

ses_A 8-24 16:41:22 Unverified 6016 drive enclosure detected :
EncID:2000050CC00BE51: channel A
ses_F 8-24 16:41:22 Unverified 6016 drive enclosure detected :
EncID:2000050CC00BE9F: channel F

ses_faul 8-24 16:41:23 Unable to clear fault LED for drive 1A
ses_faul 8-24 16:41:23 Unable to clear fault LED for drive 1B
ses_faul 8-24 16:41:23 Unable to clear fault LED for drive 1C
ses_faul 8-24 16:41:23 Unable to clear fault LED for drive 1D
ses_faul 8-24 16:41:23 Unable to clear fault LED for drive 1E
ses_faul 8-24 16:41:23 Unable to clear fault LED for drive 1F
ses_faul 8-24 16:41:23 Unable to clear fault LED for drive 1G
ses_faul 8-24 16:41:24 Unable to clear fault LED for drive 1H
ses_faul 8-24 16:41:24 Unable to clear fault LED for drive 1P
ses_faul 8-24 16:41:24 Unable to clear fault LED for drive 1S
ses_faul 8-24 16:41:24 Unable to clear fault LED for drive 2A
ses_faul 8-24 16:41:24 Unable to clear fault LED for drive 2B
ses_faul 8-24 16:41:24 Unable to clear fault LED for drive 2C
ses_faul 8-24 16:41:24 Unable to clear fault LED for drive 2D
ses_faul 8-24 16:41:24 Unable to clear fault LED for drive 2E
ses_faul 8-24 16:41:24 Unable to clear fault LED for drive 2F
ses_faul 8-24 16:41:24 Unable to clear fault LED for drive 2G
ses_faul 8-24 16:41:25 Unable to clear fault LED for drive 2H
ses_faul 8-24 16:41:25 Unable to clear fault LED for drive 2P
ses_faul 8-24 16:41:25 Unable to clear fault LED for drive 2S
ses_faul 8-24 16:41:25 Unable to clear fault LED for drive 3A
ses_faul 8-24 16:41:25 Unable to clear fault LED for drive 3B
ses_faul 8-24 16:41:25 Unable to clear fault LED for drive 3C
ses_faul 8-24 16:41:25 Unable to clear fault LED for drive 3D
ses_faul 8-24 16:41:25 Unable to clear fault LED for drive 3E
ses_faul 8-24 16:41:25 Unable to clear fault LED for drive 3F
ses_faul 8-24 16:41:25 Unable to clear fault LED for drive 3G
ses_faul 8-24 16:41:26 Unable to clear fault LED for drive 3H
ses_faul 8-24 16:41:26 Unable to clear fault LED for drive 3P
ses_faul 8-24 16:41:26 Unable to clear fault LED for drive 3S
```

2. Enter command: `ses verify_6016`<Enter>

```

RM660 [1]: ses verify_6016

The connections for an attached 6016 drive enclosure must be
verified before the SCSI Enclosure Services (SES) can accurately
indicate a drive failure within that enclosure.

Select the 6016 drive enclosure to verify :

Enclosure WWN      Tiers      Channels      Verified      Errors
-----
1 - 20000050cc00be51  1, 2, 3  A,,,,,,,,,    No           No
2 - 20000050cc00be9f  1, 2, 3  F,,,,,,,,,    No           No

e - Exit this command

Enter the number of the 6016 enclosure to verify :
    
```

3. Type "1<Enter>" when prompted to select an enclosure to verify.

```

Enter the number of the 6016 enclosure to verify : 1
Front view of the 6016 drive enclosure WWN 20000050cc00be51

  |-----|
  |  *   | *   | *   |   |
  |-----|
  |-----|
  |-----|
  |-----|
  |-----|

The '*' indicates the correct positions of the disk(s) for channel A.
Do the fault LEDs on the enclosure match this diagram? (y/N) :
    
```

4. When prompted to check on fault LEDs, visually verify the amber LED's on the enclosure and make sure the correct channel is lit up. Then confirm by typing "y<Enter>". Similarly, verify the correct positions of the disks for the other channels on this enclosure.

```

Do the fault LEDs on the enclosure match this diagram? (y/N) : y
Channel A has been visually verified.
Front view of the 6016 drive enclosure WWN 20000050cc00be51
-----
| | | | |
|-----|
| * | * | * | |
|-----|
| | | | |
|-----|
| | | | |
|-----|

The '*' indicates the correct positions of the disk(s) for channel B.
Do the fault LEDs on the enclosure match this diagram? (y/N) : y
Channel B has been visually verified.
Front view of the 6016 drive enclosure WWN 20000050cc00be51
-----
| | | | |
|-----|
| * | * | * | |
|-----|
| | | | |
|-----|
| | | | |
|-----|

The '*' indicates the correct positions of the disk(s) for channel C.
Do the fault LEDs on the enclosure match this diagram? (y/N) : y
Channel C has been visually verified.
Front view of the 6016 drive enclosure WWN 20000050cc00be51
-----
| | | | |
|-----|
| | | | |
|-----|
| * | * | * | |
|-----|
| | | | |
|-----|
| | | | |
|-----|

The '*' indicates the correct positions of the disk(s) for channel D.
Do the fault LEDs on the enclosure match this diagram? (y/N) : y
Channel D has been visually verified.
Front view of the 6016 drive enclosure WWN 20000050cc00be51
-----
| | | | * |
|-----|
| | | | * |
|-----|
| | | | * |
|-----|
| | | | |
|-----|

The '*' indicates the correct positions of the disk(s) for channel E.
Do the fault LEDs on the enclosure match this diagram? (y/N) : y
Channel E has been visually verified.

6016 WWN 20000050cc00be51 has been added to the verified list.
ses_F 8-24 16:42:19 Unverified 6016 drive enclosure detected :
EncID:20000050CC00BE9F: channel F
ses_faul 8-24 16:42:21 Unable to clear fault LED for drive 1F
ses_faul 8-24 16:42:21 Unable to clear fault LED for drive 1G
.
.
ses_faul 8-24 16:42:25 Unable to clear fault LED for drive 3P
ses_faul 8-24 16:42:25 Unable to clear fault LED for drive 3S

```

- To verify connections for the second enclosure, enter command: `ses verify_6016<Enter>`

```

RM660 [1]: ses verify_6016

The connections for an attached 6016 drive enclosure must be
verified before the SCSI Enclosure Services (SES) can accurately
indicate a drive failure within that enclosure.

Select the 6016 drive enclosure to verify :

Enclosure WWN      Tiers      Channels      Verified  Errors
-----
1 - 20000050cc00be51  1, 2, 3  A,B,C,D,E,.   Yes       No
2 - 20000050cc00be9f  1, 2, 3  F,,,,,,,,.    No        No

e - Exit this command

Enter the number of the 6016 enclosure to verify :
    
```

- Type "2<Enter>" when prompted to select an enclosure to verify.

```

Enter the number of the 6016 enclosure to verify : 2

Front view of the 6016 drive enclosure WWN 20000050cc00be9f

  |-----|
  |  *   | *   | *   |
  |-----|
  |-----|
  |-----|
  |-----|
  |-----|

The '*' indicates the correct positions of the disk(s) for channel F.
Do the fault LEDs on the enclosure match this diagram? (y/N) :
    
```

- When prompted to check on fault LEDs, visually verify the amber LED's on the enclosure and make sure the correct channel is lit up. Then confirm by typing "y<Enter>". Similarly, verify the correct positions of the disks for the other channels on this enclosure.

```

Do the fault LEDs on the enclosure match this diagram? (y/N) : y
Channel F has been visually verified.
Front view of the 6016 drive enclosure WWN 20000050cc00be9f

|-----|
| * | * | * | |
|-----|
|-----|
|-----|
|-----|

The '*' indicates the correct positions of the disk(s) for channel G.
Do the fault LEDs on the enclosure match this diagram? (y/N) : y
Channel G has been visually verified.
Front view of the 6016 drive enclosure WWN 20000050cc00be9f

|-----|
|-----|
| * | * | * | |
|-----|
|-----|
|-----|

The '*' indicates the correct positions of the disk(s) for channel H.
Do the fault LEDs on the enclosure match this diagram? (y/N) : y
Channel H has been visually verified.
Front view of the 6016 drive enclosure WWN 20000050cc00be9f

|-----|
|-----|
| * | * | * | |
|-----|
|-----|
|-----|

The '*' indicates the correct positions of the disk(s) for channel P.
Do the fault LEDs on the enclosure match this diagram? (y/N) : y
Channel P has been visually verified.
Front view of the 6016 drive enclosure WWN 20000050cc00be9f

|-----|
|-----|
| * | * | * | |
|-----|
|-----|
|-----|

The '*' indicates the correct positions of the disk(s) for channel S.
Do the fault LEDs on the enclosure match this diagram? (y/N) : y
Channel S has been visually verified.

6016 WWN 20000050cc00be9f has been added to the verified list.

```

8. Verify the connections for the enclosures again, enter command: `ses verify_6016<Enter>`. Check that "Yes" is displayed in the "Verified" column for both enclosures.
9. If you only have single RM controller unit, this SES verification procedure is complete. If you have dual RM controller configuration, proceed to the next step.

10. Login to Unit 2. Make sure that no host activity is in progress. Then fail Unit 1 by typing: `dual fail`<Enter>

```
RM660 [2]: dual fail

                        Dual RM660 Configuration
                        Unit 1                               Unit 2
-----
Label                  RM660 8000[1]                       RM660 8000[2]
Status                 Failed                               Healthy

Dual communication:    established.
Ethernet communication: established.
Cache coherency:       not enabled.
Cache coherency timeout: 2
```

11. Enter command: `ses on`<Enter>

```
RM660 [2]: ses on
Starting the SES monitors...
Done.

ses_A 8-24 16:43:22 Unverified 6016 drive enclosure detected :
EncID:20000050CC00BE51: channel A
ses_F 8-24 16:43:22 Unverified 6016 drive enclosure detected :
EncID:20000050CC00BE9F: channel F

ses_faul 8-24 16:43:23 Unable to clear fault LED for drive 1A
ses_faul 8-24 16:43:23 Unable to clear fault LED for drive 1B
:
ses_faul 8-24 16:43:26 Unable to clear fault LED for drive 3P
ses_faul 8-24 16:43:26 Unable to clear fault LED for drive 3S
```

12. Enter command: `ses verify_6016`<Enter>

```
RM660 [2]: ses verify_6016
The connections for an attached 6016 drive enclosure must be
verified before the SCSI Enclosure Services (SES) can accurately
indicate a drive failure within that enclosure.
Select the 6016 drive enclosure to verify :

Enclosure WWN      Tiers      Channels      Verified  Errors
-----
1 - 20000050cc00be51  1, 2, 3  A,.....  No       No
2 - 20000050cc00be9f  1, 2, 3  F,.....  No       No
e - Exit this command

Enter the number of the 6016 enclosure to verify :
```

13. Type “1<Enter>” when prompted to select an enclosure to verify.

```

Enter the number of the 6016 enclosure to verify : 1
Front view of the 6016 drive enclosure WWN 20000050cc00be51

|-----|
| * | * | * | |
|-----|
|-----|
|-----|
|-----|

The '*' indicates the correct positions of the disk(s) for channel A.
Do the fault LEDs on the enclosure match this diagram? (y/N) :
    
```

- When prompted to check on fault LEDs, visually verify the amber LED's on the enclosure and make sure the correct channel is lit up. Then confirm by typing "y<Enter>". Similarly, verify the correct positions of the disks for the other channels on this enclosure.

```

Do the fault LEDs on the enclosure match this diagram? (y/N) : y
Channel A has been visually verified.
Front view of the 6016 drive enclosure WWN 20000050cc00be51

|-----|
| * | * | * | |
|-----|
|-----|
|-----|
|-----|

The '*' indicates the correct positions of the disk(s) for channel B.
Do the fault LEDs on the enclosure match this diagram? (y/N) : y
Channel B has been visually verified.
Front view of the 6016 drive enclosure WWN 20000050cc00be51

|-----|
|-----|
| * | * | * | |
|-----|
|-----|
|-----|

The '*' indicates the correct positions of the disk(s) for channel C.
Do the fault LEDs on the enclosure match this diagram? (y/N) : y
Channel C has been visually verified.
Front view of the 6016 drive enclosure WWN 20000050cc00be51

|-----|
|-----|
|-----|
| * | * | * | |
|-----|
|-----|

The '*' indicates the correct positions of the disk(s) for channel D.
Do the fault LEDs on the enclosure match this diagram? (y/N) : y
    
```

```

Channel D has been visually verified.
Front view of the 6016 drive enclosure WWN 20000050cc00be51
-----
|         |         |         |         |
|         |         |         |         |
|         |         |         |         |
|         |         |         |         |
|         |         |         |         |
-----

The '*' indicates the correct positions of the disk(s) for channel E.
Do the fault LEDs on the enclosure match this diagram? (y/N) : y
Channel E has been visually verified.

6016 WWN 20000050cc00be51 has been added to the verified list.
ses_F 8-24 16:44:19 Unverified 6016 drive enclosure detected :
EncID:20000050CC00BE9F: channel F
ses_faul 8-24 16:44:21 Unable to clear fault LED for drive 1F
ses_faul 8-24 16:44:21 Unable to clear fault LED for drive 1G
.
.
ses_faul 8-24 16:44:25 Unable to clear fault LED for drive 3P
ses_faul 8-24 16:44:25 Unable to clear fault LED for drive 3S
    
```

- To verify connections for the second enclosure, enter command: `ses verify_6016`<Enter>

```

RM660 [2]: ses verify_6016

The connections for an attached 6016 drive enclosure must be
verified before the SCSI Enclosure Services (SES) can accurately
indicate a drive failure within that enclosure.

Select the 6016 drive enclosure to verify :

Enclosure WWN      Tiers      Channels      Verified  Errors
-----
1 - 20000050cc00be51  1, 2, 3  A,B,C,D,E,.  Yes      No
2 - 20000050cc00be9f  1, 2, 3  F,,,,,,,,,   No      No

e - Exit this command

Enter the number of the 6016 enclosure to verify :
    
```

- Type "2<Enter>" when prompted to select an enclosure to verify.

```

Enter the number of the 6016 enclosure to verify : 2
Front view of the 6016 drive enclosure WWN 20000050cc00be9f

|-----|
| * | * | * |
|-----|
|-----|
|-----|
|-----|

The '*' indicates the correct positions of the disk(s) for channel F.
Do the fault LEDs on the enclosure match this diagram? (y/N) :

```

- When prompted to check on fault LEDs, visually verify the amber LED's on the enclosure and make sure the correct channel is lit up. Then confirm by typing "y<Enter>". Similarly, verify the correct positions of the disks for the other channels on this enclosure.

```

Do the fault LEDs on the enclosure match this diagram? (y/N) : y
Channel F has been visually verified.
Front view of the 6016 drive enclosure WWN 20000050cc00be9f

|-----|
| * | * | * |
|-----|
|-----|
|-----|

The '*' indicates the correct positions of the disk(s) for channel G.
Do the fault LEDs on the enclosure match this diagram? (y/N) : y
Channel G has been visually verified.
Front view of the 6016 drive enclosure WWN 20000050cc00be9f

|-----|
| * | * | * |
|-----|
|-----|

The '*' indicates the correct positions of the disk(s) for channel H.
Do the fault LEDs on the enclosure match this diagram? (y/N) : y

```

```

Channel H has been visually verified.
Front view of the 6016 drive enclosure WWN 20000050cc00be9f
-----
| | | | |
| | | | |
| | | | |
| * | * | * |
| | | | |
-----

The '*' indicates the correct positions of the disk(s) for channel P.
Do the fault LEDs on the enclosure match this diagram? (y/N) : y
Channel P has been visually verified.
Front view of the 6016 drive enclosure WWN 20000050cc00be9f
-----
| | | | * |
| | | | * |
| | | | * |
| | | | |
| | | | |
-----

The '*' indicates the correct positions of the disk(s) for channel S.
Do the fault LEDs on the enclosure match this diagram? (y/N) : y
Channel S has been visually verified.

6016 WWN 20000050cc00be9f has been added to the verified list.
    
```

18. Verify the connections for the enclosures again, enter command: `ses verify_6016<Enter>`. Check that “Yes” is displayed in the “Verified” column for both enclosures.
19. Enter command: `dual heal<Enter>` to restore Unit 1. Check that both units are “Healthy”, and “Dual” and “Ethernet” communications are established.

```

RM660 [2]: dual heal

                        Dual RM6608000/8500 Configuration
                        Unit 1                               Unit 2
-----
Label          RM660 8000[1]                               RM660 8000[2]
Status         Healthy                                     Healthy

Dual communication:      established.
Ethernet communication:  established.
Cache coherency:        not enabled.
Cache coherency timeout: 100
    
```

20. The SES verification procedure for dual RM controller configuration is complete.

Enclosure AL_PA and Tier Mapping Charts

Follow the instructions given in “Setting Tier Mapping Mode” on page 32 to select the correct Tier Mapping Mode for your enclosures.

The SFx016 enclosure supports 15 address ranges (1 to 15). The thumb wheel switch on the Ops panel is used to select the SEL_ID base address.

Table C-2 and Table C-3 show the mapping charts for SF4016 and SF2016 enclosures respectively for the RM660. Table C-4 and Table C-5 show the mapping charts for SF4016 and SF2016 enclosures respectively for the RM610. Each table cell represents a drive slot on the enclosure (when viewing from front) and displays its AL_PA and disk number. Up to fifteen SF4016 enclosures (or up to seven SF2016 enclosures) can be daisy-chained to one channel on the RM660 and each enclosure must have a unique enclosure ID (1..15). ID 1 should always be used for the first set of enclosures, ID 2 for the second set, ID 3 for the third set, and so on.




Table C-2 SF4016 Enclosure AL_PA and Tier Mapping Chart for RM660

Channels A and B				Channels A and B					
ID=1	EF/Disk 1A	E8/Disk 2A	E4/Disk 3A	E2/Disk 4A	ID=9	72/Disk 65A	71/Disk 66A	6E/Disk 67A	6D/Disk 68A
	E1/Disk 5A	E0/Disk 6A	DC/Disk 7A	DA/Disk 8A		6C/Disk 69A	6B/Disk 70A	6A/Disk 71A	69/Disk 72A
	DA/Disk 8B	DC/Disk 7B	E0/Disk 6B	E1/Disk 5B		69/Disk 72B	6A/Disk 71B	6B/Disk 70B	6C/Disk 69B
	E2/Disk 4B	E4/Disk 3B	E8/Disk 2B	EF/Disk 1B		6D/Disk 68B	6E/Disk 67B	71/Disk 66B	72/Disk 65B
ID=2	D9/Disk 9A	D6/Disk 10A	D5/Disk 11A	D4/Disk 12A	ID=10	67/Disk 73A	66/Disk 74A	65/Disk 75A	63/Disk 76A
	D3/Disk 13A	D2/Disk 14A	D1/Disk 15A	CE/Disk 16A		5C/Disk 77A	5A/Disk 78A	59/Disk 79A	56/Disk 80A
	CE/Disk 16B	D1/Disk 15B	D2/Disk 14B	D3/Disk 13B		56/Disk 80B	59/Disk 79B	5A/Disk 78B	5C/Disk 77B
	D4/Disk 12B	D5/Disk 11B	D6/Disk 10B	D9/Disk 9B		63/Disk 76B	65/Disk 75B	66/Disk 74B	67/Disk 73B
ID=3	CD/Disk 17A	CC/Disk 18A	CB/Disk 19A	CA/Disk 20A	ID=11	55/Disk 81A	54/Disk 82A	53/Disk 83A	52/Disk 84A
	C9/Disk 21A	C7/Disk 22A	C6/Disk 23A	C5/Disk 24A		51/Disk 85A	4E/Disk 86A	4D/Disk 87A	4C/Disk 88A
	C5/Disk 24B	C6/Disk 23B	C7/Disk 22B	C9/Disk 21B		4C/Disk 88B	4D/Disk 87B	4E/Disk 86B	51/Disk 85B
	CA/Disk 20B	CB/Disk 19B	CC/Disk 18B	CD/Disk 17B		52/Disk 84B	53/Disk 83B	54/Disk 82B	55/Disk 81B
ID=4	C3/Disk 25A	BC/Disk 26A	BA/Disk 27A	B9/Disk 28A	ID=12	4B/Disk 89A	4A/Disk 90A	49/Disk 91A	47/Disk 92A
	B6/Disk 29A	B5/Disk 30A	B4/Disk 31A	B3/Disk 32A		46/Disk 93A	45/Disk 94A	43/Disk 95A	3C/Disk 96A
	B3/Disk 32B	B4/Disk 31B	B5/Disk 30B	B6/Disk 29B		3C/Disk 96B	43/Disk 95B	45/Disk 94B	46/Disk 93B
	B9/Disk 28B	BA/Disk 27B	BC/Disk 26B	C3/Disk 25B		47/Disk 92B	49/Disk 91B	4A/Disk 90B	4B/Disk 89B
ID=5	B2/Disk 33A	B1/Disk 34A	AE/Disk 35A	AD/Disk 36A	ID=13	3A/Disk 97A	39/Disk 98A	36/Disk 99A	35/Disk 100A
	AC/Disk 37A	AB/Disk 38A	AA/Disk 39A	A9/Disk 40A		34/Disk 101A	33/Disk 102A	32/Disk 103A	31/Disk 104A
	A9/Disk 40B	AA/Disk 39B	AB/Disk 38B	AC/Disk 37B		31/Disk 104B	32/Disk 103B	33/Disk 102B	34/Disk 101B
	AD/Disk 36B	AE/Disk 35B	B1/Disk 34B	B2/Disk 33B		35/Disk 100B	36/Disk 99B	39/Disk 98B	3A/Disk 97B
ID=6	A7/Disk 41A	A6/Disk 42A	A5/Disk 43A	A3/Disk 44A	ID=14	2E/Disk 105A	2D/Disk 106A	2C/Disk 107A	2B/Disk 108A
	9F/Disk 45A	9E/Disk 46A	9D/Disk 47A	9B/Disk 48A		2A/Disk 109A	29/Disk 110A	27/Disk 111A	26/Disk 112A
	9B/Disk 48B	9D/Disk 47B	9E/Disk 46B	9F/Disk 45B		26/Disk 112B	27/Disk 111B	29/Disk 110B	2A/Disk 109B
	A3/Disk 44B	A5/Disk 43B	A6/Disk 42B	A7/Disk 41B		2B/Disk 108B	2C/Disk 107B	2D/Disk 106B	2E/Disk 105B
ID=7	98/Disk 49A	97/Disk 50A	90/Disk 51A	8F/Disk 52A	ID=15	25/Disk 113A	23/Disk 114A	1F/Disk 115A	1E/Disk 116A
	88/Disk 53A	84/Disk 54A	82/Disk 55A	81/Disk 56A		1D/Disk 117A	1B/Disk 118A	18/Disk 119A	17/Disk 120A
	81/Disk 56B	82/Disk 55B	84/Disk 54B	88/Disk 53B		17/Disk 120B	18/Disk 119B	1B/Disk 118B	1D/Disk 117B
	8F/Disk 52B	90/Disk 51B	97/Disk 50B	98/Disk 49B		1E/Disk 116B	1F/Disk 115B	23/Disk 114B	25/Disk 113B
ID=8	80/Disk 57A	7C/Disk 58A	7A/Disk 59A	79/Disk 60A					
	76/Disk 61A	75/Disk 62A	74/Disk 63A	73/Disk 64A					

= SES drives

Table C-2 SF4016 Enclosure AL_PA and Tier Mapping Chart for RM660 (continued)

	73/Disk 64B	74/Disk 63B	75/Disk 62B	76/Disk 61B					
	79/Disk 60B	7A/Disk 59B	7C/Disk 58B	80/Disk 57B					
	Channels C and D					Channels C and D			
ID=1	EF/Disk 1C	E8/Disk 2C	E4/Disk 3C	E2/Disk 4C	ID=9	72/Disk 65C	71/Disk 66C	6E/Disk 67C	6D/Disk 68C
	E1/Disk 5C	E0/Disk 6C	DC/Disk 7C	DA/Disk 8C		6C/Disk 69C	6B/Disk 70C	6A/Disk 71C	69/Disk 72C
	DA/Disk 8D	DC/Disk 7D	E0/Disk 6D	E1/Disk 5D		69/Disk 72D	6A/Disk 71D	6B/Disk 70D	6C/Disk 69D
	E2/Disk 4D	E4/Disk 3D	E8/Disk 2D	EF/Disk 1D		6D/Disk 68D	6E/Disk 67D	71/Disk 66D	72/Disk 65D
ID=2	D9/Disk 9C	D6/Disk 10C	D5/Disk 11C	D4/Disk 12C	ID=10	67/Disk 73C	66/Disk 74C	65/Disk 75C	63/Disk 76C
	D3/Disk 13C	D2/Disk 14C	D1/Disk 15C	CE/Disk 16C		5C/Disk 77C	5A/Disk 78C	59/Disk 79C	56/Disk 80C
	CE/Disk 16D	D1/Disk 15D	D2/Disk 14D	D3/Disk 13D		56/Disk 80D	59/Disk 79D	5A/Disk 78D	5C/Disk 77D
	D4/Disk 12D	D5/Disk 11D	D6/Disk 10D	D9/Disk 9D		63/Disk 76D	65/Disk 75D	66/Disk 74D	67/Disk 73D
ID=3	CD/Disk 17C	CC/Disk 18C	CB/Disk 19C	CA/Disk 20C	ID=11	55/Disk 81C	54/Disk 82C	53/Disk 83C	52/Disk 84C
	C9/Disk 21C	C7/Disk 22C	C6/Disk 23C	C5/Disk 24C		51/Disk 85C	4E/Disk 86C	4D/Disk 87C	4C/Disk 88C
	C5/Disk 24D	C6/Disk 23D	C7/Disk 22D	C9/Disk 21D		4C/Disk 88D	4D/Disk 87D	4E/Disk 86D	51/Disk 85D
	CA/Disk 20D	CB/Disk 19D	CC/Disk 18D	CD/Disk 17D		52/Disk 84D	53/Disk 83D	54/Disk 82D	55/Disk 81D
ID=4	C3/Disk 25C	BC/Disk 26C	BA/Disk 27C	B9/Disk 28C	ID=12	4B/Disk 89C	4A/Disk 90C	49/Disk 91C	47/Disk 92C
	B6/Disk 29C	B5/Disk 30C	B4/Disk 31C	B3/Disk 32C		46/Disk 93C	45/Disk 94C	43/Disk 95C	3C/Disk 96C
	B3/Disk 32D	B4/Disk 31D	B5/Disk 30D	B6/Disk 29D		3C/Disk 96D	43/Disk 95D	45/Disk 94D	46/Disk 93D
	B9/Disk 28D	BA/Disk 27D	BC/Disk 26D	C3/Disk 25D		47/Disk 92D	49/Disk 91D	4A/Disk 90D	4B/Disk 89D
ID=5	B2/Disk 33C	B1/Disk 34C	AE/Disk 35C	AD/Disk 36C	ID=13	3A/Disk 97C	39/Disk 98C	36/Disk 99C	35/Disk 100C
	AC/Disk 37C	AB/Disk 38C	AA/Disk 39C	A9/Disk 40C		34/Disk 101C	33/Disk 102C	32/Disk 103C	31/Disk 104C
	A9/Disk 40D	AA/Disk 39D	AB/Disk 38D	AC/Disk 37D		31/Disk 104D	32/Disk 103D	33/Disk 102D	34/Disk 101D
	AD/Disk 36D	AE/Disk 35D	B1/Disk 34D	B2/Disk 33D		35/Disk 100D	36/Disk 99D	39/Disk 98D	3A/Disk 97D
ID=6	A7/Disk 41C	A6/Disk 42C	A5/Disk 43C	A3/Disk 44C	ID=14	2E/Disk 105C	2D/Disk 106C	2C/Disk 107C	2B/Disk 108C
	9F/Disk 45C	9E/Disk 46C	9D/Disk 47C	9B/Disk 48C		2A/Disk 109C	29/Disk 110C	27/Disk 111C	26/Disk 112C
	9B/Disk 48D	9D/Disk 47D	9E/Disk 46D	9F/Disk 45D		26/Disk 112D	27/Disk 111D	29/Disk 110D	2A/Disk 109D
	A3/Disk 44D	A5/Disk 43D	A6/Disk 42D	A7/Disk 41D		2B/Disk 108D	2C/Disk 107D	2D/Disk 106D	2E/Disk 105D
ID=7	98/Disk 49C	97/Disk 50C	90/Disk 51C	8F/Disk 52C	ID=15	25/Disk 113C	23/Disk 114C	1F/Disk 115C	1E/Disk 116C
	88/Disk 53C	84/Disk 54C	82/Disk 55C	81/Disk 56C		1D/Disk 117C	1B/Disk 118C	18/Disk 119C	17/Disk 120C
	81/Disk 56D	82/Disk 55D	84/Disk 54D	88/Disk 53D		17/Disk 120D	18/Disk 119D	1B/Disk 118D	1D/Disk 117D
	8F/Disk 52D	90/Disk 51D	97/Disk 50D	98/Disk 49D		1E/Disk 116D	1F/Disk 115D	23/Disk 114D	25/Disk 113D

Table C-2 SF4016 Enclosure AL_PA and Tier Mapping Chart for RM660 (continued)

	81/Disk 56H	82/Disk 55H	84/Disk 54H	88/Disk 53H		17/Disk 120H	18/Disk 119H	1B/Disk 118H	1D/Disk 117H
	8F/Disk 52H	90/Disk 51H	97/Disk 50H	98/Disk 49H		1E/Disk 116H	1F/Disk 115H	23/Disk 114H	25/Disk 113H
ID=8	80/Disk 57G	7C/Disk 58G	7A/Disk 59G	79/Disk 60G					
	76/Disk 61G	75/Disk 62G	74/Disk 63G	73/Disk 64G					
	73/Disk 64H	74/Disk 63H	75/Disk 62H	76/Disk 61H					
	79/Disk 60H	7A/Disk 59H	7C/Disk 58H	80/Disk 57H					
	Channels P (Parity) and S (Spare)					Channels P (Parity) and S (Spare)			
ID=1	EF/Disk 1P	E8/Disk 2P	E4/Disk 3P	E2/Disk 4P	ID=9	72/Disk 65P	71/Disk 66P	6E/Disk 67P	6D/Disk 68P
	E1/Disk 5P	E0/Disk 6P	DC/Disk 7P	DA/Disk 8P		6C/Disk 69P	6B/Disk 70P	6A/Disk 71P	69/Disk 72P
	DA/Disk 8S	DC/Disk 7S	E0/Disk 6S	E1/Disk 5S		69/Disk 72S	6A/Disk 71S	6B/Disk 70S	6C/Disk 69S
	E2/Disk 4S	E4/Disk 3S	E8/Disk 2S	EF/Disk 1S		6D/Disk 68S	6E/Disk 67S	71/Disk 66S	72/Disk 65S
ID=2	D9/Disk 9P	D6/Disk 10P	D5/Disk 11P	D4/Disk 12P	ID=10	67/Disk 73P	66/Disk 74P	65/Disk 75P	63/Disk 76P
	D3/Disk 13P	D2/Disk 14P	D1/Disk 15P	CE/Disk 16P		5C/Disk 77P	5A/Disk 78P	59/Disk 79P	56/Disk 80P
	CE/Disk 16S	D1/Disk 15S	D2/Disk 14S	D3/Disk 13S		56/Disk 80S	59/Disk 79S	5A/Disk 78S	5C/Disk 77S
	D4/Disk 12S	D5/Disk 11S	D6/Disk 10S	D9/Disk 9S		63/Disk 76S	65/Disk 75S	66/Disk 74S	67/Disk 73S
ID=3	CD/Disk 17P	CC/Disk 18P	CB/Disk 19P	CA/Disk 20P	ID=11	55/Disk 81P	54/Disk 82P	53/Disk 83P	52/Disk 84P
	C9/Disk 21P	C7/Disk 22P	C6/Disk 23P	C5/Disk 24P		51/Disk 85P	4E/Disk 86P	4D/Disk 87P	4C/Disk 88P
	C5/Disk 24S	C6/Disk 23S	C7/Disk 22S	C9/Disk 21S		4C/Disk 88S	4D/Disk 87S	4E/Disk 86S	51/Disk 85S
	CA/Disk 20S	CB/Disk 19S	CC/Disk 18S	CD/Disk 17S		52/Disk 84S	53/Disk 83S	54/Disk 82S	55/Disk 81S
ID=4	C3/Disk 25P	BC/Disk 26P	BA/Disk 27P	B9/Disk 28P	ID=12	4B/Disk 89P	4A/Disk 90P	49/Disk 91P	47/Disk 92P
	B6/Disk 29P	B5/Disk 30P	B4/Disk 31P	B3/Disk 32P		46/Disk 93P	45/Disk 94P	43/Disk 95P	3C/Disk 96P
	B3/Disk 32S	B4/Disk 31S	B5/Disk 30S	B6/Disk 29S		3C/Disk 96S	43/Disk 95S	45/Disk 94S	46/Disk 93S
	B9/Disk 28S	BA/Disk 27S	BC/Disk 26S	C3/Disk 25S		47/Disk 92S	49/Disk 91S	4A/Disk 90S	4B/Disk 89S
ID=5	B2/Disk 33P	B1/Disk 34P	AE/Disk 35P	AD/Disk 36P	ID=13	3A/Disk 97P	39/Disk 98P	36/Disk 99P	35/Disk 100P
	AC/Disk 37P	AB/Disk 38P	AA/Disk 39P	A9/Disk 40P		34/Disk 101P	33/Disk 102P	32/Disk 103P	31/Disk 104P
	A9/Disk 40S	AA/Disk 39S	AB/Disk 38S	AC/Disk 37S		31/Disk 104S	32/Disk 103S	33/Disk 102S	34/Disk 101S
	AD/Disk 36S	AE/Disk 35S	B1/Disk 34S	B2/Disk 33S		35/Disk 100S	36/Disk 99S	39/Disk 98S	3A/Disk 97S
ID=6	A7/Disk 41P	A6/Disk 42P	A5/Disk 43P	A3/Disk 44P	ID=14	2E/Disk 105P	2D/Disk 106P	2C/Disk 107P	2B/Disk 108P
	9F/Disk 45P	9E/Disk 46P	9D/Disk 47P	9B/Disk 48P		2A/Disk 109P	29/Disk 110P	27/Disk 111P	26/Disk 112P
	9B/Disk 48S	9D/Disk 47S	9E/Disk 46S	9F/Disk 45S		26/Disk 112S	27/Disk 111S	29/Disk 110S	2A/Disk 109S
	A3/Disk 44S	A5/Disk 43S	A6/Disk 42S	A7/Disk 41S		2B/Disk 108S	2C/Disk 107S	2D/Disk 106S	2E/Disk 105S
ID=7	98/Disk 49P	97/Disk 50P	90/Disk 51P	8F/Disk 52P	ID=15	25/Disk 113P	23/Disk 114P	1F/Disk 115P	1E/Disk 116P
	88/Disk 53P	84/Disk 54P	82/Disk 55P	81/Disk 56P		1D/Disk 117P	1B/Disk 118P	18/Disk 119P	17/Disk 120P

 = SES drives

Table C-2 SF4016 Enclosure AL_PA and Tier Mapping Chart for RM660 (continued)

	81/Disk 56S	82/Disk 55S	84/Disk 54S	88/Disk 53S	17/Disk 120S	18/Disk 119S	1B/Disk 118S	1D/Disk 117S
	8F/Disk 52S	90/Disk 51S	97/Disk 50S	98/Disk 49S	1E/Disk 116S	1F/Disk 115S	23/Disk 114S	25/Disk 113S
ID=8	80/Disk 57P	7C/Disk 58P	7A/Disk 59P	79/Disk 60P				
	76/Disk 61P	75/Disk 62P	74/Disk 63P	73/Disk 64P				
	73/Disk 64S	74/Disk 63S	75/Disk 62S	76/Disk 61S				
	79/Disk 60S	7A/Disk 59S	7C/Disk 58S	80/Disk 57S				

 = SES drives

Table C-3 SF2016 Enclosure AL_PA and Tier Mapping Chart for RM660

All Channels

ID=1	E4/Disk 1	E2/Disk 2	E1/Disk 3	E0/Disk 4
	DC/Disk 5	DA/Disk 6	D9/Disk 7	D6/Disk 8
	D5/Disk 9	D4/Disk 10	D3/Disk 11	D2/Disk 12
	D1/Disk 13	CE/Disk 14	CD/Disk 15	CC/Disk 16
ID=2	CB/Disk 17	CA/Disk 18	C9/Disk 19	C7/Disk 20
	C6/Disk 21	C5/Disk 22	C3/Disk 23	BC/Disk 24
	BA/Disk 25	B9/Disk 26	B6/Disk 27	B5/Disk 28
	B4/Disk 29	B3/Disk 30	B2/Disk 31	B1/Disk 32
ID=3	AE/Disk 33	AD/Disk 34	AC/Disk 35	AB/Disk 36
	AA/Disk 37	A9/Disk 38	A7/Disk 39	A6/Disk 40
	A5/Disk 41	A3/Disk 42	9F/Disk 43	9E/Disk 44
	9D/Disk 45	9B/Disk 46	98/Disk 47	97/Disk 48
ID=4	90/Disk 49	8F/Disk 50	88/Disk 51	84/Disk 52
	82/Disk 53	81/Disk 54	80/Disk 55	7C/Disk 56
	7A/Disk 57	79/Disk 58	76/Disk 59	75/Disk 60
	74/Disk 61	73/Disk 62	72/Disk 63	71/Disk 64
ID=5	6E/Disk 65	6D/Disk 66	6C/Disk 67	6B/Disk 68
	6A/Disk 69	69/Disk 70	67/Disk 71	66/Disk 72
	65/Disk 73	63/Disk 74	5C/Disk 75	5A/Disk 76
	59/Disk 77	56/Disk 78	55/Disk 79	54/Disk 80
ID=6	53/Disk 81	52/Disk 82	51/Disk 83	4E/Disk 84
	4D/Disk 85	4C/Disk 86	4B/Disk 87	4A/Disk 88
	49/Disk 89	47/Disk 90	46/Disk 91	45/Disk 92
	43/Disk 93	3C/Disk 94	3A/Disk 95	39/Disk 96
ID=7	36/Disk 97	35/Disk 98	34/Disk 99	33/Disk 100
	32/Disk 101	31/Disk 102	2E/Disk 103	2D/Disk 104
	2C/Disk 105	2B/Disk 106	2A/Disk 107	29/Disk 108
	27/Disk 109	26/Disk 110	25/Disk 111	23/Disk 112

 = SES drives

Table C-4 SF4016 Enclosure AL_PA and Tier Mapping Charts for RM610
Channels A and B

Channels A and B				Channels A and B					
ID=1	EF/Disk 1A	E8/Disk 2A	E4/Disk 3A	E2/Disk 4A	ID=9	72/Disk 65A	71/Disk 66A	6E/Disk 67A	6D/Disk 68A
	E1/Disk 5A	E0/Disk 6A	DC/Disk 7A	DA/Disk 8A		6C/Disk 69A	6B/Disk 70A	6A/Disk 71A	69/Disk 72A
	DA/Disk 8B	DC/Disk 7B	E0/Disk 6B	E1/Disk 5B		69/Disk 72B	6A/Disk 71B	6B/Disk 70B	6C/Disk 69B
	E2/Disk 4B	E4/Disk 3B	E8/Disk 2B	EF/Disk 1B		6D/Disk 68B	6E/Disk 67B	71/Disk 66B	72/Disk 65B
ID=2	D9/Disk 9A	D6/Disk 10A	D5/Disk 11A	D4/Disk 12A	ID=10	67/Disk 73A	66/Disk 74A	65/Disk 75A	63/Disk 76A
	D3/Disk 13A	D2/Disk 14A	D1/Disk 15A	CE/Disk 16A		5C/Disk 77A	5A/Disk 78A	59/Disk 79A	56/Disk 80A
	CE/Disk 16B	D1/Disk 15B	D2/Disk 14B	D3/Disk 13B		56/Disk 80B	59/Disk 79B	5A/Disk 78B	5C/Disk 77B
	D4/Disk 12B	D5/Disk 11B	D6/Disk 10B	D9/Disk 9B		63/Disk 76B	65/Disk 75B	66/Disk 74B	67/Disk 73B
ID=3	CD/Disk 17A	CC/Disk 18A	CB/Disk 19A	CA/Disk 20A	ID=11	55/Disk 81A	54/Disk 82A	53/Disk 83A	52/Disk 84A
	C9/Disk 21A	C7/Disk 22A	C6/Disk 23A	C5/Disk 24A		51/Disk 85A	4E/Disk 86A	4D/Disk 87A	4C/Disk 88A
	C5/Disk 24B	C6/Disk 23B	C7/Disk 22B	C9/Disk 21B		4C/Disk 88B	4D/Disk 87B	4E/Disk 86B	51/Disk 85B
	CA/Disk 20B	CB/Disk 19B	CC/Disk 18B	CD/Disk 17B		52/Disk 84B	53/Disk 83B	54/Disk 82B	55/Disk 81B
ID=4	C3/Disk 25A	BC/Disk 26A	BA/Disk 27A	B9/Disk 28A	ID=12	4B/Disk 89A	4A/Disk 90A	49/Disk 91A	47/Disk 92A
	B6/Disk 29A	B5/Disk 30A	B4/Disk 31A	B3/Disk 32A		46/Disk 93A	45/Disk 94A	43/Disk 95A	3C/Disk 96A
	B3/Disk 32B	B4/Disk 31B	B5/Disk 30B	B6/Disk 29B		3C/Disk 96B	43/Disk 95B	45/Disk 94B	46/Disk 93B
	B9/Disk 28B	BA/Disk 27B	BC/Disk 26B	C3/Disk 25B		47/Disk 92B	49/Disk 91B	4A/Disk 90B	4B/Disk 89B
ID=5	B2/Disk 33A	B1/Disk 34A	AE/Disk 35A	AD/Disk 36A	ID=13	3A/Disk 97A	39/Disk 98A	36/Disk 99A	35/Disk 100A
	AC/Disk 37A	AB/Disk 38A	AA/Disk 39A	A9/Disk 40A		34/Disk 101A	33/Disk 102A	32/Disk 103A	31/Disk 104A
	A9/Disk 40B	AA/Disk 39B	AB/Disk 38B	AC/Disk 37B		31/Disk 104B	32/Disk 103B	33/Disk 102B	34/Disk 101B
	AD/Disk 36B	AE/Disk 35B	B1/Disk 34B	B2/Disk 33B		35/Disk 100B	36/Disk 99B	39/Disk 98B	3A/Disk 97B
ID=6	A7/Disk 41A	A6/Disk 42A	A5/Disk 43A	A3/Disk 44A	ID=14	2E/Disk 105A	2D/Disk 106A	2C/Disk 107A	2B/Disk 108A
	9F/Disk 45A	9E/Disk 46A	9D/Disk 47A	9B/Disk 48A		2A/Disk 109A	29/Disk 110A	27/Disk 111A	26/Disk 112A
	9B/Disk 48B	9D/Disk 47B	9E/Disk 46B	9F/Disk 45B		26/Disk 112B	27/Disk 111B	29/Disk 110B	2A/Disk 109B
	A3/Disk 44B	A5/Disk 43B	A6/Disk 42B	A7/Disk 41B		2B/Disk 108B	2C/Disk 107B	2D/Disk 106B	2E/Disk 105B
ID=7	98/Disk 49A	97/Disk 50A	90/Disk 51A	8F/Disk 52A	ID=15	25/Disk 113A	23/Disk 114A	1F/Disk 115A	1E/Disk 116A
	88/Disk 53A	84/Disk 54A	82/Disk 55A	81/Disk 56A		1D/Disk 117A	1B/Disk 118A	18/Disk 119A	17/Disk 120A
	81/Disk 56B	82/Disk 55B	84/Disk 54B	88/Disk 53B		17/Disk 120B	18/Disk 119B	1B/Disk 118B	1D/Disk 117B
	8F/Disk 52B	90/Disk 51B	97/Disk 50B	98/Disk 49B		1E/Disk 116B	1F/Disk 115B	23/Disk 114B	25/Disk 113B
ID=8	80/Disk 57A	7C/Disk 58A	7A/Disk 59A	79/Disk 60A					
	76/Disk 61A	75/Disk 62A	74/Disk 63A	73/Disk 64A					

= SES drives

Table C-4 SF4016 Enclosure AL_PA and Tier Mapping Charts for RM610 (continued)

	73/Disk 64B	74/Disk 63B	75/Disk 62B	76/Disk 61B					
	79/Disk 60B	7A/Disk 59B	7C/Disk 58B	80/Disk 57B					
	Channels C and D					Channels C and D			
ID=1	EF/Disk 1C	E8/Disk 2C	E4/Disk 3C	E2/Disk 4C	ID=9	72/Disk 65C	71/Disk 66C	6E/Disk 67C	6D/Disk 68C
	E1/Disk 5C	E0/Disk 6C	DC/Disk 7C	DA/Disk 8C		6C/Disk 69C	6B/Disk 70C	6A/Disk 71C	69/Disk 72C
	DA/Disk 8D	DC/Disk 7D	E0/Disk 6D	E1/Disk 5D		69/Disk 72D	6A/Disk 71D	6B/Disk 70D	6C/Disk 69D
	E2/Disk 4D	E4/Disk 3D	E8/Disk 2D	EF/Disk 1D		6D/Disk 68D	6E/Disk 67D	71/Disk 66D	72/Disk 65D
ID=2	D9/Disk 9C	D6/Disk 10C	D5/Disk 11C	D4/Disk 12C	ID=10	67/Disk 73C	66/Disk 74C	65/Disk 75C	63/Disk 76C
	D3/Disk 13C	D2/Disk 14C	D1/Disk 15C	CE/Disk 16C		5C/Disk 77C	5A/Disk 78C	59/Disk 79C	56/Disk 80C
	CE/Disk 16D	D1/Disk 15D	D2/Disk 14D	D3/Disk 13D		56/Disk 80D	59/Disk 79D	5A/Disk 78D	5C/Disk 77D
	D4/Disk 12D	D5/Disk 11D	D6/Disk 10D	D9/Disk 9D		63/Disk 76D	65/Disk 75D	66/Disk 74D	67/Disk 73D
ID=3	CD/Disk 17C	CC/Disk 18C	CB/Disk 19C	CA/Disk 20C	ID=11	55/Disk 81C	54/Disk 82C	53/Disk 83C	52/Disk 84C
	C9/Disk 21C	C7/Disk 22C	C6/Disk 23C	C5/Disk 24C		51/Disk 85C	4E/Disk 86C	4D/Disk 87C	4C/Disk 88C
	C5/Disk 24D	C6/Disk 23D	C7/Disk 22D	C9/Disk 21D		4C/Disk 88D	4D/Disk 87D	4E/Disk 86D	51/Disk 85D
	CA/Disk 20D	CB/Disk 19D	CC/Disk 18D	CD/Disk 17D		52/Disk 84D	53/Disk 83D	54/Disk 82D	55/Disk 81D
ID=4	C3/Disk 25C	BC/Disk 26C	BA/Disk 27C	B9/Disk 28C	ID=12	4B/Disk 89C	4A/Disk 90C	49/Disk 91C	47/Disk 92C
	B6/Disk 29C	B5/Disk 30C	B4/Disk 31C	B3/Disk 32C		46/Disk 93C	45/Disk 94C	43/Disk 95C	3C/Disk 96C
	B3/Disk 32D	B4/Disk 31D	B5/Disk 30D	B6/Disk 29D		3C/Disk 96D	43/Disk 95D	45/Disk 94D	46/Disk 93D
	B9/Disk 28D	BA/Disk 27D	BC/Disk 26D	C3/Disk 25D		47/Disk 92D	49/Disk 91D	4A/Disk 90D	4B/Disk 89D
ID=5	B2/Disk 33C	B1/Disk 34C	AE/Disk 35C	AD/Disk 36C	ID=13	3A/Disk 97C	39/Disk 98C	36/Disk 99C	35/Disk 100C
	AC/Disk 37C	AB/Disk 38C	AA/Disk 39C	A9/Disk 40C		34/Disk 101C	33/Disk 102C	32/Disk 103C	31/Disk 104C
	A9/Disk 40D	AA/Disk 39D	AB/Disk 38D	AC/Disk 37D		31/Disk 104D	32/Disk 103D	33/Disk 102D	34/Disk 101D
	AD/Disk 36D	AE/Disk 35D	B1/Disk 34D	B2/Disk 33D		35/Disk 100D	36/Disk 99D	39/Disk 98D	3A/Disk 97D
ID=6	A7/Disk 41C	A6/Disk 42C	A5/Disk 43C	A3/Disk 44C	ID=14	2E/Disk 105C	2D/Disk 106C	2C/Disk 107C	2B/Disk 108C
	9F/Disk 45C	9E/Disk 46C	9D/Disk 47C	9B/Disk 48C		2A/Disk 109C	29/Disk 110C	27/Disk 111C	26/Disk 112C
	9B/Disk 48D	9D/Disk 47D	9E/Disk 46D	9F/Disk 45D		26/Disk 112D	27/Disk 111D	29/Disk 110D	2A/Disk 109D
	A3/Disk 44D	A5/Disk 43D	A6/Disk 42D	A7/Disk 41D		2B/Disk 108D	2C/Disk 107D	2D/Disk 106D	2E/Disk 105D
ID=7	98/Disk 49C	97/Disk 50C	90/Disk 51C	8F/Disk 52C	ID=15	25/Disk 113C	23/Disk 114C	1F/Disk 115C	1E/Disk 116C
	88/Disk 53C	84/Disk 54C	82/Disk 55C	81/Disk 56C		1D/Disk 117C	1B/Disk 118C	18/Disk 119C	17/Disk 120C
	81/Disk 56D	82/Disk 55D	84/Disk 54D	88/Disk 53D		17/Disk 120D	18/Disk 119D	1B/Disk 118D	1D/Disk 117D
	8F/Disk 52D	90/Disk 51D	97/Disk 50D	98/Disk 49D		1E/Disk 116D	1F/Disk 115D	23/Disk 114D	25/Disk 113D

Table C-4 SF4016 Enclosure AL_PA and Tier Mapping Charts for RM610 **(continued)**

ID=8	80/Disk57P	7C/Disk58P	7A/Disk59P	79/Disk60P	= SES drives
	76/Disk61P	75/Disk62P	74/Disk63P	73/Disk64P	
	73/Disk64S	74/Disk63S	75/Disk62S	76/Disk61S	
	79/Disk60S	7A/Disk59S	7C/Disk58S	80/Disk57S	

Table C-5 SF2016 Enclosure AL_PA and Tier Mapping Charts for RM 610

All Channels

ID=1	E4/Disk 1	E2/Disk 2	E1/Disk 3	E0/Disk 4
	DC/Disk 5	DA/Disk 6	D9/Disk 7	D6/Disk 8
	D5/Disk 9	D4/Disk 10	D3/Disk 11	D2/Disk 12
	D1/Disk 13	CE/Disk 14	CD/Disk 15	CC/Disk 16
ID=2	CB/Disk 17	CA/Disk 18	C9/Disk 19	C7/Disk 20
	C6/Disk 21	C5/Disk 22	C3/Disk 23	BC/Disk 24
	BA/Disk 25	B9/Disk 26	B6/Disk 27	B5/Disk 28
	B4/Disk 29	B3/Disk 30	B2/Disk 31	B1/Disk 32
ID=3	AE/Disk 33	AD/Disk 34	AC/Disk 35	AB/Disk 36
	AA/Disk 37	A9/Disk 38	A7/Disk 39	A6/Disk 40
	A5/Disk 41	A3/Disk 42	9F/Disk 43	9E/Disk 44
	9D/Disk 45	9B/Disk 46	98/Disk 47	97/Disk 48
ID=4	90/Disk 49	8F/Disk 50	88/Disk 51	84/Disk 52
	82/Disk 53	81/Disk 54	80/Disk 55	7C/Disk 56
	7A/Disk 57	79/Disk 58	76/Disk 59	75/Disk 60
	74/Disk 61	73/Disk 62	72/Disk 63	71/Disk 64
ID=5	6E/Disk 65	6D/Disk 66	6C/Disk 67	6B/Disk 68
	6A/Disk 69	69/Disk 70	67/Disk 71	66/Disk 72
	65/Disk 73	63/Disk 74	5C/Disk 75	5A/Disk 76
	59/Disk 77	56/Disk 78	55/Disk 79	54/Disk 80
ID=6	53/Disk 81	52/Disk 82	51/Disk 83	4E/Disk 84
	4D/Disk 85	4C/Disk 86	4B/Disk 87	4A/Disk 88
	49/Disk 89	47/Disk 90	46/Disk 91	45/Disk 92
	43/Disk 93	3C/Disk 94	3A/Disk 95	39/Disk 96
ID=7	36/Disk 97	35/Disk 98	34/Disk 99	33/Disk 100
	32/Disk 101	31/Disk 102	2E/Disk 103	2D/Disk 104
	2C/Disk 105	2B/Disk 106	2A/Disk 107	29/Disk 108
	27/Disk 109	26/Disk 110	25/Disk 111	23/Disk 112

 = SES drives

Troubleshooting & Problem Solving

The SFx016 family of enclosures include a processor and associated monitoring and control logic to enable it to diagnose problems within the enclosure's power, cooling and drive systems.

The Enclosure Services Processor is housed along with the Ops Panel in the rear of the enclosure.

The sensors for power and cooling conditions are housed within the Power Supply/Cooling modules. There is independent monitoring for each module.

RM660 Doesn't Recognize the Drive Enclosure

1. Check that the correct Tier Mapping Mode has been selected (see "Tier Mapping for Enclosures" on page 64).
2. Check that the SFP cables from the RM660's to the drive enclosure are fitted correctly.
3. Check the ID settings on your enclosures. All enclosures daisy-chained to the same RM660 channel must have a unique ID number.
4. Check that all drive modules are well seated, the Status LED on all installed drive modules are Green and the Fault LEDs are off.
5. Check that there is a valid FC-AL signal present at the I/O connector. If there is no signal present check that the cable has not been inverted during installation.
6. Check the maximum cable length has not been exceeded (see specifications in "SFx016 Enclosure Technical Specifications" on page 321).

LED Test Mode

The Alarm Mute push-button can be used to test the LEDs on the Ops Panel. When the Mute push-button is held, all LEDs will be illuminated if there are no faults present.

Status LED Indicators

Green LEDs are always used for good or positive indication, flashing green/amber if non-critical conditions exist. Amber LEDs indicate there is a critical fault present within the module.

Power Supply/Cooling Module

The Power Supply/Cooling LEDs are shown in Figure C-46. Under normal conditions the LEDs should all be lit constant green. If a problem is detected the color of the relevant LED will change to amber.




Figure C-46 PSU LEDs

Ops Panel

The Ops Panel displays the aggregated status of all the modules. The Ops Panel LEDs are shown in Figure C-47 and defined in Table C-6.




Figure C-47 Ops Panel LEDs

Table C-6 Ops Panel LED States

Ops Panel LEDs							
Power (Green)	PSU/ Cooling/ Temp (Amber)	System (Amber)	Address Mode Error (Amber)	FC Loop Speed	Hub Mode Selected	Other Associated LEDs or Alarm	State Description
On	Off	Off	Off	Off	Off		5V Aux present, overall power failed
On	On	On	On	On	On	Single beep, then double beep	Ops Panel power On (5s) test state
On	Off	Off	Off				Power On, all functions good
On	On	Off				PSU LEDs or Fan LEDs	Any PSU fault or Fan fault
On	On	Flash					Over or Under temperature

Table C-6 Ops Panel LED States **(continued)**

Ops Panel LEDs							
Power (Green)	PSU/ Cooling/ Temp (Amber)	System (Amber)	Address Mode Error (Amber)	FC Loop Speed	Hub Mode Selected	Other Associated LEDs or Alarm	State Description
On	Off	On				ESI LED on LRC	ESI processor A Failed
On	Off	On				ESI LED on LRC	ESI processor B Failed
On	Flashing	Flashing				PSU Removed	PSU removed and System power redundancy check option set. No indication if option not set.
On	Off	Flashing					No SES Drives installed
On	Off	On				None	Unknown (invalid or mixed) LRC module type installed, or 12C Bus Failure (inter ESI processor), or Backplane autostart watchdog failed.
On	Flashing	On				Intermittent audible alarm	Ops to ESI Communications failed
On			Flashing				Invalid address mode setting (change thumb wheel to valid ranges)
On				On			2Gb FC drive loop speed selected
On					On		RAID ONLY Host side Hub mode enabled

Troubleshooting

The following sections describe common problems, with possible solutions, which can occur with your drive enclosure.

System Faults

Symptom	Cause	Action
1. The SYSTEM LED will illuminate AMBER on the Ops Panel module.	The ESI processor has detected an internal fault (e.g. failure of an internal communications path)	1. Check for other AMBER LED indications on the Power Supply/Cooling modules. If there is a PSU error present there may be a communications problem with that Power Supply/ Cooling module. Remove and then re-fit the module, if the problem persists then replace the module.
2. Audible alarm sound.		2. Check for other AMBER LED indications on the drive modules. If none are evident then there may either be an ESI processor problem or a backplane problem.
		3. Change the Ops Panel module (see "Replacing the Ops Panel" on page 319).

Note: See also section "Thermal Shutdown" on page 315.

Power Supply/Cooling Faults

Symptom	Cause	Action
1. Ops Panel FAULT LED is amber.	1. Any power fault. 2. A fan failure. 3. A thermal condition which could cause PSU overheating.	1. Check power On/Off switch on rear of Power Supply/Cooling module is turned ON. (not accessible on later models)
2. An amber LED on one or more Power Supply/ Cooling modules.		2. Check AC Mains connections to Power Supply/Cooling module is live.
3. Audible alarm sounds.		3. Disconnect the Power Supply/Cooling module from AC power and remove the module from the system, re-install module. If problem persists, replace Power Supply/Cooling Module. 4. Reduce the ambient temperature.

Thermal Control

The SFx016 enclosure uses extensive thermal monitoring and takes a number of actions to ensure component temperatures are kept low and also to minimize acoustic noise. Air flow is from front to rear of the enclosure.

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 power supplies installed.</p>	<ol style="list-style-type: none"> 1. Check the installation for any airflow restrictions at the front and rear of the enclosure. A minimum gap of 1" (25mm) at the front and 2" (50mm) at the rear is recommended. 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

Symptom	Cause	Action
1. Ops Panel FAULT LED is amber.	If the internal temperature measured in the airflow through the enclosure exceeds a pre-set threshold a thermal alarm will sound.	1. Check local ambient environment temperature is below the upper 40°C specification.
2. An amber LED on one or more Power Supply/Cooling Modules.		2. Check the installation for any airflow restrictions at the front and rear of the enclosure. A minimum gap of 1" (25mm) at the front and 2" (50mm) at the rear is recommended.
3. Audible alarm sounds.		3. Check for restrictions due to dust build-up, clean as appropriate.
4. Air temperature exiting PSU is above 55°C.		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

Symptom	Cause	Action
<p>1. ALL amber LEDs on the Ops Panel and on ALL drive bays are flashing.</p> <p>2. Audible alarm sounds almost continuously and cannot be muted.</p>	<p>At a higher threshold than the Thermal Alarm (this should already have been activated) the Enclosure is programmed to shutdown in order to protect itself and the disk drives from damage.</p> <p>OR - All fans have failed.</p> <p>OR - Only 1 fan operating and the internal temperature is 40°C or above.</p>	<p>1. Check for airflow restrictions.</p> <p>2. Check Power Supply/ Cooling module faults.</p> <p>3. Check for excessive local temperatures.</p>
<p>Important: The Enclosure will SHUTDOWN 10 seconds after the above Symptoms are observed. This will leave the following indications active.</p>		
<p>1. FAULT LED illuminated amber.</p> <p>2. Enclosure powered off.</p>	<p>1. Clear the source of the overheating.</p> <p>2. Leave for a period to cool down.</p> <p>3. Remove AC Mains power from the enclosure for at least 30 seconds to reset the shutdown condition.</p> <p>4. Re-start enclosure using normal operating procedure.</p> <p>5. Check for re-occurring cooling faults (especially fan failure).</p>	

FC-AL Drive Module Faults

Disk drive status is indicated by a green Status LED and an amber Fault LED mounted on the front of each drive module, providing the following indications:

Table C-7 LED Functions

State	Green	Amber
No drive fitted	Off	Off
Drive Power ON	On	Off
Drive Activity	Flash	Off
Drive Fault	On	On

Drive activity - LED may be off for a length of time during power up.

Dummy Drive Modules

Dummy drive modules must be installed in all unused drive bays to maintain a balanced air flow.

Auto Start Failure

Unless otherwise selected at installation time, all drives in the enclosure are automatically spun up after power is applied. If this has not occurred, there is a power problem (an alarm and power fault indication would normally be active).

Note: The SYSTEM LED will flash green/amber.

See also “Recovering from Drive Failures” on page 121 for information on drive failure recovery.

Dealing with Hardware Faults

Ensure that you have obtained a replacement module of the same type *before* removing any faulty module.

Warning: If the SFx016 enclosure is powered up and you remove any module, replace it immediately. If the enclosure is used with modules or blank modules missing for an extended period of time, the enclosure can overheat, causing power failure and data loss. Such use will invalidate the warranty.

- Replace a faulty drive with a drive of the same type and capacity.
- All drive bays must be fitted with a drive module or a dummy module in order to maintain a balanced air flow.
- All the supplied plug-in power supply units, electronics modules and blank modules must be in place for the air to flow correctly around the enclosure.

Continuous Operation During Replacement

Depending on how the system is set up, if a disk drive fails, it can normally be replaced without interrupting the use of the system (see “Recovering from Drive Failures” on page 121 for information on drive failure recovery).

In addition, each enclosure contains two Power Supply/Cooling modules, either of which can maintain power and cooling to the enclosure while the other is replaced.

Replacing a Power Supply/Cooling Module

Warning: Do not remove the failed Power Supply/Cooling module unless you have a replacement unit ready for insertion.

If a power supply unit or a fan has failed, you must replace the whole power supply/cooling module.

1. Make sure you identify the faulty Power Supply/Cooling module correctly, from the two modules installed.
2. Switch off the PSU module and disconnect its power cord.
3. Squeeze the two latches together to disengage (Figure C-48).




Figure C-48 Disengage Latches on PSU Module

4. Swing down the handle to cam the module out of the bay and continue to pull the module out.
5. On the new module, check for damage especially to the rear connector. Make sure the switch is turned off.

Caution: Handle the module carefully and avoid damaging the connector pins. Do not install the module if any pins appear to be bent.

6. Disengage the latches.
7. Slide the module into the bay (Figure C-49).




Figure C-49 PSU Modules in Rear Panel

8. Cam the module home by swinging up the handle. Make sure the latches are engaged.
9. Connect the power cord to the module and turn on the power switch.

Note: The alarm will sound until the new Power Supply/Cooling module is operating correctly.

Replacing the Ops Panel

The Ops Panel is an integral part of the enclosure chassis assembly and should only be replaced by trained personnel.

Replacing the I/O Module

Warning: Do not remove this module unless a replacement module is available

1. Make note of how the SFP cables are connected to the failed I/O module. Then disconnect the SFP cables.
2. Using both hands, squeeze open the two latches (Figure C-50).




Figure C-50 LRC I/O Module Latch Operation

3. Pull and swing open the latches to cam the module out of the bay.

4. Grip the latch handles and pull the module out of the bay.
5. Open the latches on the new module. Slide the module into the bay until the latches engage automatically.

Make sure the module orientation is correct (Figure C-51).




Figure C-51 LRC I/O Modules Orientation

6. Cam the module home by manually closing the latches. A click should be heard as the latch engages.
7. Reconnect the SFP cables to the I/O module.

Replacing a Drive Module

Caution: Drive spin down

Damage can occur to a drive if it is removed while still spinning. We recommend that you perform **All** steps of the following procedure to ensure that the drive is spun down prior to removal.

1. Disengage the anti-tamper lock.
2. Release the module handle by pressing the latch in the handle towards the right.
3. Gently withdraw the drive module approximately 1" (25mm) and wait 30 seconds. This will allow the drive to spin down prior to full removal.

- 4. Withdraw the module from the drive bay and install a replacement module in accordance with the instructions in “Installing the Drive Enclosures in Rack” on page 248.

SFx016 Enclosure Technical Specifications

Note: Specifications are subject to change without notice.

Dimensions

Rack-mount: Height 5.28" (134mm);
 Width 17.56" (446mm);
 Depth 19.69" (500mm)

Tower: Height 19.72" (501mm);
 Width 9.06" (230mm), including mounting feet;
 Depth 20.59" (523mm)

Weight

Fully Loaded: rack-mount - 77 lbs (35kg)
 tower - 83.6 lbs (38kg)

Empty Enclosure (Rack): 19.8 lbs (9 kg)

Tower Conversion Kit: 6.6 lbs (3 kg)

PSU/Cooling Module: 8.8 lbs (4 kg)

ESI/Ops Panel Module: 1.98lbs (<0.9kg)

Power Supply Unit (PSU)

Voltage, auto-ranging:	100 - 120 / 200 - 240 VAC
Frequency:	47-63 Hz
Inrush Current:	50A @ 260VAC
Power Factor:	>0.98
Harmonics	Meets EN61000-3-2

Power Cord

(minimum requirements)

Cord Type:	SV or SVT, 18 AWG minimum, 3 conductor
Plug:	250V, 10A
Socket:	IEC 320, C-14, 250V, 10A

Environment

Ambient Temperature & Humidity:

	Temperature Range	Relative Humidity	Maximum Wet Bulb
Operational	10°C to 40°C	20% to 80% non condensing	23°C
Non- Operational	0°C to 50°C	8% to 80% non condensing	27°C
Storage	1°C to 60°C	8% to 80% non condensing	29°C
Shipping	-40°C to +60°C	5% to 100% non precipitating	29°C

Altitude, Operational:	0 to 10,000 feet (3047 m)
Altitude, Non-Operational:	-1,000 to 40,000 feet (-305 to 12, 192m)
Shock, Operational:	Vertical axis 5g peak 1/2 sine, 10ms
Shock, Non-Operational:	30g 10ms 1/2 sine
Vibration, Operational:	0.21grms 5-500 Hz Random
Vibration, Non-Operational:	1.04grms 2-200 Hz Random
Vibration, Relocation:	0.3g 2-200 Hz sine
Acoustics:	Less than 6.0 B LwA - operating at 20°C
Safety & Approvals:	CE, UL, cUL
EMC:	EN 55022 (CISPR - A), FCC A
Orientation & Mounting:	19" Rack mount (3EIA Units)
	Rack: Back pressure not exceeding 5 pascals (0.5mm water gauge)

Interfaces

Drive interface:	See drive module specification
Attachment:	1 loop of 16 drives (SF6016 / SF2016) 2 FC-AL dual loops of 8 drives (SF4016) Passive Backplane with 2 Loop Resiliency Circuit (LRC) I/O Modules FC-AL SFP HSSDC-2 cables (Maximum external length: 10m) FC-AL SFP optical cables (Maximum daisy chain cable length: <i>see SFP manufacturing specification</i>)

FC-AL LRC I/O Module Specification

Module:	Speed 2Gb/1Gb switchable Creates connections to a single loop of 16 drives or dual loops of 8 drives 4 external ports, SFP GBIC connectors
Mounting:	Rear Single Bay 3 and/or 4
Connectors:	4 x 2Gb SFP module with copper HSSDC-2 connector, maximum cable length 33feet (10m) 1 x RS-232, 6-pin mini DIN (factory use only)
LED functions:	Port LED (4 off) - Green: Rx good, Off: no signal ESI/LRC module fault LRC -- amber: fault
Power dissipation:	5A @ 3.3V, 1A @ 5V

Drive Module Specification

Module Dimensions:	Height 1.15" (29.1mm) Width 4.19" (106.55mm) Depth 8.15" (207mm)
Weight:	1.94 lbs (0.88kg) for 1.0" 36GB drive
Operating Temperature:	5°C to 40°C (when installed in an SFx016 enclosure with dual Power Supply/Cooling modules)
Power Dissipation:	18.7 Watts maximum

Software Enclosure Services (SES) Support

The enclosure has a sophisticated self-monitoring and reporting function which conforms to ANSI SES specifications. This reports on such topics as:

- Enclosure temperature
- Fan speed
- Drive condition
- Operator panel status

Spare Parts and Accessories

- Chassis (including backplane)
- 19-inch rack mounting rail kit
- AC Power Supply/Cooling module
- Ops panel module
- Drive module and key
- Front dummy module
- LRC I/O module
- Various types of external FC-AL signal cables
- SFP module, optical or copper
- (Country specific) power cords
- this User Guide

Complete List of Commands at OEM Level

This appendix covers the Command Line Interface (CLI) commands for OEM (Original Equipment Manufacturer) access. Description and usage examples are given for each command.

Available Commands:

ALARM
APC_UPS
API
AV
CACHE
COMMENT
CONSOLE
DATE
DEBUG
DEFAULTS
DISK
DISK_CLI
DUAL
FAULTS
HELP
HOST
HOST_CLI
LICENSE
LOG
LOGIN
LOGOUT
LUN
MIRROR
NETWORK
PASSWORD

RESTART
ROUTE
SAVE
SES
SETTINGS
SHOWALL
SHUTDOWN
STATS
TELNET
TFTP
TIER
TIME
UPTIME
USER
VERSION
WHOAMI
ZONING

ALARM

Use the ALARM command to control the functions pertaining to AVR System Alarm.

ALARM

ALARM parameter

Command Syntax Detail:

Parameter:

Description:

TEST

Tests the AVR system alarm.

Example:

```
RM660[1]: alarm test
Testing AVR System Alarm
Press Enter to Discontinue Test
```

APC_UPS

Use the APC_UPS to display the status of, enable or disable the APC UPS trap monitor, and to clear any outstanding APC UPS traps.

APC_UPS
APC_UPS parameter
Command Syntax Detail:
Parameter:
Description:

MONITOR=ON
Enables the APC UPS monitor.

MONITOR=OFF
Default setting. Disables the APC UPS monitor.

CLEAR_FAULTS
This parameter deletes all pending APC UPS faults from the fault list. All APC UPS events that disabled writeback caching are cleared.

Example

The following example displays the default status, enables monitoring, and clears apc_ups fault list:

```
RM660[1]: apc_ups
APC UPS SNMP trap monitor is off.
No APC UPS faults detected via SNMP trap.
RM660[1]: apc_ups monitor=on
RM660[1]: apc_ups clear_faults
```

API

Use the API command to display and change the enabling/disabling of API connections current (temporarily) enabled or disabled.

API parameter

Command Syntax Detail:

PARAMETER:

Description:

STATS

Displays the collected statistics on API connections.

CLEARSTATS

Resets the collected statistics on API connections.

DISABLE

Temporarily disables the establishment of connections to the API Server. Users at remote locations will be unable to establish a new API connection until an API ENABLE command is issued.

ENABLE

To (re-)enable the establishment of connections to the API Server.

Usage Guidelines:

This command only provides control over API connections during the current power cycle. To "permanently" disable or enable API connections (that is, across power-cycles), refer to the NETWORK [[API_SERVER=ON|OFF] command.

Default setting for this command at power-on is ENABLED.

Example:

```
RM660[1]: api
API Server connections are currently -- ENABLED --
```

AV

Use the AV commands with the appropriate parameters to display information about or to change the Audio/Visual settings of the system. Use the parameters to tune the system and the disks for better performance and a lower latency.

AV

AV parameter

Command Syntax Detail:

PARAMETER:

Description:

FASTAV=ON

Enables the disk fast Audio/Video read option for streaming data. If the FASTAV parameter is ON, the system starts 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 per-LUN basis. Use in combination with the LUN=x parameter to change the settings for a single LUN.

FASTAV=OFF

Default setting. Disables the disk fast Audio/Video read option for 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.

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 0.

DISKAV=ON

Enables the disk Audio/Video options for streaming data. When ON, the system adjusts disk parameters to minimize the latency for data transfers by disabling non-essential features which may impact performance.

DISKAV=OFF

Default setting. Disables the disk Audio/Video options for streaming data.

TIMELIMIT=x

Specifies the recovery time limit for the drives when DISKAV is enabled. The maximum amount of time, in one millisecond increments, that a disk can use for the data error recovery. The parameter is located in bytes 10 and 11 of the Read-write error recovery mode page on the disks. The recovery time limit of each disk is set to its default value when DISKAV=OFF.

Valid range is 0 to 65535.

Default is 65535.

ARRE=ON

Specifies the Automatic Read Reallocation Enabled (ARRE) bit setting on the drives when DISKAV is enabled. When ON, the disks automatically reallocate defective data blocks during read operations.

NOTE: Automatic Write Reallocation Enabled (AWRE) is always enabled.

The parameter is located in bit 6, byte 2 of the Read-write error recovery mode page on the disks.

ARRE=OFF

Specifies the Automatic Read Reallocation Enabled (ARRE) bit setting on the drives when DISKAV is enabled. When OFF, the disks do not reallocate defective data blocks during read operations. Automatic Write Reallocation Enabled (AWRE) is always enabled. The parameter is located in bit 6, byte 2 of the Read-write error recovery mode page on the disks. The ARRE of each disk is set to its default value (OFF) when DISKAV=OFF.

Default is OFF.

ERR=ON

Default setting.

Specifies setting of the Early Error Recovery (EER) bit on the drives when DISKAV is enabled. When enabled, the disks perform the most expedient form of error recovery first. When disabled, the disks use an error recovery procedure that minimizes the risk of mis-detection or mis-correction.

The parameter is located in bit 3, byte 2 of the Read-write error recovery mode page on the disks. The EER of each disk is set to its default value when DISKAV=OFF.

ERR=OFF

Default setting.

Specifies setting of the Early Error Recovery (EER) bit on the drives when DISKAV is enabled. When enabled, the disks perform the most expedient form of error recovery first. When disabled, the disks use an error recovery procedure that minimizes the risk of mis-detection or mis-correction.

The parameter is located in bit 3, byte 2 of the Read-write error recovery mode page on the disks. The EER of each disk is set to its default value(ON) when DISKAV=OFF.

READRETRY=X

Default setting is 1.

Specifies the read retry count setting on the drives when DISKAV is enabled. This field indicates the number of times a disk attempts its recovery algorithm during a read operation.

This parameter is located in byte 3 of the Read-write error recovery mode page on the disks. The retry count of each disk is set to its default value (1) when DISKAV=OFF.

Valid range is 0 to 255.

WRITERETRY=x

Specifies the write retry count setting on the drives when DISKAV is enabled; indicating the number of times a disk attempts its recovery algorithm during a write operation.

This parameter is located in byte 8 of the Read-write error recovery mode page on the disks.

The retry count of each disk is set to its default value when DISKAV=OFF.

Valid range is 0 to 255.

Default is 1.

ORDEREDQUEUE=x

Enables the RM controller to use ordered tags when communicating with the drives. The value, 'x', correlates to the percentage of ordered tags, where:

- 0 indicates no ordered tags,
- 1 indicates 100% ordered tags,
- 2 indicates 50% ordered tags,
- 3 indicates 33% ordered tags,
- 4 means 25% ordered tags,
- etc...

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

Enables the Read Continuous 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.

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

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.

Enabling this feature automatically enables FASTAV.

RC=OFF

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.

RCTIMEOUT=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 0.

LUN=x

Use in combination with the FAST_AV and RC parameters, in order to specify which LUN is to be changed. By default, if no LUN is specified with this parameter, then all the LUNs in the system are updated.

Valid LUNs are 0 to 127.

Default is all LUNs.

DISKUPDATE

Use this parameter to recheck all of the mode parameters for the disks in the system. This allows you to update the disk mode parameters after changing several of the AV parameters instead of changing them one at a time.

Usage Guidelines:

CAUTION! Changing the disk parameters can adversely affect the I/O operation of the system. Adjust only when the system is idle.

NOTE: Use the CACHE command to change the writeback and prefetch settings for each LUN.

Example:

```
RM660[1]: av
```

Current LUN Audio/Visual settings					
LUN	Label	Audio/Visual settings		Maximum	
		FastAV	Continuous	Read Caching	Write Prefetch
0		Off	Off	On	x 1
1		Off	Off	On	x 1
2		Off	Off	On	x 1

Disk Audio/Visual settings are: Disabled (Using disk defaults)

```

Early Error Recovery:      Enabled
Automatic Read Reallocation: Disabled
Read Retry Count:         1
Write Retry Count:        1
Recovery Time Limit:      65535

```

```

Ordered Tag Count:        0
Unit Attention:           Enabled
FASTAV Timeout:           50
RC Timeout:                8
Fail Check Condition:     Disabled

```

CACHE

Use the CACHE command to display the current cache settings for each LUN in the system and adjust the cache settings of the LUNs to change the performance of the system.

These cache parameters are identical to the cache parameters in the SCSI caching mode page.

CACHE parameter

Command Syntax Detail:

PARAMETER:

Description:

LUN=X

Specifies the specific LUN(s) to be affected when used in combination with any of the following other parameters:DEFAULTS, MAX, MF, PREFETCH, WRITEBACK

If no LUN is specified, then all LUNs will be updated by default.

Valid LUN values for X are 0 to 127.

Default is ALL LUNs.

WRITEBACK=ON

Enables write back caching which allows the system to increase the performance of write I/O requests by storing the data in cache and saving the data to the disks at a later time. Default is ON.

WRITEBACK=OFF

Disables write back caching.

PREFETCH=x

Sets the prefetch that occurs on read commands.

If the MF parameter is OFF, the system prefetchs the number of blocks specified by this parameter after every read command.

If the MF parameter is ON, then the system multiplies the transfer length of the command by this parameter to determine how much data is prefetched.

A prefetch value of less than 8 is recommended when the MF parameter is ON.

Valid range is 0 to 65535.

Default is 1.

MAX=x

Sets the maximum prefetch ceiling in blocks for prefetches on read commands; sets an upper limit on prefetching when the MF parameter is ON. The system automatically limits the amount of prefetching if the system is running low on resources.

Valid range is 0 to 65535.

Default is 65535.

MF=ON

Default setting that enables the Multiplication Factor bit. The system multiplies the transfer length of the command by the PREFETCH parameter to determine how much data is prefetched.

MF=OFF

Disables the Multiplication Factor bit. The system prefetches the number of blocks specified by the PREFETCH parameter after every read command.

SIZE=X

Sets the cache segment size in Kbytes, for the system which allows you to adjust the performance of the system by changing the cache segment size to match the size of the host I/O requests.

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 the best performance, the cache segment size should be larger than the average host I/O request size.

Use the STATS LENGTH command to determine the average host I/O request size.

Do not use this command under heavy I/O conditions because the system temporarily halts all I/O requests while the changes are taking effect.

Valid segment sizes are 64, 128, 256, 512, 1024, and 2048.

Default is 128.

DEFAULTS

Loads the default settings for all of the cache parameters for the specified LUNs but does NOT change the WRITELIMIT or SIZE parameters.

NOTE: The DEFAULTS value is applied first to the specified LUNs when this command is used in combination with any of the following other parameters: WRITEBACK, PREFETCH,

MAX, MF

WRITELIMIT=X

Specifies the maximum percentage of the cache that can be used for write back caching. The system forces all writeback requests to be flushed to the disks immediately if the percentage of writeback data in the cache exceeds this value.

Range is 0 to 100.

Default is 75.

Example:

```
RM660[1]: cache
```

Current Cache settings							
LUN	Write Caching	Maximum Prefetch	MF Bit	Prefetch Ceiling	Read Priority	Write Priority	
0	Enabled	x 1	On	65535	high	high	
1	Enabled	x 1	On	65535	high	high	
2	Enabled	x 1	On	65535	high	high	

```
Writeback Limit: 75%
```

```
2560.0 Mbytes of Cache Installed
(32768 Segments of 64 Kbytes)
```

COMMENT

Use the COMMENT command to echo a message to the screen and save it to the LOG and the syslog if it is enabled.

COMMENT text

Usage Guidelines:

Any printable text can be entered on the command line (except help or ?). See example.

Example:

```
RM660[1]: comment comment
RM660[1]:
CLI_MAIN 8-6 08:41:55 comment
```

```
RM660[1]: comment help
```

COMMENT Echoes a message to the screen and saves it to the LOG.

This command is used to echo a message to the CLI. The message is saved in the LOG and is also sent to syslog if it is enabled. Any printable text can be entered on the command line.

<message>

Any printable text can be entered in the message.

```
RM660[1]: comment help help
```

COMMENT Echoes a message to the screen and saves it to the LOG.

This command is used to echo a message to the CLI. The message is saved in the LOG and is also sent to syslog if it is enabled. Any printable text can be entered on the command line.

There is no help available for parameter HELP.

CONSOLE

Use the `CONSOLE` command to display information about the current serial console settings of this unit or to change the baud rate for the serial console.

`CONSOLE`

`CONSOLE` parameter

Command Syntax Detail:

`PARAMETER:`

Description:

`BAUD`

Displays a list of console baud rates and allows you to set a new console baud rate. You also have the option of escaping from command and making no changes by selecting (e). Available baud rates: 9600, 19200, 38400, 57600, 115200.

Example:

The following example shows the current setting:

```
RM660[1]: console
```

Serial console baud rate is 115200 baud.

The following example displays the current setting and allows you to change the setting:

```
RM660[1]: console baud
```

Select the new serial console baud rate from choices below:

- 1 - 9600
- 2 - 19200
- 3 - 38400
- 4 - 57600
- 5 - 115200

DATE

Use the DATE command to display or change the current system date.

DATE

DATE parameter

Command Syntax Detail:

PARAMETER:

Description:

mm dd yyyy

Changes the system date to the new value indicated by <mm> <dd> <yyyy>, where:

mm: indicates the month in the range 1..12

dd: indicates the day in the range 1..31

yyyy: indicates the year in the range 2000..2104

Example:

To simply display the system date:

```
RM660[1]: date
```

Current date is 8-06-2004

To change the system date to Sept. 10, 2041:

```
RM660[1]: date 9 10 2041
```

The system error traps for invalid dates:

```
RM660[1]: date 13 31 3131
```

Invalid parameter: 13

```
RM660[1]: date 9 31 2005
```

9 31 2005 is an invalid date.

```
RM660[1]: date 9 30 1995
```

Invalid parameter: 1995

The system allows you to back date. The example below shows a date change to a year prior to current year (2004):

```
RM660[1]: date 8 20 2003  
Current date is 8-20-2003
```

DEBUG

Use the DEBUG command to enable or disable the system debug messages.

The debug command is provided for diagnostic purposes and should only be used by qualified service personnel.

NOTE: OEM access level does NOT have these privileges. Contact SGI for further assistance.

DEFAULTS

Use the DEFAULTS command to restore the system to its default configuration. The system halts all I/O requests, deletes all the LUNs, and restores all the parameters back to their defaults values.

WARNING! This is a destructive operation which deletes all the data stored in the system.

The system then asks if it should erase all the configuration information stored on the disks. This erasing of the disks prevents the system from retrieving the backup copies of the configuration settings from the disks after the system is restarted.

After the defaults are loaded, the system asks if it should begin reconfiguring by scanning for the disks. New LUNs can be created after the disks have been added back to the system.

DISK

Use the DISK command to display and change the configuration settings for the disks in the system and monitor the status of the disk channels. The command displays the current disk configuration settings and the status of each disk channel.

The INFO= parameter can be used to display all of the information about a disk in the system.

The LIST parameter displays a list of the disks installed in the system and indicates how many were found.

DISK

DISK parameter

Command Syntax Detail:

PARAMETER:

Description:

INFO=tc

Displays the information and status about a specific disk in the system. The disk is specified by its physical tier and channel locations, 'tc', where:

't' indicates the tier in the range <1..125>, and

'c' indicates the channel in the range <ABCDEFGHPS>.

LIST

Displays a list of all the disks installed in the system and indicates how many were found of each type.

RLS=tc

Requests and displays the Read Link Error Status Block information for the specified drive.

The RLS information consists of the following items:

Link -> Link Failure Count

Sync -> Loss of Synchronization Count

Signal -> Loss of Signal Count
 PSPE -> Primitive Sequence Protocol Error
 ITW -> Invalid Transmission Word
 CRC -> Invalid CRC Count
 F7Init -> Lip F7 initiated count
 F7Rec -> Lip F7 received count
 F8Init -> Lip F8 initiated count
 F8Rec -> Lip F8 received count

The disk is specified by its physical tier and channel locations, 'tc', where:

't' indicates the tier in the range <1..125>, and

'c' indicates the channel in the range <ABCDEFGHPS>.

If neither the tier nor the channel are specified, the RLS information is requested from all drives.

If only the channel is specified, the RLS information is requested from all drives on the specified channel.

If only the tier is specified, the RLS information is requested from all the drives on the specified tier.

LIPINFO

Displays the last LILP payload for all disk channels.

DEFECTLIST=tc

Displays the number of defects in the defect list for the specified disk. The defect list contains all the physical sectors on the disk that the drive has identified as bad, and to which the disk's hardware prevents access.

The list is classified into two types:

The permanent list consists of the bad sectors that are identified by the disk manufacturer;

the grown list consists of the bad sectors that are found after the disk has left the factory (and which can be added to at any time).

The disk is specified by its tier and channel locations, 'tc', where:

't' indicates the tier in the range <1..125>, and

'c' indicates the channel in the range <ABCDEFGHPS>.

STATUS

Displays the loop status of each disk channel and a count of the fibre channel errors encountered on each channel.

STATUSCLEAR

Resets the fibre channel error counts on each disk channel.

FAIL=tc

Instructs the system to fail the specified disk at the physical tier: tier (t) in the range <1..125>, and channel (c) in the range <ABCDEFGHPS>.

When a non-SPARE disk is specified and it is failing
the disk won't 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.

REBUILD=tc

Instructs the system to start a rebuild operation on a (presumably) already failed disk. A rebuild operation restores a failed disk to a healthy status once it completes. Note that this operation can take several hours to complete depending on the size of the disk and the speed of the rebuild operation. The speed of the rebuild operation can be adjusted with the DELAY and EXTENT parameters of the TIER command.

In addition, the rebuild operation can be stopped, or paused and resumed with the TIER STOP, TIER PAUSE, and TIER RESUME commands.

The TIER AUTOREBUILD command can be used to automate the rebuild process.

Note that SPARE disks are handled slightly differently from other disks, in that SPARES that are not in use as an active replacement for a failed disk elsewhere in the system are simply returned to a normal healthy status by this command; SPARES that are in use are already considered healthy and are not rebuilt.

The failed disk to be rebuilt is specified by its physical:
tier (t) in the range <1..125>, and channel (c) in the range <ABCDEFGHPS>.

REBUILDVERIFY=ON

Instructs the system to 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. Note: This feature will increase the time it takes for rebuilds to finish.

REBUILDVERIFY=OFF

Instructs the system NOT to send SCSI Write with Verify commands to the disks when rebuilding failed disks. Default setting.

REPLACE=tc

Instructs the system to replace the specified failed disk with a spare disk. A replace operation is used to temporarily replace a failed disk with a healthy spare disk. This 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 of the TIER command.

The failed disk to be replaced is specified by its physical:
tier (t) in the range <1..125>, and
channel (c) in the range <ABCDEFGHPS>.

(Note that spare disks themselves cannot be replaced with this command).

SCAN

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

TIMEOUT=x

Sets the disk timeout (in seconds) for an I/O request. Valid range is 1 to 512 seconds. Default is 21 seconds.

PLOGI=ON

Forces the initiator to issue a PLOGI/PRLI after a loop initialization. This parameter must be set to ON if SA2016/SA4016 are used.

PLOGI=OFF

Default setting. is OFF. The initiator uses ADISC.

SPINDOWN=tc

Sends a stop command to the specified disk in the system. The stop command causes the disk to spin down, which allows it to be removed or transported.

The disk to be stopped is specified by its physical:
tier (t) in the range <1..125>, and
channel (c) in the range <ABCDEFGHPS>.

SPINUP=tc

Sends a start command to the specified disk in the system. The start command will cause a disk which has been spundown with the SPINDOWN command to spin up again. The disk to be started is specified by its physical:
tier (t) in the range <1..125>, and
channel (c) in the range <ABCDEFGHPS>.

DIAG=tc

Performs a series of diagnostics tests on the specified disk.
The disk is specified by its physical:
tier (t) in the range <1..125>, and
channel (c) in the range <ABCDEFGHPS>.

REASSIGN=tc 0xh

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..125>, and
channel (c) in the range <ABCDEFGHPS>.

0xh is the hexadecimal value of the LBA (Logical Block Address) to be reassigned.

LLFORMAT=tc

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..125>, and
channel (c) in the range <ABCDEFGHPS>.

AUTOREASSIGN=ON

Bad blocks are reassigned when a medium error occurs on a healthy tier.

AUTOREASSIGN=OFF

Bad blocks are NOT reassigned when a medium error occurs on a healthy tier.

MAXCMDS=x

Sets the maximum command queue depth to a tier of disks.
Range: 1 to 32 commands per tier.
Default: 32 commands.
Settings:

6 for SA4016/SA2016 in singlet mode
2 for SA4016/SA2016 in couplet mode
32 for everything else

Example:

```
RM660[1]: disk
    Disk Channel Status
    Disk Channel A healthy.
    Disk Channel B healthy.
    Disk Channel C healthy.
    Disk Channel D healthy.
    Disk Channel E healthy.
    Disk Channel F healthy.
    Disk Channel G healthy.
    Disk Channel H healthy.
    Disk Channel P healthy.
    Disk Channel S healthy.

    Disk 1A is failed and replaced by spare 1S.
    Disk 2A is failed and replaced by spare 2S.
    Disk 3A is failed and replaced by spare 3S.

    Disk rebuild verify:      Disabled
    Disk command timeout:    27 seconds
    Disk maximum commands:   32
    Disk PLOGI:              Disabled
    Auto Reassign:           Enabled

    Disk commands outstanding: 1
```

DISK_CLI

Use the DISK_CLI command to send a CLI command to any of the disk channels as specified.

CAUTION! This command is provided for diagnostic purposes and should only be used by qualified service personnel.

Note: The DISK UART command is invalid for OEM level access. Contact DDN for further assistance.

DISK_CLI PARAMETERS

Command Syntax Detail:

PARAMETER:

Description:

DISK=C

Specifies which disk channel(s) (C); more than one disk channel may be listed in the range of A through H, P, S or ALL.

MINUTES TIMEOUT=m

Specifies the number of minutes, m, to wait for the disk G/F(s) to complete the supplied disk CLI command, in the range 1..120. If a minutes value is not supplied, the default timeout of 120 minutes is used.

disk_cli_command

Indicates the disk CLI command and parameters to send to the specified disk channels.

Example:

```
RM660[1]: disk_cli disk=all minutes=60
```

```
Waiting for disk response(s); Timeout=60 minutes.
```

```
The Disk CLI command did not complete in time.
```

```
+-----+
| Disk CLI Command Results:
|           Disk Channel:
|           A  B  C  D  E  F  G  H  P  S
|           -----
| NOT RUN   :
| TIMED OUT :  x  x  x  x  x  x  x  x  x  x
| FAILED    :
| SUCCEDED  :
+-----+
```

DISK_CLI: Command failed at the Galeforce(s).

DUAL

Use the DUAL command to display information about the dual system configuration. Use the appropriate parameter to change the related configuration.

DUAL

DUAL parameter

Command Syntax Detail:

PARAMETER:

Description:

FAIL

Fails the partner unit in the system.

HEAL

Restores the partner unit in the system to a healthy status.

COHERENCY=ON

Enables the cache coherency between the two units.

COHERENCY=OFF

Disables the cache coherency between the two units.

TIMEOUT=x

Sets the dual cache coherency timeout for cache node requests. The timeout value is given in seconds. A value of zero allows for only one retry. Timeout value should be less than the host timeout (HOST TIMEOUT=x).

The valid range for 'x' is 0 to 255 seconds. Default is 2 seconds.

LABEL=x

Changes the label assigned to each unit, to uniquely identify each unit in the system. Valid values for 'x' are 1 and 2. The CLI prompt for each unit is built by adding a colon and a space at the end of the label.

Each unit can have a label up to 31 characters long.

Entering DEFAULT restores the label of the unit to its default setting.

Example:

```
RM660[1]: dual
```

```
Dual RM660 Configuration
```

```
Unit 1
```

```
Unit 2
```

```
-----  
Label RM660[1]          RM660[2]  
Status                Healthy          Healthy
```

```
Dual communication:    established.
```

```
Ethernet communication: established.
```

```
Cache coherency:      established.
```

```
Cache coherency timeout: 0
```

FAULTS

Use the FAULTS command to display all current faults, warnings and errors in the system, providing a convenient way to quickly check the status of the system.

FAULTS

FAULTS parameter
Command Syntax Detail:
PARAMETER:
Description:
SFP

Displays the current status of the host and disk SFPs. 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.

Example:

```
RM660[1]: faults
```

```
      Current System Faults
-----
Disk 1A is failed and replaced by spare 1S.
Disk 2A is failed and replaced by spare 2S.
Disk 3A is failed and replaced by spare 3S.
```

```
--- Faults detected! ---
```

```
RM660[1]: faults sfp
```

```
Disk channel A : SFP transmitter OK, SFP signal OK
Disk channel B : SFP transmitter OK, SFP signal OK
Disk channel C : SFP transmitter OK, SFP signal OK
Disk channel D : SFP transmitter OK, SFP signal OK
Disk channel E : SFP transmitter OK, SFP signal OK
Disk channel F : SFP transmitter OK, SFP signal OK
Disk channel G : SFP transmitter OK, SFP signal OK
Disk channel H : SFP transmitter OK, SFP signal OK
Disk channel P : SFP transmitter OK, SFP signal OK
Disk channel S : SFP transmitter OK, SFP signal OK
```

Host port 1 : SFP transmitter OK, SFP signal OK
Host port 2 : SFP transmitter OK, SFP signal OK
Host port 3 : SFP transmitter OK, SFP signal OK
Host port 4 : SFP transmitter OK, SFP signal OK

HELP

Use the HELP command to display help information about system commands. By default, this command displays a list of all the available system commands that are available to your access level. This command also displays detailed information about a specific command if one is specified. Information about the specific parameters of a command can be displayed if the command is specified followed by the parameters of interest.

Usage:

```
HELP ?  
HELP ? <command>  
HELP ? <command> <parameter, parameter, ...>  
<command> HELP ?
```

```
<command> HELP ? <parameter, parameter, ...>
```

Command Syntax Detail:

PARAMETER:

Description:

SHORT

Displays a shorter form of help for the command(s) and/or parameter(s). This flag only applies to the current HELP command's execution. It must appear immediately adjacent to the HELP command in order to be honored with it.

Example:

To display detailed help on the help command:

```
RM660[1]: HELP HELP
```

To display detailed help on any parameters beginning with a 'P' for the help command:

```
HELP HELP
```

[SHORT]

Displays a shorter form of help for the command(s) and/or parameter(s). This flag only applies to the current HELP command's execution. It must appear immediately adjacent to the HELP command in order to be honored with it.

HOST

Use the HOST command to display information about the host fibre channel ports and to change the configuration settings for the host fibre channel ports in the system and monitor their status.

The command displays:

the current settings and status for each host port

a list of the host users currently logged into the system.

HOST

HOST parameter

Command Syntax Detail:

PARAMETER:

Description:

LIPINFO

Shows the last LILP payload for all host ports.

STATUS

Displays the loop status of each host port and a count of the fibre channel errors encountered on each port.

STATUSCLEAR

Resets the fibre channel error counts on each port.

TIMEOUT=x

Sets the host command timeout for an I/O request to the value specified by 'x'. Valid range for 'x' is 1 to 512 seconds. Default is 30 seconds.

ID=x

Changes the hard loop ID of a host port. The supplied value, 'x', is the fibre channel AL PA value which is used by the host port. The system selects a soft ID if the hard loop ID is already taken by another device. Refer to the manual for a list of valid loop IDs. Enter this parameter as an 8-bit hex value. Default is EF.

WWN=x

WWN=0

WWN=DEFAULT

Overrides the system ID and specifies a different World Wide Name for a host port. Enter the new World Wide Name value as one of:

- a 64-bit hex value (x)
- default value of 0 (0)
- default value of 0 (DEFAULT)

PORT=x

PORT=ALL

Specifies the specific host port(s) to be affected when used in combination with any of the following other parameters: ID, TIMEOUT, UART, WWN. If PORT is left unspecified, the user is prompted for choice of host ports. Valid port values (x) are 1 to 4. Default is ALL host ports.

ARRAYPARITY=ON

Enables host array parity checking on all the host ports. This feature is normally used for testing only.

ARRAYPARITY=OFF

Default setting. Disables host array parity checking on all the host ports.

SPEED

Sets and changes the port speed on the host port(s). The user is prompted for the desired speed as well as for the choice of the host port(s). Note: When displaying the speed settings with the HOST command, the following acronyms are used:

- Gbps - Gigabits per second
- NA - Not Applicable

Example:

RM610[2]: host

Host Port Configuration

Host Port	Hard Loop ID	Current Loop ID	Port Speed Desired/Actual	Timeout seconds	World Wide Name	Loop Status
1	EF	16DEF	1Gbps/-NA-	30	21000001FF0202B8	Not connected
2	EF	16EEF	1Gbps/-NA-	30	22000001FF0202B8	Not connected
3	EF	169EF	1Gbps/-NA-	30	23000001FF0202B8	Not connected
4	EF	16AEF	1Gbps/-NA-	30	24000001FF0202B8	Not connected

Current Logins

User	Port	Frame	S_ID	World Wide Name	Login

LICENSE

Use the LICENSE commands to display and change the unit's current licensing information.

The licensing information describes whether or not a given feature is licensed, whether that license is permanent or temporary, and if temporary, how many days the license has before expiration.

The licensing information display includes the license status for the following licensable features:

direction -- for Host Port licensing.

directDISCOVER -- for Tier licensing.

Note that:

'*' indicates a tier that is currently available on the system.

'+' indicates an item that was changed for a REQUEST.

You may edit a license request which is activated as the new current licensing information when the you install the appropriate licensing key value.

LICENSE

LICENSE parameter

Command Syntax Detail:

PARAMETER:

Description:

INFO

Review the Firmware version, OEM value, unit IEEE Serial Number, and any pending Requested licensing information.

REQUEST

Edit a set of Requested licensing information and then use it to create a formal License REQUEST.

INSTALL

Install a new License Key in order to activate the pending License Request.

CLEAR
Clears a pending License request.

Example:

RM660[1]: license

RM660#1: Current License Map

=====

Port	Licensed	Expires/Permanent	Days Remaining
1	YES	PERMANENT	N/A
2	YES	PERMANENT	N/A
3	YES	PERMANENT	N/A
4	YES	PERMANENT	N/A

Tier Licensed Tier Licensed Tier Licensed Tier Licensed Tier Licensed

*1	YES	2	YES	3	YES	4	YES	5	YES
6	YES	7	YES	8	YES	9	YES	10	YES
11	YES	12	YES	13	YES	14	YES	15	YES
16	YES	17	YES	18	YES	19	YES	20	YES
21	YES	22	YES	23	YES	24	YES	25	YES
26	YES	27	YES	28	YES	29	YES	30	YES
31	YES	32	YES	33	YES	34	YES	35	YES
36	YES	37	YES	38	YES	39	YES	40	YES
41	YES	42	YES	43	YES	44	YES	45	YES
46	YES	47	YES	48	YES	49	YES	50	YES
51	YES	52	YES	53	YES	54	YES	55	YES
56	YES	57	YES	58	YES	59	YES	60	YES
61	YES	62	YES	63	YES	64	YES	65	YES
66	YES	67	YES	68	YES	69	YES	70	YES
71	YES	72	YES	73	YES	74	YES	75	YES
76	YES	77	YES	78	YES	79	YES	80	YES
81	YES	82	YES	83	YES	84	YES	85	YES
86	YES	87	YES	88	YES	89	YES	90	YES
91	YES	92	YES	93	YES	94	YES	95	YES
96	YES	97	YES	98	YES	99	YES	100	YES
101	YES	102	YES	103	YES	104	YES	105	YES
106	YES	107	YES	108	YES	109	YES	110	YES
111	YES	112	YES	113	YES	114	YES	115	YES
116	YES	117	YES	118	YES	119	YES	120	YES
121	YES	122	YES	123	YES	124	YES	125	YES

LOG

Use the LOG command to display a log of previous system messages.

This log is saved in non-volatile memory and will automatically roll over when full.

LOG parameter

Command Syntax Detail:

PARAMETER:

Description:

CLEAR

Clears the log of all previous messages.

CHECKCONDITION

CHECKCONDITION=MORE

Displays the Check Condition Log. Use the option MORE to display additional information concerning the check condition.

CHECKCLEAR

Clears the Check Condition Log.

Example:

```
RAID58 RM610[2]: log
```

```
Current log size: 60756 characters.
```

```
==== START of System Message LOG Contents: ====
```

```
CLI_MAIN 8-23 15:25:29 Message Log cleared.
```

```
CLI_MAIN 8-23 15:27:01 cleared log monday 08-23-04 3:30PM
```

```
DataDirect Networks Silicon Storage Appliance Model 3000  
Firmware Version: 5.02
```

```
Telnet Services enabled.
```

```
API Services enabled.
```

```
Loading System
```

```
Checking Cache
```

```
Booting Disks
Checking Disks
Loading users and zoning
Checking Dual Communication
Checking Licensing
Dual Mode: Setup Communications success.
Dual_Int 8-24 11:30:33 Warning: Dual mode - resynchronizing seq_cnt,
old=0000, new=0001
Loading DMT tasks

Warning: Improper shutdown of the system detected.
        Always use SHUTDOWN or RESTART

System Online

Ethernet communication established(00)
Cache coherency established.
BIT_MON 8-24 12:07:25 Host port 1 : SFP signal OK
Host_Int 8-24 12:07:26 Login Anonymous WWN:20FD006069806046, port:1
S_ID:FFFC01
Host_Int 8-24 12:07:26 Logout Anonymous WWN:20FD006069806046, port:1
S_ID:FFFC01
BIT_MON 8-24 12:09:29 Host port 3 : SFP signal OK
Host_Int 8-24 12:09:31 Login Anonymous WWN:20FD006069806046, port:3
S_ID:FFFC01
Host_Int 8-24 12:09:31 Logout Anonymous WWN:20FD006069806046, port:3
S_ID:FFFC01
BIT_MON 8-24 12:09:35 Host port 4 : SFP signal OK
Host_Int 8-24 12:09:36 Login Anonymous WWN:20FD006069806046, port:4
S_ID:FFFC01
Host_Int 8-24 12:09:36 Logout Anonymous WWN:20FD006069806046, port:4
S_ID:FFFC01
BIT_MON 8-24 12:09:43 Host port 2 : SFP signal OK
Host_Int 8-24 12:09:45 Login Anonymous WWN:20FD006069806046, port:2
S_ID:FFFC01
Host_Int 8-24 12:09:45 Logout Anonymous WWN:20FD006069806046, port:2
S_ID:FFFC01
Host_Int 8-25 08:57:25 Login Anonymous WWN:210100E08B2719A8, port:4
S_ID:152A00
Host_Int 8-25 08:57:25 Login Anonymous WWN:210100E08B2719A8, port:2
S_ID:152A00
Host_Int 8-25 08:57:25 Login Anonymous WWN:210100E08B2719A8, port:3
S_ID:152A00
Host_Int 8-25 08:57:25 Login Anonymous WWN:210100E08B2719A8, port:1
S_ID:152A00
```

```

Host_Int 8-25 11:16:03 Login Anonymous WWN:210000E08B1021EF, port:3
S_ID:010F00
Host_Int 8-25 11:16:03 Login Anonymous WWN:210000E08B1021EF, port:4
S_ID:010F00
Host_Int 8-25 11:16:04 Login Anonymous WWN:210000E08B1021EF, port:1
S_ID:010F00
Host_Int 8-25 11:16:04 Login Anonymous WWN:210000E08B1021EF, port:2
S_ID:010F00
Time_Int 8-25 12:04:22 Timeout: SCSI cmd:2A LUN 6 DMT_301 Lane:0 T:300
a: 634758 l: 4 21/21 01,00 W:RSP
Anonymous WWN:210100E08B2719A8 port:3 lane:0 OX_ID:29E0
0 a: 6347 w:y l:y h:1/0,1 r:0/0 w:0/0 ea:0,0 DLG
DMT_301 8-25 12:04:22 Command Aborted: SCSI cmd:2A LUN 6 DMT_301
Lane:0
T:300
a: 634758 l: 4 00/21 01,01 W:RDY AB
Anonymous WWN:210100E08B2719A8 port:3 lane:0 OX_ID:29E0
Host_Int 8-25 13:30:19 Login Anonymous WWN:210000E08B1021EF, port:3
S_ID:010F00
Host_Int 8-25 13:30:19 Login Anonymous WWN:210000E08B1021EF, port:4
S_ID:010F00
Host_Int 8-25 13:30:20 Login Anonymous WWN:210000E08B1021EF, port:1
S_ID:010F00
Host_Int 8-25 13:30:20 Login Anonymous WWN:210000E08B1021EF, port:2
S_ID:010F00

```

LOGIN

Use the LOGIN command to log into a new CLI or Telnet session (such as at a specific security level). You are prompted for a password.

LOGIN PARAMETER

Command Syntax Detail:

PARAMETER:

Description:

<login name>

Indicates which login level the user wishes to log in under.

Example:

The example below shows erroneous logon and a successful logon:

```
RM660[1]: login
Enter a login name: eom
Enter the password:
Sorry.
RM660[1]: login
Enter a login name: oem
Enter the password:
    Successful Telnet session login.
        New owner          : oem.
        New security level: OEM.
```

LOGOUT

Use the LOGOUT command to log out of a CLI or Telnet session.

For the CLI, control is returned to the general purpose user level.

For the Telnet, the current session is disconnected.

NOTE: When you successfully login as another user or at a different security level, you are automatically logged out.

Example:

```
RM660[1]: whoami
```

Telnet session:

```
Current owner      : user.
```

```
Current security level: General User.
```

```
RM660[1]: login
```

```
Enter a login name: admin
```

```
Enter the password:
```

```
Successful Telnet session login.
```

```
    New owner      : admin.
```

```
    New security level: Administrative.
```

```
RM660[1]: login oem
```

```
Enter the password:
```

```
Successful Telnet session login.
```

```
    New owner      : oem.
```

```
    New security level: OEM.
```

LUN

Use the LUN command to add, delete, format, and monitor LUNs in the system.

LUN

LUN PARAMETER

Command Syntax Detail:

PARAMETER:

Description:

CONFIG

Display the configuration information about all the valid LUNs in the system.

LIST

Display a list of all valid LUNs in the system. The list shows the capacity, owner, status and serial number of each LUN.

ADD=x

Creates a new LUN. The system prompts the user for all the necessary information to create the LUN and indicates if the LUN was successfully added to the system. The system can support up to 128 LUNs. Each LUN can be subdivided into up to 64 smaller equally-sized LUN segments. The LUN segments of a LUN are managed together. The LUN to be added can be specified by 'x', where 'x' is in the range 0..127.

DEL

DELETE

DELETE=x

Delete a LUN, 'x', from the system and all of the data in the LUN. The LUN to be deleted can be specified by 'x', where 'x' is in the range 0..127.

FORMAT=x

FORMAT=x.y

Performs a destructive initialization of a LUN by over-writing all the data on the LUN with zeroes. The LUN to be formatted can be specified by 'x', where 'x' is in the range 0..127. A sub LUN in a LUN group can be specified by 'x.y', where 'x' is the LUN group in the range 0..127 and 'y' is the sub LUN of the group in the range 0..63. If all of the LUNs in a LUN group are unformatted then the entire LUN group will be formatted even if only one sub LUN was specified.

PAUSE

Pauses all the format, verify and mirror operations.

RESUME

Releases all the paused format, verify and mirror operations.

STOP

Aborts all active format, verify and mirror operations. Note that it does not, however, change the Background Verify setting for the affected LUN(s). Use 'LUN VERIFY=OFF' to change the Background Verify setting for a LUN.

LOCK=x

Locks a LUN in the data cache, keeps all of the data for the LUN in the cache for faster access. Up to 50%% of the data cache can be used for locking LUNs. The LUN to be locked can be specified by 'x', where 'x' is in the range 0..127.

UNLOCK=x

Unlocks a LUN and release the cache locked by the LUN. The LUN to be unlocked can be specified by 'x', where 'x' is in the range 0..127.

MOVE=x

Changes the ownership of a LUN from one unit to another when the units are in dual mode. Overrides all the configuration checks and allow the user to move any LUN. The system displays a list of all the tiers and other LUNs that need to be moved in order to move the LUN specified. The LUNs should not be moved while any unit in the system has active format, verify, rebuild or mirror operations. The LUN to be moved can be specified by 'x', where 'x' is in the range 0..127.

LABEL=x

Changes the label of the LUN. A LUN label can be up to 12 characters long. The LUN to be labeled can be specified by 'x', where 'x' is in the range 0..127.

DELAY=x

Sets the system verify delay value to 'x'. The verify delay value determines how long a verify operation pauses 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).

Note that DELAY=0 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. The default is 30.

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 Mbytes. Default is 16 Mbytes.

VERIFY

Displays the current background verify settings for all LUNs in the system.

VERIFY=x**VERIFY=x.y**

Turns ON background verify for LUN 'x', where 'x' is in the range 0..127. A sub LUN in a LUN group can be specified by 'x.y', where 'x' is the LUN group in the range 0..127 and 'y' is the sub LUN of the group in the range 0..63.

VERIFY=ON

Prompts the user for a list of LUNs on which background verify will be turned either ON. The 'VERIFY=ON' argument will turn on background verify for the specified LUN(s), optionally running in continuous mode.

VERIFY=OFF

Prompts the user for a list of LUNs on which background verify will be turned OFF. The 'VERIFY=OFF' command, however, only turns off the Background Verify setting for the specified LUN(s). Therefore, any verifies 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 'LUN STOP'.

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. The command LUN RELEASE can be used to release any SCSI reservations on a LUN.

RELEASE=x**RELEASE=x.y**

Releases 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..127. A sub LUN in a LUN group can be specified by 'x.y', where 'x' is the LUN group in the range 0..127 and 'y' is the sub LUN of the group in the range 0..63.

START

Starts all the LUNs that have been stopped by a SCSI START/STOP request. This parameter is not related to the STOP parameter.

Example:

RM660[1]: lun

Logical Unit Status									
LUN	Label	Owner	Status	Capacity (Mbytes)	Block Size	Tiers	Tier	list	
0		1	Critical	[GHS 30000	512	4	1	2	3 4
1		1	Critical	1800	512	1	5		
2		1	Critical	30000	512	1	5		

System verify extent: 128 Mbytes

System verify delay: 10

System Capacity 6720588 Mbytes, 6658788 Mbytes available.

MIRROR

Use the MIRROR command to display, create, edit, and destroy IMGs from the system, as well as breaking, merging, adding and deleting LUNs within the IMG.

MIRROR

MIRROR parameter

Command Syntax Detail:

PARAMETER:

Description:

CREATE

Creates a new IMG in the system. The system prompts the user for all the necessary information to create the IMG and indicates if the IMG was successfully added to the system.

DESTROY

Permanently deletes an IMG from the system. This deletes only the IMG itself, but not the data on the individual LUNs. NOTE: All Backup LUNs in the IMG must be broken before a delete can occur.

ADD

Adds a new Backup LUN to the IMG. All data on the Backup LUN is destroyed and then it is data synced with the Primary LUN.

DEL

DELETE

Deletes a broken Backup LUN from the IMG. All data on the Backup LUN is preserved.

BREAK

Breaks a Backup LUN from the IMG where the Backup LUN will no longer be mirroring the Primary LUN. All data on the Backup LUN is preserved.

MERGE

Joins a broken Backup LUN back to the IMG, by copying all of the Primary LUNs data onto the Backup.

PAUSE
Pauses the merge operations.

RESUME
Releases the paused merge operations.

STOP
Abort the merge operations.

Usage Guidelines:

Parentheses surrounding a Backup LUN is an indication that the LUN is broken from the IMG.

The MERGE parameter can be used to add the Backup LUN back into the IMG allowing it to mirror the Primary LUN once again.

All data from the Primary LUN is copied over onto the Backup LUN so they are again identical.

Example:

```
RM660[1]: mirror
```

```

                                Internal Mirrored Groups (IMG)
      Primary                    Backup LUNs
IMG  LUN  1st  2nd  3rd  4th  5th  6th  7th Owner Status
-----
                                -----No IMGs found-----
Verifying LUN 0, 91.9%, Continuous.
Verifying LUN 1, 78.5%, Continuous.
Verifying LUN 2, 77.7%, Continuous.
                                System Capacity 560040 Mbytes, 0 Mbytes available.
```

NETWORK

Use the NETWORK command to display the current network settings of the system and use the appropriate parameter(s) to change the settings.

network

network parameter

Command Syntax Detail:

PARAMETER:

Description:

PING=aaa.bbb.ccc.ddd

Attempts to ping the specified destination with a single packet.

PRIVATE

Displays the MAC address for the private network device.

IP=aaa.bbb.ccc.ddd

Changes the IP address of the system. The system must be restarted before the changes will take effect.

NETMASK=aaa.bbb.ccc.ddd

Changes the netmask of the system. The system must be restarted before the changes will take effect.

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 that GATEWAY= (with no Internet address) will clear out the current gateway.

TELNET=ON

Specifies that the Telnet capability is active. The system must be restarted before the changes will take effect. Note: To affect telnet session availability only temporarily during the current power-cycle, refer to the TELNET ENABLE/DISABLE command.

TELNET=OFF

Specifies that the Telnet capability is not active. The system must be restarted before the changes will take effect. Note: To affect telnet session availability only temporarily during the current power-cycle, refer to the TELNET ENABLE/DISABLE command.

TELNETPORT=x

Changes the Telnet port number for this unit to that specified by 'x'. The system must be restarted before the changes will take effect. Valid ports are 0 to 32768. Note however, that the results may be unpredictable if the port number chosen is already in use (on this unit) by either the GUI or SYSLOG facilities. Default port is 23.

API_SERVER=ON

Specifies whether or not the API Server capability is active. The RSM controller manager interface 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, refer to the API ENABLE/DISABLE command.

API_SERVER=OFF

Specifies whether or not the API Server capability is active. The RSM controller manager interface 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, refer to the API ENABLE/DISABLE command.

API_PORT=x

Changes the API Server port number for this unit to that specified by 'x'. The system must be restarted before the changes will take effect. Valid ports are 0 to 32768. Note however, that the results may be unpredictable if the port number chosen is already in use (on this unit) by either the TELNET or SYSLOG facilities. Default port is 8008.

SNMP=ON

Specifies the SNMP functionality is active. The system must be restarted before the changes will take effect.

SNMP=OFF

Specifies the SNMP functionality is inactive. The system must be restarted before the changes will take effect.

LIMIT_SNMP=ON

Specifies the SNMP functionality will only report component-level information, or all levels of information.

LIMIT_SNMP=OFF

Default setting, that specifies the SNMP functionality will not report component-level information, or all levels of information. Default is OFF.

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.

SYSLOG=ON

Specifies the Syslog capability is active.

SYSLOG=OFF

Specifies the Syslog capability is inactive.

SYSLOGIP=aaa.bbb.ccc.ddd

Changes the destination IP address for syslog packets. If the unit is coupled with a second unit and setup for dual mode, both units in the system will share the same syslog destination IP address but each unit can specify a different destination port.

SYSLOGPORT=x

Changes the destination port number for syslog packets for this unit. If the unit is coupled with a second unit and setup for dual mode, both units in the system will share the same syslog destination IP address but each unit can specify a different destination port. Valid ports are 0 to 32768. Note however, that 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 is 514.

USAGE

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

Example:

```
RM660[1]: network
```

```
Network Configuration
=====
```

```
Gateway:
```

```
-- None --
```

```
NetMask:                255.255.255.0
MAC Address RM660 #1:    00:01:ff:03:02:1e
IP Address  RM660 #1:    192.168.0.168
IP Address  RM660 #2:    192.168.0.169

Services
-----
Telnet:                 ENABLED
Telnet Port RM660 #1:   23
Telnet Port RM660 #2:   23

API Server:             ENABLED
API Server Port RM660 #1: 8008
API Server Port RM660 #2: 8008

SNMP:                   DISABLED
SNMP Trap IP Address:   0.0.0.0
Limited SNMP:          ISABLED

Syslog:                 ENABLED
Syslog IP Address:      192.168.0.88
Syslog Port RM660 #1:   514
Syslog Port RM660 #2:   514
```

PASSWORD

Use the PASSWORD command to change the CLI's (and Telnet's) administrative and general purpose login name(s) and password(s). Login names may be no more than 10 characters in length. Passwords must be between 8 and 20 characters in length.

PASSWORD parameter

Command Syntax Detail:

PARAMETER:

Description:

DEFAULTS

Reset all the login names and passwords to their default values.

RESTART

Use the RESTART command to perform a restart of the system. 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.

RESTART parameter

Command Syntax Detail:

PARAMETER:

Description:

DELAY=x

Performs a restart of the unit in 'x' minutes, any number between 0 and 255.

DUAL

Restarts both this unit and the other unit.

KILL

Stops a timed restart that is in progress.

ROUTE

Use the ROUTE command to display and update the current IP routing table of the unit.

ROUTE parameter

Command Syntax Detail:

PARAMETER:

Description:

ADD=aaa.bbb.ccc.ddd GATEWAY=aaa.bbb.ccc.ddd

Adds a gateway to the network routing tables, where 'aaa.bbb.ccc.ddd' represents a standard Internet address. Up to 6 permanent routes can be added to the routing table

DEL

DELETE=aaa.bbb.ccc.ddd GATEWAY=aaa.bbb.ccc.ddd

Deletes a gateway from the network routing table, where 'aaa.bbb.ccc.ddd' represents a standard Internet address.

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. If an empty gateway value is provided, then the current gateway is cleared.

Example:

To indicate that the machine with Internet address 91.0.0.3 is the gateway to the destination network 90.0.0.0:

```
RM660[1]: ROUTE ADD=90.0.0.0 GATEWAY=91.0.0.3
```

To display the current IP routing table:

```
RM660[1]: route
```

```
Gateway:  -- None  --
```

```
Permanent Routing Table:
```

```
=====
```

```
destination      gateway
```

```
-----  
-- None --      -- None --  
-----  
  
Current Routing Tables:  
=====
```

ROUTE NET TABLE					
destination	gateway	flags	Refcnt	Use	
Interface					
192.168.0.0	192.168.0.168	101	0	0	fei0

ROUTE HOST TABLE					
destination	gateway	flags	Refcnt	Use	
Interface					
127.0.0.1	127.0.0.1	5	0	0	lo0

SAVE

Use the SAVE command to save the parameter blocks; the system configuration is saved to non-volatile memory. Backup copies of the system configuration are also saved on the disks. The system automatically saves and updates the backup copies when changes are made to the system configuration or status.

SAVE parameter

Command Syntax Detail:

PARAMETER:

Description:

STATUS

Displays the current status of the system parameters.

Example:

To display the current status:

```
RM660[1]: RM660[1]: save status
```

```

                System Parameters Status
Number    Revision    Updates    Last Update
    0         500      94796      WED AUG 11 14:33:25 2004
    1         500       2021      FRI AUG 06 15:11:45 2004

```

SES

Use the SES command to display a list of all current enclosure faults detected by the SCSI Enclosure Services (SES) as well as provide a means to access SES specific functions such as disk, channel, lun, or tier visual identification.

SES parameter

Command Syntax Detail:

PARAMETER:

Description:

ON

Save the SES state to the parameter blocks, and startup the SES monitors.

OFF

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

SHOWDEVICES

Display all the SES devices on all channels.

SHOWALL

Display all configuration information for all the SES devices on all channels.

SHOW=tc

Display the configuration information and the status information returned from an SES Enclosure Status page for the SES device for the specified drive:

tier (t); valid range <1..125>.

channel (c) valid range <ABCDEFGHPS>.

IDDISK=tc

Commands the system to provide a visual indication of the specified drive, 'tc', where 't' is the tier and 'c' is the channel. The valid range for 't' is <1..125>. The valid range for 'c' is <ABCDEFGHPS>.

IDTIER=t

Commands the system to provide a visual indication of the specified tier, 't'. The valid range for 't' is <1..125>.

IDCHANNEL=c

Commands the system to provide a visual indication of the specified channel, 'c'. The valid range for 'c' is <ABCDEFGHPS>

IDLUN=n

Commands the system to provide a visual indication of the specified LUN, 'n'. The valid range for 'n' is <0..the max. number of LUNs> (which is set at 128 in the current system).

ID=OFF

Clears the visual indication task that is currently in progress and restores the system to its original visual state.

M_WAIT=x

Sets the SES device monitoring rate for the system to 'x'. The monitoring rate is given in seconds. The valid range for 'x' is <4..90>. The default monitoring rate value is 6 seconds.

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

M_WAIT

Displays the current value for the SES device monitoring rate. The value is given in seconds.

VERIFY_6016

Verifies the connections of a 6016 drive enclosure that is currently attached to the system. The connections must be verified before the SCSI Enclosure Services (SES) can accurately indicate a drive fault through SES commands.

VERIFY_6012REM

Verifies the connections of a SF6012REM drive enclosure that is currently attached to the system. The connections must be verified before the SCSI Enclosure Services (SES) can accurately indicate a drive fault through SES commands.

RESET_VERIFIED_ENCLOSURES

Removes all of the verified drive enclosures from the verified enclosures list for this controller.

Example:

To show all configuration information for all the SES devices on all channels:

```
RM660[1]: ses showall
```

Information about SES module at 1A

Timestamp : 09:29:25
Number of sub enclosures : 0
Sub enclosure ID : 0
Number of elements supported : 7
Enclosure WWN : 20000050CC00BDC9
Enclosure Vendor ID : XYRATEX
Product Identification : RS1600-FC2
Product Revision : 38
Supported elements :
Device : 16 00
Power Supply : 02 00
Cooling Element : 02 00
Temperature Sensor : 01 00
Audible Alarm : 01 00
SES Encl Srvc Electronics : 02 00
Display : 01 00

Information about SES module at 1B

Timestamp : 09:29:21
Number of sub enclosures : 0
Sub enclosure ID : 0
Number of elements supported : 7
Enclosure WWN : 20000050CC00BDC9
Enclosure Vendor ID : XYRATEX
Product Identification : RS1600-FC2
Product Revision : 38
Supported elements :
Device : 16 00
Power Supply : 02 00
Cooling Element : 02 00
Temperature Sensor : 01 00
Audible Alarm : 01 00
SES Encl Srvc Electronics : 02 00
Display : 01 00

Information about SES module at 1C

Timestamp : 09:29:26
Number of sub enclosures : 0
Sub enclosure ID : 0
Number of elements supported : 7
Enclosure WWN : 20000050CC00BEB9
Enclosure Vendor ID : XYRATEX
Product Identification : RS1600-FC2

```

Product Revision                : 38
Supported elements :
Device                          : 16 00
Power Supply                    : 02 00
Cooling Element                 : 02 00
Temperature Sensor              : 01 00
Audible Alarm                   : 01 00
SES Encl Srvc Electronics      : 02 00
Display                         : 01 00

```

Information about SES module at 1D

```

Timestamp                       : 09:29:21
Number of sub enclosures        : 0
Sub enclosure ID                : 0
Number of elements supported    : 7
Enclosure WWN                   : 20000050CC00BEB9
Enclosure Vendor ID             : XYRATEX
Product Identification          : RS1600-FC2
Product Revision                : 38
Supported elements :

```

```

Device                          : 16 00
Power Supply                    : 02 00
Cooling Element                 : 02 00
Temperature Sensor              : 01 00
Audible Alarm                   : 01 00
SES Encl Srvc Electronics      : 02 00
Display                         : 01 00

```

Information about SES module at 1E

```

Timestamp                       : 09:29:23
Number of sub enclosures        : 0
Sub enclosure ID                : 0
Number of elements supported    : 7
Enclosure WWN                   : 20000050CC00BECB
Enclosure Vendor ID             : XYRATEX
Product Identification          : RS1600-FC2
Product Revision                : 38
Supported elements :

```

```

Device                          : 16 00
Power Supply                    : 02 00
Cooling Element                 : 02 00
Temperature Sensor              : 01 00
Audible Alarm                   : 01 00
SES Encl Srvc Electronics      : 02 00
Display                         : 01 00

```

Information about SES module at 1F

Timestamp : 09:29:21
Number of sub enclosures : 0
Sub enclosure ID : 0
Number of elements supported : 7
Enclosure WWN : 20000050CC00BECD
Enclosure Vendor ID : XYRATEX
Product Identification : RS1600-FC2
Product Revision : 38
Supported elements :
Device : 16 00
Power Supply : 02 00
Cooling Element : 02 00
Temperature Sensor : 01 00
Audible Alarm : 01 00
SES Encl Srvc Electronics : 02 00
Display : 01 00

Information about SES module at 1G

Timestamp : 09:29:24
Number of sub enclosures : 0
Sub enclosure ID : 0
Number of elements supported : 7
Enclosure WWN : 20000050CC00BD95
Enclosure Vendor ID : XYRATEX
Product Identification : RS1600-FC2
Product Revision : 38
Supported elements :
Device : 16 00
Power Supply : 02 00
Cooling Element : 02 00
Temperature Sensor : 01 00
Audible Alarm : 01 00
SES Encl Srvc Electronics : 02 00
Display : 01 00

Information about SES module at 1H

Timestamp : 09:29:25
Number of sub enclosures : 0
Sub enclosure ID : 0
Number of elements supported : 7
Enclosure WWN : 20000050CC00BD95
Enclosure Vendor ID : XYRATEX
Product Identification : RS1600-FC2

```
Product Revision                : 38
Supported elements :
Device                          : 16 00
Power Supply                    : 02 00
Cooling Element                 : 02 00
Temperature Sensor              : 01 00
Audible Alarm                   : 01 00
SES Encl Srvc Electronics       : 02 00
Display                         : 01 00
```

Information about SES module at 1P

```
Timestamp                       : 09:29:24
Number of sub enclosures        : 0
Sub enclosure ID                : 0
Number of elements supported    : 7
Enclosure WWN                   : 20000050CC00BD77
Enclosure Vendor ID             : XYRATEX
Product Identification          : RS1600-FC2
Product Revision                : 38
Supported elements :
```

```
Device                          : 16 00
Power Supply                    : 02 00
Cooling Element                 : 02 00
Temperature Sensor              : 01 00
Audible Alarm                   : 01 00
SES Encl Srvc Electronics       : 02 00
Display                         : 01 00
```

Information about SES module at 1S

```
Timestamp                       : 09:29:25
Number of sub enclosures        : 0
Sub enclosure ID                : 0
Number of elements supported    : 7
Enclosure WWN                   : 20000050CC00BD77
Enclosure Vendor ID             : XYRATEX
Product Identification          : RS1600-FC2
Product Revision                : 38
Supported elements :
```

```
Device                          : 16 00
Power Supply                    : 02 00
Cooling Element                 : 02 00
Temperature Sensor              : 01 00
Audible Alarm                   : 01 00
SES Encl Srvc Electronics       : 02 00
Display                         : 01 00
```

SETTINGS

Use the SETTINGS command to display and change the CLI's and Telnet's various session control settings.

SETTINGS parameter

Command Syntax Detail:

PARAMETER:

Description:

DEFAULTS

Resets all the CLI and telnet session control settings to the default values.

LINES=x

Displays or 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. The user is 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 entirely.

Valid range is 0 to 512 lines, where 0 indicates that no paging is to be performed on the output information.

Default is 0.

Example:

This example shows the current session control settings:

```
RM660[1]: settings defaults
```

```
RM660[1]: settings
```

```
Current Session Control Settings
```

```
-----
```

```
Lines per page:    10
```

```
RM660[1]:
```

This example resets the defaults:

```
RM660[1]: settings
```

```
Current Session Control Settings
```

```
-----
```

```
Lines per page:    0 (No paging - continuous scroll)
```

SHOWALL

Use the SHOWALL command to display the output from a collection of essential CLI commands that, taken all together, present a comprehensive view of the system's current configuration. In order, this command shows the results from each of the following commands:

VERSION, VERSION AVR, TIER, TIER CONFIG, LUN, LUN CONFIG, CACHE, AV, DISK, DISK LIST, DUAL, HOST, HOST STATUS, NETWORK, STATS, STATS DELAY, STATS TIERDELAY, LOG.

Example:

```
RM610[2]: showall
```

```
System Configuration Information Summary:
```

```
Version Information:
```

```
-----
```

```
Silicon Graphics RM610
Firmware Version: 5.10
```

```
Firmware date: Sep  3 2004, 16:19:37
IEEE ULA Number: 000202B8
Bootrom Version: 1.08
Platform: SGI RM610
```

```
Port 1:  Chip Rev: 1  Firmware: 1.61
Port 2:  Chip Rev: 1  Firmware: 1.61
Port 3:  Chip Rev: 1  Firmware: 1.61
Port 4:  Chip Rev: 1  Firmware: 1.61
```

```
Channel A:  Chip Rev: 1  Firmware: 1.54
Channel B:  Chip Rev: 1  Firmware: 1.54
Channel C:  Chip Rev: 1  Firmware: 1.54
Channel D:  Chip Rev: 1  Firmware: 1.54
Channel P:  Chip Rev: 1  Firmware: 1.54
Channel S:  Chip Rev: 1  Firmware: 1.54
```

```
RCM API Version: 1.20
```

```
AVR Version Information:
```

```
-----
```

04-00151-102

ATmega163

3.6864MHZ

01.80.00

11.18.2003

System Uptime (YY:DDD:HH:MM): 00:047:06:00

Tier Information:

Tier	Owner	Capacity (Mbytes)	Space Available (Mbytes)	Disk Status	LUN List
1	1	560048	48	ABCDPS	0 1
2	1	560048	48	ABCDPS	2 3
3	2	560048	48	ABCDP	4 5
4	2	560048	48	ABCDP	6 7
5		560048	560048	ABCDP	
6		560048	560048	ABCDP	

All disks are healthy.

Automatic disk rebuilding: Enabled
 System rebuild extent: 32 Mbytes
 System rebuild delay: 30

System Capacity 3360288 Mbytes, 1120288 Mbytes available.

Data compatibility mode is enabled (default).
 This system is fully compatible across all platforms.

Tier Configuration

Tier	Owner	Total LUNs	Disk Status	Installed Disks	Healthy Disks	F R	Sp H Sp A	Spare Owner	Spare Used on	Repl Spare from
1	1	2	ABCDPS	ABCDPS	ABCDP		Y Y			

2	1	2	ABCDPS	ABCDPS	ABCDP	Y	Y
3	2	2	ABCDP	ABCDP	ABCDP		
4	2	2	ABCDP	ABCDP	ABCDP		
5		0	ABCDP	ABCDP	ABCDP		
6		0	ABCDP	ABCDP	ABCDP		

All disks are healthy.

Automatic disk rebuilding: Enabled

Other unit host busy: 0.

Other unit disk busy: 0.

Current mapping mode: 0: 'Standard Enclosure'

LUN Information:

LUN	Label	Owner	Status	Capacity (Mbytes)	Block Size	Tiers	Tier list
0	0	1	Ready	280000	512	1	1
1	1	1	Ready	280000	512	1	1
2	2	1	Ready	280000	512	1	2
3	3	1	Ready	280000	512	1	2
4	4	2	Ready	280000	512	1	3
5	5	2	Ready	280000	512	1	3
6	6	2	Ready	280000	512	1	4
7	7	2	Ready	280000	512	1	4

System verify extent: 16 Mbytes

System verify delay: 30

System Capacity 3360288 Mbytes, 1120288 Mbytes available.

Logical Unit Configuration

LUN	Capacity (Blocks)	Block Size	LUN Offset	Tier Start	Tier End	Tier list

D: Complete List of Commands at OEM Level

```

0 222E0000 512 0 0 88B7FFF 1
1 222E0000 512 0 88B8000 1116FFFF 1
2 222E0000 512 0 0 88B7FFF 2
3 222E0000 512 0 88B8000 1116FFFF 2
4 222E0000 512 0 0 88B7FFF 3
5 222E0000 512 0 88B8000 1116FFFF 3
6 222E0000 512 0 0 88B7FFF 4
7 222E0000 512 0 88B8000 1116FFFF 4

```

System Capacity 3360288 Mbytes, 1120288 Mbytes available.

LUNs currently being Verified

LUN Time	Label	Continuous	Progress	Passes	Last Verify	Completion
-------------	-------	------------	----------	--------	-------------	------------

```

-----
0 0 No 0 0
1 1 No 0 0
2 2 No 0 0
3 3 No 0 0
4 4 No 0 0
5 5 No 0 0
6 6 No 0 0
7 7 No 0 0

```

There are NO LUNs being verified currently.

System verify extent: 16 Mbytes
System verify delay: 30

Cache Information:

```

-----
Current Cache settings
LUN Write Maximum MF Prefetch Read Write
Caching Prefetch Bit Ceiling Priority Priority
-----
0 Disabled x 1 On 65535 high high
1 Disabled x 1 On 65535 high high

```

2	Disabled	x	1	On	65535	high	high
3	Enabled	x	1	On	65535	high	high
4	Disabled	x	1	On	65535	high	high
5	Disabled	x	1	On	65535	high	high
6	Disabled	x	1	On	65535	high	high
7	Disabled	x	1	On	65535	high	high

Writeback Limit: 75%

1536.0 Mbytes of Cache Installed
(8192 Segments of 128 Kbytes)

Total Cache Nodes: 8192
Allocation Size: 20316160, Free: 18895340 93.0%

Dual Nodes Required: 8509
Dual Nodes Allocated: 8509 100.0%
Dual Nodes Active: 0 0.0%

Writeback Nodes: 0 0.0%
Rebuild Nodes: 0 0.0%
Verify Nodes: 0 0.0%

Locked Nodes: 0 0.0%
Lock Node Limit: 4096 50.0%

Writeback Threshold: 2048 25%
Writeback BDB Limit: 4096 50%
Writeback Limit: 6144 75%

Audio/Visual Information:

Current LUN Audio/Visual settings

LUN	Label	FastAV	Read Continuous	Write Caching	Maximum Prefetch
0	0	Off	Off	Off	x 1
1	1	Off	Off	Off	x 1
2	2	Off	Off	Off	x 1
3	3	Off	Off	On	x 1
4	4	Off	Off	Off	x 1
5	5	Off	Off	Off	x 1
6	6	Off	Off	Off	x 1
7	7	Off	Off	Off	x 1

Disk Audio/Visual settings are: Disabled (Using disk defaults)
Early Error Recovery: Enabled
Automatic Read Reallocation: Disabled
Read Retry Count: 1
Write Retry Count: 1
Recovery Time Limit: 65535

Ordered Tag Count: 0
Unit Attention: Enabled
FASTAV Timeout: 50
RC Timeout: 8
Fail Check Condition: Disabled

Disk Information:

Disk Channel Status

Disk Channel A healthy.
Disk Channel B healthy.
Disk Channel C healthy.
Disk Channel D healthy.
Disk Channel P healthy.
Disk Channel S healthy.

All disks are healthy.

Disk rebuild verify: Disabled
Disk command timeout: 27 seconds
Disk maximum commands: 32
Disk PLOGI: Disabled
Auto Reassign: Enabled

Disk commands outstanding: 0

Disks Installed

	Vendor	Product ID	Mbytes	Rev	Serial Number
Disk 1A	SEAGATE	ST3146807FC	140014	0006	3HY8FP6N
Disk 1B	SEAGATE	ST3146807FC	140014	0006	3HY8EXZC

Disk	1C	SEAGATE	ST3146807FC	140014	0006	3HY8GK8N
Disk	1D	SEAGATE	ST3146807FC	140014	0006	3HY8BB75
Disk	1P	SEAGATE	ST3146807FC	140014	0006	3HY8GK54
Disk	1S	SEAGATE	ST3146807FC	140014	0006	3HY8FQY8
Disk	2A	SEAGATE	ST3146807FC	140014	0006	3HY8GGFK
Disk	2B	SEAGATE	ST3146807FC	140014	0006	3HY8DXAT
Disk	2C	SEAGATE	ST3146807FC	140014	0006	3HY8C4DK
Disk	2D	SEAGATE	ST3146807FC	140014	0006	3HY8DYBE
Disk	2P	SEAGATE	ST3146807FC	140014	0006	3HY8GCB9
Disk	2S	SEAGATE	ST3146807FC	140014	0006	3HY8BZT6
Disk	3A	SEAGATE	ST3146807FC	140014	0006	3HY8FK9K
Disk	3B	SEAGATE	ST3146807FC	140014	0006	3HY8GK27
Disk	3C	SEAGATE	ST3146807FC	140014	0006	3HY8G9KH
Disk	3D	SEAGATE	ST3146807FC	140014	0006	3HY8GJP8
Disk	3P	SEAGATE	ST3146807FC	140014	0006	3HY8GFTE
Disk	4A	SEAGATE	ST3146807FC	140014	0006	3HY8FLK7
Disk	4B	SEAGATE	ST3146807FC	140014	0006	3HY8BVJS
Disk	4C	SEAGATE	ST3146807FC	140014	0006	3HY8DBPB
Disk	4D	SEAGATE	ST3146807FC	140014	0006	3HY8DB72
Disk	4P	SEAGATE	ST3146807FC	140014	0006	3HY8DKLS
Disk	5A	SEAGATE	ST3146807FC	140014	0006	3HY8BKSG
Disk	5B	SEAGATE	ST3146807FC	140014	0006	3HY8C8AW
Disk	5C	SEAGATE	ST3146807FC	140014	0006	3HY8B9AY
Disk	5D	SEAGATE	ST3146807FC	140014	0006	3HY8DAD8
Disk	5P	SEAGATE	ST3146807FC	140014	0006	3HY8DCHN
Disk	6A	SEAGATE	ST3146807FC	140014	0006	3HY8DXAJ
Disk	6B	SEAGATE	ST3146807FC	140014	0006	3HY8D418
Disk	6C	SEAGATE	ST3146807FC	140014	0006	3HY8D5H3
Disk	6D	SEAGATE	ST3146807FC	140014	0006	3HY8DAV7
Disk	6P	SEAGATE	ST3146807FC	140014	0006	3HY8D6S0
Found:	32	SEAGATE	ST3146807FC	140014	0006	

SHUTDOWN

Use the SHUTDOWN command to perform a shutdown of the system.

It prepares the system to be shutdown. All hosts and users actively using this system should be safely shutdown before using this command. The system halts all I/O requests and save the data to the disks. The system also asks if the disks should be spundown. Disks should be spundown before they are moved. The unit can be safely turned off after using this command.

SHUTDOWN parameter

Command Syntax Detail:

PARAMETER:

Description:

RESTART=x

Performs a Hard Restart of the unit by cycling the power, where 'x' indicates the number of seconds before the unit powers up again, any number between 1 and 127. If a time is not specified, the default delay is 15 seconds.

Note: If this parameter is used in conjunction with the DUAL parameter, the RESTART will only affect this unit (and not the other unit).

DELAY=x

Performs a shutdown of the unit in 'x' minutes; any number between 0 and 255.

DUAL

Shutsdown both this unit and the other unit.

KILL

Stops a timed shutdown that is in progress.

STATS

Use the STATS command to display the performance statistics for the host ports, disk channels, and the cache memory, including the read and write performance of each of the host ports.

STATS

STATS parameter

Command Syntax Detail:

PARAMETER:

Description:

DELAY

Displays a histogram of the time it takes for the host and disk I/O requests to complete in 100 msec. intervals.

HOSTDELAY

Displays a histogram of the time delay between when the last data transfer is set ready and the host command completes. The host ready delay information is shown in 100 msec. intervals.

TIERDELAY

TIERDELAY=n

Displays a histogram of the time it takes for the disk I/O requests to complete for all the disks in the specified tier (n). If no tier is specified, then all valid tiers are displayed. The histogram is displayed in 100 msec. intervals.

LENGTH

Displays a histogram of the length of the host I/O requests in 16 Kbyte intervals.

OFFSET

Displays a histogram of the offset of the host I/O requests into the cache segments. 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.

DISK

Displays a histogram of which disks in the system have taken an unusually long time to complete an I/O request. The count is incremented for a disk if the disk took 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 the 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 or it may have problems.

DUAL

Displays the statistics for the dual mode messages.

CLEAR

Resets all the statistics back to zero.

REPEAT=OFF

REPEAT=MBS

REPEAT=IOS

Enables and disables the repeating statistics display, where:

MBS - displays MB/s,

IOS - displays IO/s,

OFF - turns off (both) the repeating displays.

Usage Guidelines:

Refer to the first example below in regards to the following field explanations:

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 because of prefetching.

Prefetches shows the percentage of host read I/O requests to the disks which are due to prefetching.

The read and write performance of the disks is displayed at the bottom of the screen. Disk Pieces shows the total number of disk I/O requests from the host ports. The system combines 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.

BDB Pieces is the number of host I/O blocking deblocking requests.

Cache Writeback Data shows the percentage of the cache which contains writeback data which must be written to the disks.

Cache Rebuild Data shows the percentage of the cache in use for rebuild operations.

Cache Data Locked shows the percentage of the cache which is locked by the locked LUNs.

Example:

RM660[1]: stats

```

                                System Performance Statistics
      All Ports      Port 1      Port 2      Port 3      Port 4
Read MB/s:           0.0         0.0         0.0         0.0         0.0
Write MB/s:          49.4        12.3        12.6        12.2        12.4
Total MB/s:          49.4        12.3        12.6        12.2        12.4

Read IO/s:           0          0          0          0          0
Write IO/s:          47         12         12         11         12
Total IO/s:          48         12         12         12         12

Read Hits:           0.0%        0.0%        0.0%        0.0%        0.0%
Prefetch Hits:      0.0%        0.0%        0.0%        0.0%        0.0%
Prefetches:         0.0%        0.0%        0.0%        0.0%        0.0%
Writebacks:         100.0%       100.0%      100.0%      100.0%      100.0%
Rebuild MB/s:       0.0         0.0         0.0         0.0         0.0
Verify MB/s:        0.0         0.0         0.0         0.0         0.0

      Total      Reads      Writes      Pieces      Reads      Writes
Disk IO/s:         86         0         86         1: 47990142 51554446
Disk MB/s:         52.5        0.0        52.5        2: 12956138 2065934
Disk Pieces: 206995200 110044464 96950736        3: 9490271 75159
BDB Pieces:                54382798        4: 693547 3214067
                                5: 717928 611104
Cache Writeback Data: 5.3%        6: 112751 50032
Rebuild/Verify Data: 0.0%      0.0%        7: 84665 32207
Cache Data locked:  0.0%        8: 4781 3075190

```

Example:

The following shows command length statistics example:

RM660[1]: stats length

```

                                Command Length Statistics
      Port 1      Port 2      Port 3      Port 4
Length      Reads  Writes  Reads  Writes  Reads  Writes  Reads  Writes
Kbytes
> 0         0      0      0      0      0      0      0      0
> 16        0      0      0      0      0      0      0      0
> 32        0      0      0      0      0      0      0      0
> 48        0      0      0      0      0      0      0      0
> 64        0      0      0      0      0      0      0      0
> 80        0      0      0      0      0      0      0      0
> 96        0      0      0      0      0      0      0      0
> 112       0      0      0      0      0      0      0      0
> 128       0      0      0      0      0      0      0      0

```

D: Complete List of Commands at OEM Level

> 144	0	0	0	0	0	0	0	0
> 160	0	0	0	0	0	0	0	0
> 176	0	0	0	0	0	0	0	0
> 192	0	0	0	0	0	0	0	0
> 208	0	0	0	0	0	0	0	0
> 224	0	0	0	0	0	0	0	0
> 240	0	0	0	0	0	0	0	0

TELNET

Use TELNET command to display or to change whether remote Telnet sessions are currently (temporarily) ENABLEd or DISABLEd.

NOTE: This command only provides control over Telnet sessions during the current power cycle. To 'permanently' disable or enable Telnet sessions (i.e. across power-cycles), the user is referred to the NETWORK TELNET=ON | OFF command.

WARNING! The default setting for this command at power-on is ENABLEd. You are 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 DISABLE command, so that remote users cannot log in to the system in the middle of an administrative command.

TELNET parameter

Command Syntax Detail:

PARAMETER:

Description:

STATS

Displays the collected statistics on Telnet sessions.

CLEARSTATS

Clears the collected statistics on Telnet sessions.

DISABLE

Temporarily disables the establishment of remote Telnet sessions during the current power cycle. Users at remote locations will be unable to start a new Telnet session until after a corresponding TELNET ENABLE command is issued. To disable the Telnet functionality completely and maintain that setting through a power-cycle, refer to the NETWORK TELNET=ON/OFF command.

ENABLE

(Re-)enable the establishment of remote Telnet sessions during the current power cycle. Note that basic Telnet functionality for the system must already have been enabled with the NETWORK TELNET=ON/OFF command for this parameter to have any effect.

Example:

```
RM660[1]: telnet
Telnet Sessions are currently -- ENABLED --
```

```
A Telnet Session is currently in progress.
```

```
RM660[1]: telnet stats
```

```
Telnet Session Statistics
```

```
=====
```

	Time	Date
	-----	-----
System Boot-Up Completion	: 03:17:02	08/19/2004
System's Current	: 09:22:02	08/24/2004
Telnet Services Initiation	: 03:16:59	08/19/2004

```
Telnet Sessions have been : -- ENABLED --
since : 03:16:59 08/19/2004
```

```
Telnet Sessions have been :
```

```
Enabled : 1 time.
```

```
Disabled : 1 time.
```

```
Exited gracefully (via LOGOUT) : 22 times.
```

```
Disconnected at the remote site : 6 times.
```

```
KILLED by the administrator : 0 times.
```

```
Telnet session attempts : 30
```

```
SUCCESSFUL Telnet Sessions : 29
```

```
UNsuccessful Telnet session attempts : 1
```

```
Consecutive UNsuccessful Telnet session attempts : 0
```

```
Telnet session was already active : 1
```

```
Telnet sessions were disabled : 0
```

```
User supplied invalid login information : 0
```

```
Other (miscellaneous) : 0
```

TFTP

Use the TFTP command to perform a TFTP update to the unit.

TFTP parameter

Command Syntax Detail:

PARAMETER:

Description:

<IP_address>

Supplies the TFTP server the IP Address from which to transfer a copy of the filename containing the desired software upgrade. Note that the user is prompted for this information if it is not explicitly given. <IP_address> is a string in the form: aaa.bbb.ccc.ddd,

<Filename>

Supplies the TFTP server the IP Address from which to transfer a copy of the filename containing the desired software upgrade. Note that the user is prompted for this information if it is not explicitly given.

<Filename> is a string value.

<RESTART>

If RESTART is entered with all the other parameters after a successful TFTP download, the unit will restart without prompting for any information.

Usage Guidelines:

The unit must be able to enter FAIL mode in order to update firmware without losing data. If the unit is setup for dual mode, cache coherency is disabled.

TIER

Use the TIER command to display the current status and configuration of the tiers in the system.

TIER

TIER parameter

Command Syntax Detail:

PARAMETER:

Description:

MAP

Shows the current tier mapping mode for the disks in the array.

DISPLAYMAP

Shows the current tier mapping for each tier of disks in the array.

CHANGEMAP

Changes the current tier mapping for the disks in the array. This changes the position of the tiers in the system to conform with the layout of different disk modules. This parameter should only be used when the system is first configured. The system must be restarted before the changes will take effect.

AUTOREBUILD=ON

Default setting; enables the automatic disk rebuild when a failed disk is replaced by a new disk. A disk will only be replaced by a spare disk if it fails and AUTOREBUILD is ON.

AUTOREBUILD=OFF

Disables the automatic disk rebuild when a failed disk is replaced by a new disk.

EXTENT=x

Sets the system rebuild extent in Mbytes. The rebuild extent determines how much data can be rebuilt before the rebuild and format operations must pause. This parameter slows down the rebuild and format operations so they will not affect the performance of the system. Increasing the EXTENT will allow more data to be rebuilt in a single pass. The valid range for 'x' is 1..128 Mbytes. Default is 32 Mbytes.

DELAY=x

Sets the system rebuild delay. The rebuild delay determines how long a rebuild or format operation will pause after it reaches the rebuild extent. This parameter slows down the rebuild and format operations so they will not affect the performance of the system. DELAY=0 will remove all delays so the rebuild and format operations will go as fast as possible but this will slow down the performance of the system. This system rebuild delay value is given in 100 millisecond increments. The valid range for 'x' is 0..1000. The default is 30.

PAUSE

Pauses any ongoing rebuild operations. Note that rebuild operations can be started either automatically (see the AUTOREBUILD option), or with the DISK REBUILD command.

RESUME

Releases any paused rebuild operations and allow them to continue. Note that rebuild operations can be started either automatically (see the AUTOREBUILD option), or with the DISK REBUILD command.

STOP

Aborts any ongoing rebuild operations. Note that rebuild operations can be started either automatically (see the AUTOREBUILD option), or with the DISK REBUILD command.

CONFIG

Displays the detailed disk configuration information for all of the tiers. See the general TIER command help description for an explanation of the 'Owner' and 'Disk Status' fields. 'Total LUNs' indicates the number of LUNs which currently reside on the tier. Note that the health indication for the spare channel under the 'Healthy Disks' heading is an indication of the health of the spare disk (if any) which 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. '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. Note that 'RES-#' will appear under this heading while a replacement operation is underway to indicate that unit '#' currently has the spare reserved. 'Spare Used on' indicates the tier (if any) on which 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 on this tier.

Usage Guidelines:

Tiers are automatically added to the system when the disks are detected.

Tier ownership is determined when the first LUN is added to a tier, and is unassigned when the last LUN resident on a tier is removed.

A tier is automatically uninstalled if it is both not in use by any of the LUNs and if all of the disks in the tier are removed or moved to another location.

The status of the disks in each tier is shown as a single character in the range ABCDEFGHPS which indicates a healthy status for the disk on the channel corresponding to that letter, where:

- 'S' indicates the spare channel,
- 'P' indicates the parity channel, and all remaining characters indicate data channels.
- 'e' indicates the disk is served by an enclosure.
- '.' indicates that the disk was failed by the system.
- 'r' indicates that the disk was failed by the system and replaced by a spare disk.
- '!' indicates that the disk has is in the wrong location.
- '?' indicates that the disk has failed the diagnostics tests.
- '*' indicates that the disk is an enclosure device.
- ' ' indicates that the disk is not installed.

The speed of the rebuild and format operations can be adjusted with the DELAY and EXTENT parameters.

Example:

```
RM660[1]:
RM660[1]: tier
```

Tier		Owner	Capacity (Mbytes)	Space Available (Mbytes)	Disk Status	LUN List
1	1		560040	440040	ABCDEFGHIHPS	0 1 2

All disks are healthy.

```
Automatic disk rebuilding: Enabled
System rebuild extent:      128 Mbytes
System rebuild delay:      30
```

System Capacity 560040 Mbytes, 440040 Mbytes available.

TIME

Use the TIME command to display and change the current system time.

TIME parameter

Command Syntax Detail:

PARAMETER:

Description:

hh:mm:ss

Changes the system time to the new value indicated by 'hh:mm:ss', where:

hh: indicates the hour in the range 0..23.

mm: indicates the minute in the range 0..59.

ss: indicates the second in the range 0..59.

Example:

To simply display the system time:

```
RM660[1]: time
```

```
Current time is 8:16:27
```

To change the system time to noon:

```
RM660[1]: TIME 12:00:00
```

```
RM660[1]: time
```

```
Current time is 12:00:00
```

UPTIME

Use the UPTIME command to display the system's Uptime value (also known as Power on Hours) which is the total time the system has been operational since release from the factory.

The Uptime value is displayed as YY:DDD:HH:MM

YY
DDD
HH
MM

Number of years
Number of days
Number of hours
Number of minutes

UPTIME

Command Syntax Detail:

There are no parameters.

Example:

RM660[1]: uptime

System Uptime (YY:DDD:HH:MM): 00:348:06:06

USER

Use the USER command to display and change the system security settings for all the host users accessing the system by specifying their LUN and port mappings:

The LUN mapping indicates which internal LUNs the user has access to and where the internal LUN appears to the user (i.e. as its external LUN).

The Port mapping indicates those ports through which the user may gain access to the unit.

The host users are identified by their 64-bit World Wide Name. The system can store the configuration of up to 512 host users. Any user accessing the system without a World Wide Name entry in the configuration table is considered anonymous and is granted the default LUN zoning for the host port they are connected to. The default LUN zoning for each host port can be changed with the ZONING command.

Known host users may be assigned either their own unique LUN mapping, or the default LUN zoning for the host port they are connected to. A listing of the system's internal LUNs can be displayed with the LUN command.

USER parameter

Command Syntax Detail:

PARAMETER:

Description:

ADD

Adds and configures a new host user for the system. Users who are currently logged in anonymously may be specified simply by their current ID value; otherwise a World Wide Name value is required to identify the new host user.

AUDIT

Displays the current host user auditing setting. When auditing is turned ON, the system displays a message whenever a host user logs into or out. Anonymous users are included.

AUDIT=ON

Enables the host user auditing.

AUDIT=OFF

Disables the host user auditing; default setting.

CONNECTIONS

Displays all the current host and anonymous user connections. Note that only LUNs which currently exist in the system are displayed in the LUN Zoning/Mapping list for each user.

DEL**DELETE****DELETE=x**

Deletes an existing host user, 'x', from the system. If no host user is specified, the user is prompted for one.

EDIT**EDIT=x**

Edits the LUN mapping for an existing host user, 'x'. If no host user is specified, the user is prompted for one.

SHOWMAP**SHOWMAP=x****SHOWMAP=ALL**

Displays the detailed user LUN mapping information for the host user(s) designated by 'x', or ALL. 'x' will display the LUN mapping information for a single host user; ALL will display the LUN mapping information for all host users. If no host user is specified, the user is prompted for one. Note that all LUNs that are mapped for the user(s) are displayed, regardless of whether or not the LUNs currently exist in the system.

Example:

```
RM660[1]: user
```

```
Host Port Users Summary:
```

Unit ID	Port-Map User Name	World Wide Name	1	2	Zoning Method
0	eng_test2	210100E08B2F366F	1234	1234	000,000
1	SUN1	210000E08B0F894E	1234	1234	000,002
2	SUN2	210100E08B2F894E	1234	1234	000,003

```
User auditing is disabled.
```

```
There are 22 users currently connected.
```

VERSION

Use the VERSION command to display the firmware version information about the current version of the system hardware and firmware.

VERSION

VERSION parameter

Command Syntax Detail:

PARAMETER:

Description:

RCM_API

Displays the revision of the RCM API that is currently supported by the system firmware. The RCM API revision number refers to the pertinent revision of the GUI API Interface document.

Example:

The following example displays the firmware version of the system:

```
RM660[1]: version
  Silicon Graphics RM660
  Firmware Version: 5.10

  Firmware date: Sep  3 2004, 16:39:31
  IEEE ULA Number: 00030478
  Bootrom Version: 1.08
```

WHOAMI

Use the WHOAMI command to display the owner of the current CLI or Telnet session.

Example:

```
RM660[1]: whoami
```

Telnet session:

```
Current owner      : oem.  
Current security level: OEM.
```

ZONING

Use the ZONING command to display and change the default LUN zoning for each host port.

ZONING

ZONING parameter

Command Syntax Detail:

PARAMETER:

Description:

DEFAULT

DEFAULT=x

DEFAULT=ALL

Restores the zoning on a host port to the default settings. If no host port is specified, the user is prompted for one.

EDIT

EDIT=x

Edits the default LUN zoning on each host port. If no host port is specified, the user is prompted for one.

Usage Guidelines:

This command is used to change the default LUN zoning for the anonymous users on each host port. LUN zoning indicates which internal LUNs the user has access to and where the internal LUN will appear to the user. The users are identified by their 64-bit World Wide Name. The system can store the configuration of up to 512 users. Any user accessing the system without having a World Wide Name entry in the configuration table is considered anonymous and is only granted the default LUN zoning for the host port they are connected to.

NOTE: Users can be added to and configured for the system using the USER command.

Example:

RM660[1]: zoning

Port Zoning Summary:

Port	World Wide Name	LUN Zoning			
		(External LUN, Internal LUN)			
1	21000001FF03021E	000,000	001,001		
2	22000001FF03021E				
3	23000001FF03021E	000,000	001,001	002,002	003,003
4	24000001FF03021E	004,004			

Glossary

Arbitrated Loop Physical Address.

Each Fibre Channel Arbitrated Loop can accommodate up to 126 devices. Every device on the loop must have a unique AL_PA. The RM610/RM660 uses an address value of 1.

ASCII

American Standard Code for Information Interchange. A 7-bit binary code (0's, 1's) used to represent letters, numbers, and special characters such as \$,!, and /. Supported by almost every computer and terminal manufacturer.

ATA (Advance Technology Attachment)

A disk drive interface standard based on a 16-bit bus and dealing with the power and data signal interfaces between the motherboard and the integrated disk controller and drive. The ATA "bus" only supports two devices - master and slave.

Attribute

Setting that controls access to a specific file. Often used to protect important files (such as the Registry files) from accidental change or deletion. Set using the ATTRIB command in MS-DOS.

Backplane

A printed circuit board incorporated in the chassis assembly to provide logic level signal, and low voltage power distribution paths.

Bay

The slot that a unit or media device fits into.

Byte

A group of binary digits stored and operated upon as a unit. A byte may have a coded value equal to a character in the ASCII code (letters, numbers), or have some other value meaningful to the computer. In user documentation, the term usually refers to 8-bit units or characters.

1 kilobyte (K) is equal to 1,024 bytes or characters; 64K indicates 65,536 bytes or

characters.

Cable

Throughout this SA2016 System user guide this term is used in accordance with the preferred US context of: “an insulated flexible electric wire used for the transmission of data signals between computer equipment.”

Note: Cable is UK preferred terminology for either a power cord or a data cable:

Cache Memory

Cache memory is memory on the RM610/RM660 used for intermediate storage of read and write data. By using cache, you can increase system performance because the data for a read from the host may already be in the cache from a previous operation (thus eliminating the need to access the drive itself), and a write operation is completed once it is written to the cache, rather than to the drives. When you create a logical unit, you can specify various caching parameters for the LUN's. If you need to change any caching parameters after LUN creation, use the `CACHE` command. The following cache options are available:

- **Write Caching:** Allows write operations from the host to be stored in the RM controller's cache memory. The use of write caching increases overall performance because a write operation from the host is completed when data is put in the cache, instead of when the data is actually written to the drive itself.
- **Write Cache Coherence:** Allows both units in a couplet configuration to simultaneously access the same drives. This enables users connected to either RM controller unit to read/write any LUN in the system.

Cache Segment Size

See *Segment Size*.

Channel

Each RM controller uses ten independent drive channels to manage data distribution: eight data channels (A through H), one parity channel (P), and one optional spare channel (S). See also *Tier*.

Character

A representation, coded in binary digits, of a letter, number, or other symbol.

Characters Per Second

A data transfer rate generally estimated from the bit rate and the character length. For

example, at 2400 bps, 8-bit characters with Start and Stop bits (for a total of ten bits per character) will be transmitted at a rate of approximately 240 characters per second (cps).

Chassis

A sheet metal enclosure incorporating a Backplane PCB, an integral *Ops Panel* and a module runner system. The chassis contains a number of 'Bays', each of which can accommodate a plug in module. There are sixteen *drive* carrier bays at the front and five bays at the rear which house *power supply/cooling* and *SCM I/O modules*.

Configure

To set up a hardware device and its accompanying software.

Couplet RM Controller

Two RM controller units can be set up as couplet RM controller. That is, if one RM controller fails, the other RM controller in the pair takes over the failed RM controller's functions, and the RM controller continues to operate. You can then replace the failed RM controller, often without shutting down the RM controller, to resume normal operation.

This feature involves more than just the RM controller. It concerns the entire data connection, the route data takes from the host system to the RM controller. If any part of the connection fails (for example, if the cable connecting the two units fails), the RM controller redundancy feature reroutes I/O to the remaining good connection.

The couplet RM controller feature will control the data flow to the RM controller pairs independent of the operating system. This function keeps track of the current status of the connections and can perform the switch-over without any changes in the operating system.

Each RM controller is assigned specific LUN's to service. If one RM controller fails, the other takes over the failed RM controller's assigned LUN's.

Data Communications

A type of communications in which computers and terminals are able to exchange data over an electronic medium.

Disk (drive, carrier, module)

A SATA disk drive mounted in a carrier. You can have up to sixteen disk drive carrier modules in each SA2016 System enclosure.

Enclosure

The chassis assembly which houses the plug-in modules that make up the SA2016 System storage subsystem.

ESI/Ops module

A unit used to monitor and control all elements of the Enclosure. The **ESI/Operators (Ops)** panel is supplied as an integral part of the RS-1602 series Enclosure core product

Hot plugging

A device with the capability of being connected to a subsystem without interrupting the power supplies to that subsystem.

Hot Spare

A hot spare is a drive containing no data that acts as a standby in case a drive fails in a tier. The hot spare drive adds another level of redundancy to your disk array. If a drive fails, the hot spare takes over for the failed drive until you replace it. Once you replace the failed drive, data is copied from the spare drive to the replacement drive, the hot spare then returns to a standby status. A hot spare drive is not dedicated to a specific tier, but instead can be used for any failed drive in the disk array with the same or smaller capacity. The hot spare drives on the RM controller are global hot spares.

Hot swap

Hot swapping is the term used for manually swapping a failed disk unit with a replacement while the SA2016 System subsystem is in normal use.

Hz (Hertz)

A frequency measurement unit used internationally to indicate cycles per second.

Initialize

To prepare a hardware device for use.

LED

Light Emitting Diode. A small light displayed on the cabinet, disk units and power supply units.

Logical Unit

A logical unit (called a “LUN” for logical unit number) is the basic structure you create on the disk array to retrieve and store your data. On the RM controller, a LUN can be as small as part of a tier (see definition of Tier on page page 423) or as big as the whole disk array. Each logical unit is seen by the host operating system as *one* drive. The host system accesses each logical unit in the same way it accesses a new disk drive, and you must follow the same procedures on the LUN (creating file systems, volumes, and so on) that you would use to access a standard disk drive.

Each LUN has several properties:

- LUN number - This is a number (from 0 to 127, for up to 128 LUNs in total) assigned to the LUN when you create the LUN. It is an internal number that is seen by the RM controller.
- Number of Segments - Each LUN can be subdivided into up to 64 equally-sized segments, giving a total of 8192 LUN segments. The LUN segments of a LUN are managed together and share the same characteristics.
- Capacity - Each LUN has a fixed amount of space. The capacity is set when you create the logical unit.
- Number of tiers - LUNs can be created on one or more tiers. The number of tiers is determined when you create the logical unit.
- Tier selection - LUN can be created on any tier(s) and can be specified when you create the logical unit.
- Block size
- LUN label - which can contain up to 12 characters

LRC

Loop Resiliency Circuit. Circuits within the I/O modules which provide loop resiliency in the event of a drive failing or being unplugged.

LUN

Logical Unit Number. *See Logical Unit.*

Module (power supply, drive, I/O)

A module is a power supply, disk drive or electronics unit held in a carrier that plugs into a bay inside the enclosure. An SA2016 System enclosure can contain sixteen **drive** modules, two **power supply/cooling modules** and one **SCM I/O** modules.

Parallel Transmission

The transfer of data characters using parallel electrical paths for each bit of the character, for example, 8 paths for 8-bit characters. Data is stored in computers in parallel form, but may be converted to serial form for certain operations. *See Serial Transmission.*

Parity

Parity is additional information stored along with the data that allows the RM controller to reconstruct lost data.

Power Cord

Throughout this SA2016 System user guide this term is used in accordance with the preferred US context of: “an insulated flexible electric wire fitted with connectors at each end and used for the transmission of electrical power to computer equipment.

Protocol

A system of rules and procedures governing communications between two or more devices. Protocols vary, but communicating devices must follow the same protocol in order to exchange data. The format of the data, readiness to receive or send, error detection and error correction are some of the operations that may be defined in protocols.

RAID

Redundant Array of Independent Drives.

Rebuild

Rebuild is the process of using data and parity on all operational drives within a tier to regenerate the data on a replacement drive or a hot spare drive.

Rebuild occurs when you replace a failed drive in a degraded logical unit (a degraded logical unit has suffered a drive failure but is still operable). The rate of rebuild is determined by the `TIER DELAY` setting. The setting defines how long to wait before rebuilding the next block of data. The smaller the delay, the faster the rebuild occurs, but the slower system I/O is serviced.

Reconstruction

See *Rebuild*.

Redundant

Not essential.

SAN

Storage Area Network is a computer network dedicated to data storage. It uses Fibre Channel transport and switch technologies to provide new storage capabilities, including server clustering and scalable storage expansion.

SCM I/O module (Serial ATA Control I/O module)

A plug-in module providing FC-AL channel external cable interface with 16 (Serial or Parallel) *ATA drives*.

Segment Size

Segment size is the size of a cache node. It is given in KBytes. You can adjust the performance of the system by changing the cache segment size to match the size of the host I/O requests. 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 the best performance, the cache segment size should be larger than the average host I/O request size.

Serial Transmission

The transfer of data characters one bit at a time, sequentially, using a single electrical path. See *Parallel Transmission*.

SES

“SCSI Enclosure Services” is an industry-standard protocol used to evaluate and detect a number of different environmental conditions including the presence of drives in the enclosure, to manage fan control, and to measure usage status.

Tier

Tiers are the basic building blocks of the RM controller. One tier contains up to ten drives: four or eight data drives (Channels A through H), one parity drive (Channel P), and an optional spare drive (Channel S). Drives that have the same AL_PA across all ten (or six) channels are put on the same tier. The tier that contains the drives with the highest AL_PA value is recognized as Tier #1. You may connect up to 125 tiers to the RM controller. 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.

Index

A

- AC PSU, 146, 148
- access to data, controlling, 39, 79
- address, FC-AL, 295
- administrator login, 49
- AL_PA
 - displaying drive's, 65
- AL_PA value, Fibre Channel, 52
- ALARM command, 329
- Alarm Mute push-button, 159, 214, 215
- alarm, enclosure audible, 238, 247
- anonymous user, adding, 101
- anti-static wrist or ankle strap, 168
- anti-tamper lock, 158, 199, 227
- APC UPS trap monitor, 110
- APC_UPS command, 330
- API command, 331
- API connections, 69, 110
- API server port number, changing RM controller, 69
- array
 - configuring, 36
 - controlling access to, 39
- Audible Alarm, 151, 159, 214, 215, 216, 218
- audit, setting user, 100
- AV
 - setting disk, 86
 - setting system, 85
- AV command, 332

B

- Backplane, 146, 168, 169, 220, 227, 228
- bay, 147, 166, 167
- blank I/O module, 146, 158
- blank plates, enclosure, 246

C

- cable
 - connections
 - SF2016, 264, 268, 277, 278
 - SF4016, 259, 263, 273, 275
 - SF6016, 256, 270
 - connections, RM660, 22
 - requirements, 12
- cable, connecting, 22
- cache
 - coherency, 79
 - defaults setting, load, 84
 - disk writeback, setting, 85
 - locking a LUN in, 88
 - memory, definition, 418
 - multiplication factor, setting, 84
 - prefetch, setting, 84
 - maximum, 84
 - segment size
 - definition, 418
 - setting, 83
 - writeback limit, 84
 - writeback, setting, 83, 84

- CACHE command, 337
- channel
 - definition, 24, 25, 418
 - location, physical, 64
- chassis, 146, 169
- check condition log, 108
- COMMENT command, 110, 340
- connections, current user, 100
- CONSOLE command, 341
- cooling
 - enclosure, 246
 - RM610, 10
 - RM660, 6
- create LUNs, 36
- current owner of session, checking, 51

D

- data, controlling access to, 39, 99
- DATE command, 342
- date, setting system, 72
- DEBUG command, 344
- DEFAULTS command, 345
- defaults command, 73
- delete a LUN, 77
- device address, 295
- Disk, 207
- disk
 - AV, setting, 86
 - channel
 - definition, 24, 25, 418
 - scanning, 62
 - status, 55
 - defect list, show, 56
 - error messages, 134
 - fast AV, setting, 86
 - find all, 56

- information, 56
- LILP payload, display, 56
- listing all, 56
- location, physical, 63
- ports, RM610, 11
- rebuild
 - pause, 79
 - starting, 56, 78, 79
 - verify, 78
- RLS information, 57
- scan, 56
- status, 58
- status, show, 56
- timeout, setting, 85
- writeback cache, setting, 85

- DISK command, 346
- disk drive, 147, 167
- Disk Drives LEDs, 207
- DISK_CLI command, 352
- drive
 - failure recovery, 121
- drive carrier, 195
- drive carrier handles, 158
- Drive Carrier Module, 147, 156, 167, 194, 207, 210, 219, 228
- drive enclosure, see enclosure
- drive error messages, 134
- drive failure recovery, 121
- drive module LED
 - SFx016, 315
- Drive spin down, 227
- drive status indicators, 157
- DUAL command, 354
- dual command, 79
- dummy carrier module, 146, 158, 199, 219

E

- emulated devices, 209
- Enclosure, 147, 159, 165, 168, 217, 219
- enclosure
 - address, device, 295
 - anti-tamper lock, 245
 - audible alarm, 238, 247
 - bay numbering, 236
 - blank plates, 246
 - connection, verify SF6016, 64
 - cooling, maintain sufficient, 246
 - device address, 295
 - display failure on, 63, 109
 - drive module, 244
 - specifications, 324
 - dummy drive modules, 246
 - failure, recovery, 125
 - fan, 240
 - ID, 295
 - ID switch, 238
 - installation, 247
 - IO module, 241
 - replacement, 319
 - LED
 - AC input fail, 309
 - cooling, 240, 309
 - DC output fail, 309
 - drive module, 245, 315
 - ESI, 238
 - fan fault, 309
 - Ops panel, 238, 309
 - states, 310
 - power supply, 240, 309
 - system, 316
 - system fault, 238, 309
 - loop terminate, set, 238
 - Ops panel, 237
 - LED, 309
 - replacement, 319
 - switch, configuration, 238
 - power on, 284
 - power supply, 240
 - replacement, 317
 - problem solving, 308
 - spare parts, list of, 325
 - status monitoring rate, 98
 - technical specifications, 321
 - enclosure bay numbering convention, 167
 - Enclosure Services Processor, 209
 - Enclosure System Interface/Operators Panel, 150
 - equipment rail, installing, 18
 - error recovery, 119
 - ESD, 168, 220
 - ESD precautions, 227
 - ESI processor, 215
 - Ethernet, connection, 114
 - event log, 108

F

- failure
 - drive fixing, 121
 - enclosure, fixing, 125
 - fan, fixing, 120
 - list of, display, 109
 - notifications, SES, 63
 - power supply, fixing, 119
 - recovery, 119
- fan
 - enclosure, 240
 - failure recovery, 120
 - RM610, 10
 - RM660, 6
- fan failure, 216
- fast AV, setting disk, 86
- FAULTS command, 356

- faults command, 109
- FC loop status on host ports, 54
- FC_AL signal, 210
- Fibre Channel AL_PA value, 52
- firmware, update, 104
- format LUNs, 36

G

- gateway, 70
- gateway, changing RM controller, 67

H

- hard loop ID, 52
- heading part no., 284
- HELP command, 47, 358
- host
 - adapter, 12
 - connections, RM660, 22
 - delay statistics, 92
 - ID, setting, 52
 - lipinfo, setting, 54
 - port array parity check, 55
 - port settings, 51
 - port speed, 55
 - port status, 51
 - ports, display LILP payload, 54
 - ports, FC loop status, 54
 - ports, RM610, 11
 - timeout, setting, 99
 - utility, 51
 - WWN, setting, 53
- HOST command, 359
- hot spare, definition, 420

I

- ID, enclosure, 295
- IMG
 - pause, 112
- installation
 - enclosure, 248
 - RM controller, 14-44
- Internal Mirrored Groups, see IMG
- I/O connectors
 - RM610, 11
- IP address, changing RM controller, 66

L

- labeling a LUN, 75
- L-bracket, 18
- LED, 150, 151, 159, 210, 211, 215, 216, 218, 219
- LED indicators, 153
- LED, enclosure, see enclosure
- LICENSE command, 362
- LILP payload for host ports, 54
- locking a LUN, 88
- LOG command, 364
- log command, 108
 - check condition log, 108
- logical unit, definition, 420
- login
 - command, 49
 - current user, 51, 100
 - manage remote, 105
- LOGIN command, 367
- LOGOUT command, 368
- logout command, 49
- loop ID, 52
- loop resiliency circuit, 421

LRC, 421

LUN

- controlling access to, 39
- creating, 74
- definition, 420
- deleting, 77
- formatting, 36
 - pause, 75
- labeling, 75
- location, physical, 64
- locking, unlocking, 88
- mapping for users, display, 100
- mapping, defining, 40
- maximum number supported, 30, 420
- moving, 76
- ownership of, 79
- place holder, 39, 99
- read-only, setting, 41, 44, 101, 102
- reconstruction, 422
 - see also disk rebuild
- status, checking, 60

LUN command, 369

M

M_WAIT command, 98

mapping

- defining LUN, 40
- tier, 32, 64
- users' LUN, display, 100

message log, 108

message, echo to screen, 110

messages

- cache error, 128
- disk boot-up error, 130
- disk error, 134
- fault, 127
- host error, 139
- informational, 128

- rebuild error, 141
- SDRAM error, 128
- SES, 142
- timeout, 128
- warning, 127

mirror

- add, 112
- break, 112
- create, 112
- delete, 112
- destroy, 112
- merge, 112
- resume, 112
- stop, 112

MIRROR command, 373

mirror command, 111

moving a LUN, 76

multiplication factor, setting, 84

N

netmask, changing RM controller, 67

NETWORK command, 375

network command, 66, 67

O

operational time, system, 109

Operators Panel, 146, 150

Ops Panel, 150, 159, 166, 169, 203, 211, 214, 215, 224

Ops Panel LEDs, 205, 211

Ops Panel Switch Functions, 152, 192

owner of session, checking, 51

ownership of LUNs and tiers, 79

P

- parity, definition, 421
- PASSWORD command, 379
- password command, 50
- PATA/SATA disk drive, 209
- pause LUN format, 75
- pause rebuild, 79
- performance statistics, 90
- permission denied, 50
- physical location
 - channel, 64
 - drive, 63
 - LUN, 64
 - tier, 64
- ping, 67
- place holder LUN, 39, 99
- plug-in module, 146, 165
- Port Bypass Circuits, enclosure, 241
- Power Down, 207
- Power On, 203
- power supply
 - failure recovery, 119
 - RM610, 10
 - RM660, 5
- power supply cord, 170, 200, 221, 228
- Power Supply/Cooling, 146, 150, 166, 169, 203, 209, 215, 216, 220, 221, 222, 228
- PowerLUN, 29
- prefetch, setting, 84
- PSU LEDs, 204, 211
- PSU voltage operating ranges, 150

R

- rack installation

- RM controller, 18
- Rack mounting, 146, 147
- rack mounting rail kit, 228
- RAID Head, 209
- RAID level, definition, 422
- read-only LUN, setting, 41, 44, 101, 102
- rebuild
 - automatic, 56, 78
 - definition, 422
 - delay, setting, 97
 - extent, setting, 97
 - manual, initiate, 78, 79
 - verify, 78
- recovery
 - drive failure, 121
 - enclosure, failure, 125
 - fan failure, 120
 - power supply, 119
- remote management, 114
- replacing a module, 220
- RESTART command, 380
- restart command, 71
- resume LUN format, 75
- resume rebuild, 79
- RM controller
 - API server port number, changing, 69
 - cache coherency, 79
 - cache error messages, 128
 - configuration, defaults, 73
 - configuring, 28
 - couplet, 419
 - disk boot-up error messages, 130
 - disk error messages, 134
 - failing the other, 79, 81
 - Fault messages, 127
 - firmware update, 104
 - gateway, 70
 - gateway, changing, 67
 - healing the other, 79

- host error messages, 139
- host port array parity check, 55
- host port ID, 52
- host port speed, 55
- host port WWN, 53
- IP address, changing, 66
- label, 81
- management interface, 45
- messages, informational, 128
- name, 81
- netmask, changing, 67
- package, 14
- ping with single packet, 67
- rack mounting, 17
- rebuild error messages, 141
- remote management, 114
- restarting, 71
- restore other, 81
- routing table, 70
- SDRAM error messages, 128
- SES messages, 142
- shutting down, 71
- SNMP
 - information, report, 68
- SNMP functionality, enabling, 68
- SNMP trap IP address, changing, 68
- syslog capability, enabling, 68
- syslog packets IP address, changing, 69
- syslog packets port number, changing, 69
- Telnet capability, enabling, 68
- Telnet port number, changing, 68
- tier
 - mapping, change, 32
- time, changing system, 72
- Timeout messages, 128
- Warning messages, 127
- RM610
 - components, 8
 - cooling modules, 10
 - disk ports, 11
 - fans, 10
 - host ports, 11
 - I/O connectors, 11
 - power supplies, 10
 - technical specifications, 231
- RM660
 - cable connections, 22
 - cooling modules, 6
 - fans, 6
 - power supplies, 5
 - technical specifications, 229
- ROUTE command, 381
- route command, 70
- routing table, 70
- RS-232 terminal set up, 31
- S**
- SAN, definition, 422
- SATA Control Interface Adaptor (SCM I/O) module, 153
- SATA disk drive, 153
- SAVE command, 383
- SCM (I/O) module, 172
- SCM I/O module, 153, 190, 210, 224, 226
- SCM module, 172
- SCSI Enclosure Services, 423
- security, setting data, 39, 99
- segment size
 - definition, 423
 - setting, 83
- SEL_ID base address, 191
- Serial ATA (SATA) Drive Carrier module, 146
- Serial ATA Control (SCM) Input/Output module, 146
- serial ATA protocol, 173
- serial port, set up, 31
- SES, 194, 324

- definition, 423
- device monitoring rate, 98
- display device information, 63
- enclosure, 281
- utility, 63, 98

SES command, 384

SETTINGS command, 390

settings, CLI session, 114

SFP cable

- connecting to SF2016, 264, 268, 277, 278
- connecting to SF4016, 259, 263, 273, 275
- connecting to SF6016, 256, 270

SHOWALL command, 391

SHUTDOWN command, 398

shutdown command, 71

SNMP

- information, report, 68

SNMP functionality, enabling RM controller, 68

SNMP trap IP address, changing RM controller, 68

spare drive, 77

specifications, see technical specifications

speed, host port, 55

starting drives, 206

statistics

- clear, 91
- command delay, 92
- command length, 95
- delayed disk command counts, 94
- dual message, 94
- host command offsets, 95
- host delay, 92
- performance, 90
- repeat, 96
- reset, 91
- telnet, 105
- tier delay, 93

STATS command, 399

stats command, 90

- stop LUN format, 75
- stop rebuild, 79
- streaming data, minimize latency for, 86
- syslog capability, enabling RM controller, 68
- syslog packets IP address, changing RM controller, 69
- syslog packets port number, changing RM controller, 69

SYSTEM LED, 215

system uptime, display, 109

T

technical specifications

- enclosure, 321
- RM610, 231
- RM660, 229

Telnet, 114

- capability, enabling RM controller, 68
- kill login, 107
- manage login, 105
- terminate session, 107

TELNET command, 403

Telnet port number, changing RM controller, 68

TFTP command, 405

throughput, system, 2

tier

- adding, 62
- AL_PA mapping chart, enclosure, 295
- configuration, checking, 34, 57
- definition, 24, 25
- delay statistics, 93
- display configuration, 58
- location, physical, 64
- mapping mode, selecting, 32, 64
- ownership of, 79
- removing, 62
- status, checking, 34, 57

TIER command, 406
TIME command, 409
time, system, 72
timeout, setting disk, 85
troubleshooting, enclosure, 308

U

unlocking a LUN, 88
UPTIME command, 410
uptime command, 109
user
 adding new, 102
 command, 102
 connected, 100
 current login, 51
 login, 49
 LUN mapping information, 100
USER command, 411

V

VERSION command, 413
version information, 103
Visible and Audible Alarms, 159
visual indication
 channel, 64
 drive, 63
 LUN, 64
 tier, 64

W

WHOAMI command, 51, 414
write cache mirroring, definition, 418
write caching, definition, 418

writeback cache, setting, 83, 84

Z

ZONING command, 415
zoning, configuring, 101

