

Quantum[®]

StorNext 3.5.3 User's Guide

Quantum StorNext



Document Title, 6-67339-01 Rev A, June 2011, Product of USA.

Quantum Corporation provides this publication “as is” without warranty of any kind, either express or implied, including but not limited to the implied warranties of merchantability or fitness for a particular purpose. Quantum Corporation may revise this publication from time to time without notice.

COPYRIGHT STATEMENT

© 2011 Quantum Corporation. All rights reserved.

Your right to copy this manual is limited by copyright law. Making copies or adaptations without prior written authorization of Quantum Corporation is prohibited by law and constitutes a punishable violation of the law.


TRADEMARK STATEMENT

Quantum, the Quantum logo, DLT, DLTtape, the DLTtape logo, Scalar, and StorNext are registered trademarks of Quantum Corporation, registered in the U.S. and other countries.

Backup. Recovery. Archive. It’s What We Do., the DLT logo, DLTSage, DXi, DXi-Series, Dynamic Powerdown, FastSense, FlexLink, GoVault, MediaShield, Optyon, Pocket-sized. Well-armored, SDLT, SiteCare, SmartVerify, StorageCare, Super DLTtape, SuperLoader, and Vision are trademarks of Quantum.

LTO and Ultrium are trademarks of HP, IBM, and Quantum in the U.S. and other countries. All other trademarks are the property of their respective companies.

Specifications are subject to change without notice.



Contents

Chapter 1	Introduction	1
	About StorNext File System	1
	About StorNext Storage Manager	1
	About Distributed LAN Clients	2
	Purpose of This Guide	3
	How This Guide is Organized	3
	Notes, Cautions, and Warnings	4
	Document Conventions	4
Chapter 2	StorNext GUI Overview	7
	Accessing the StorNext GUI	7
	The StorNext Home Page	10
	System Monitors	10
	StorNext Home Page Drop-Down Menus	14
	StorNext System Status	17
	StorNext Admin Alerts	18
	StorNext Server Status	20
	Home and Help Links	20
	Application Link	20
	The SNFS and SNSM Home Pages	21

The SNFS Home Page	21
SNFS Home Page Drop-Down Menus	22
The SNSM Home Page	24
SNSM Home Page Drop-Down Menus	24
The Configuration Wizard	27

Chapter 3	SNFS-Only GUI Overview	31
	Accessing the SNFS GUI	31
	The SNFS Home Page	33
	The File System Monitor	34
	Drop-down Menus and Options	36
	Shortcut Menu Options	37
	Home and Help Links	38
	StorNext Server Status	38
	The SNFS Configuration Wizard	40

Chapter 4	Common StorNext Tasks	43
	Entering the StorNext License	43
	Entering a License String in the .dat File	50
	Controlling User Access	51
	Changing the Admin Password	52
	Adding a New User	53
	Modifying an Existing User	55
	Deleting an Existing User	57
	Starting and Stopping StorNext Components	57
	Accessing StorNext Logs	59
	Scheduling StorNext Events	62
	Viewing a Schedule	63
	Adding a New Schedule	65
	Modifying an Existing Schedule	67
	Deleting an Existing Schedule	68
	Resetting a Schedule	69
	Setting Up E-mail Notification	70
	Cancelling SNSM Requests	75

Setting up Alternate Retrieval Locations	76
Modifying or Disabling This Feature	78

Chapter 5	Backing Up StorNext	79
	Types of StorNext Software Backups	79
	Setting up Email Notification For Backup	80
	Performing a StorNext Software Backup	80
	Managing the Backup Policy	82

Chapter 6	Managing the File System	85
	Working With File Systems	85
	Adding a File System.	86
	Creating a File System From SNFS.	95
	Modifying a File System	97
	Deleting a File System.	99
	Managing File System Operations	100
	Disk Device Labeling	100
	Making Global Changes	104
	Working with the fsnameservers File	108
	Making or Unmaking a File System.	109
	Starting and Stopping the File System	111
	Mounting or Unmounting a File System	112
	Working With Disks.	113
	Adding Disks.	113
	Deleting Disks	115
	Defragmenting a Disk	116
	Working With Stripe Groups	116
	Adding a Stripe Group	117
	Modifying a Stripe Group	121
	Deleting a Stripe Group	124
	Working With Affinities	125
	Adding an Affinity	125
	Adding an Affinity Through SNFS	131
	Setting the Affinity in a Directory	133
	Modifying an Affinity	135
	Deleting an Affinity.	136

File System Configuration Restrictions	137
Performing a Metadata Dump	137
Using the SNSM File System Functions	138
Storing Files	139
Changing a File Version	141
Recovering a File	142
Recovering a Directory	143
Retrieving a File	145
Retrieving a Directory	146
Freeing Disk Blocks	146
Moving Files to New Media	148
Modifying a File's Attributes	149
Understanding Resource Allocation	150
About File System Expansion	150
About Stripe Group Movement	151
Expansion and Movement Steps	152
Checking the File System	152
Performing File System Expansion	154
Performing Stripe Group Movement	161
Launching the Movement Wizard	162
Completing Metadata Stripe Group Movement	170
Reusing LUNs in a Data Stripe Group After a Move	171
Reusing LUNs in a Metadata Stripe Group After a Move	172

Chapter 7	Managing Libraries	173
	Adding a Library	174
	Starting the Add Library Wizard	174
	Adding a SCSI Library	176
	Adding an ACSLS Network Library	180
	Adding a DAS Network Library	182
	Adding a Vault Library	186
	Modifying a Library	187
	Deleting a Library	188
	Rescanning a Library	189
	Auditing a Library	190
	Changing the Library State	191

Chapter 8	Managing Drives and Disks	193
	Working with Tape Drives	193
	Adding a Tape Drive	194
	Modifying a Tape Drive	199
	Deleting a Tape Drive	201
	Changing a Drive State	201
	Cleaning a Tape Drive	202
	Working with Drive Pools	203
	Adding a Drive Pool	204
	Modifying a Drive Pool	205
	Deleting a Drive Pool	207
	Managing Disk Space	208
	Changing Watermark Parameters	209
Chapter 9	Managing Media	211
	Adding Media to a Configured Library	211
	Adding a Vault	216
	Removing and Moving Media	220
	Removing Media	220
	Moving Media	225
	Using the SNSM Media Functions	228
	Moving Media Manually	228
	Mounting and Dismounting Media	231
	Removing Media From the Storage Manager	232
	Assigning Blank Media to a Policy Class	233
	Transcribing Media	234
	Changing Media Attributes	236
	Reclassifying a Media Class Grouping	237
	Cleaning Media	238
	Cancelling the Eject Media Process	240
Chapter 10	Managing Storage Disks	241
	Storage Disk Deduplication	242
	Adding a Storage Disk	243

Modifying a Storage Disk	247
Deleting a Storage Disk	248
Changing a Storage Disk State	249
Cleaning a Storage Disk	251

Chapter 11	Data Migration Management	253
	Policy Classes and Relationships	254
	Stub Files	254
	Disk-to-Disk Relocation	255
	Adding a Storage Policy	255
	Adding a Policy Class Through SNSM Without a Relation Point	262
	Adding a Relation Point to a Policy Class	267
	Modifying a Policy Class	269
	Deleting a Policy Class	270
	Applying a Policy Class	271

Chapter 12	StorNext Reports	273
	SNFS Reports	274
	The Backup Information Report	274
	The Drive State Information Report	275
	The File Information Report	278
	The Library Information Report	281
	The Library Space Used Report	283
	The Media Information Report	284
	The Policy Class Information Report	289
	The Relation Information Report	292
	The Request Information Report	294
	The Scheduler Information Report	295
	The Storage Disk Information Report	297

The Directory Affinity Report 300

The File System Statistics Report 303

The Stripe Group Statistics Report 305

The File System Client Report 307

The File System LAN Client Report 310

Chapter 13 Service Management 313

Using Health Check 313

 Running a Health Check 314

 Viewing the Health Check History 315

 Viewing the Health Check Results 316

Using State Capture 317

 Capturing the Current System State 317

 Downloading a Previous System Capture 318

 Deleting a Previous System Capture 319

Using the System Status Tool 320

Chapter 14 Customer Assistance 325

Quantum Technical Assistance Center 325

Appendix A Operating Guidelines 327

The Reserved Space Parameter 327

Distributed LAN Server/Client Network and Memory Tuning 328

 Distributed LAN Server and Client Network Tuning 329

 Distributed LAN Server Memory Tuning 330

Configuring LDAP 332

 Using LDAP 332

 UNIX File and Directory Modes 332

 LDAP Refresh Timeout 333

 User ID Mapping Precedence 333

Setting Up Restrictive ACLs 333

Default Single-Path I/O Retry Behavior 334

Event Handles for fsm.exe on a Windows Metadata Server	334
FSBlockSize, Metadata Disk Size, and JournalSize Settings	335
Allowance for Special Characters.	337
Disk Naming Requirements	338
General Operating Guidelines and Limitations	338


Appendix B	HA Failover	351
-------------------	--------------------	------------

Appendix C	Using The Command Line Interface	353
	Labeling Disk Devices.	354
	Modifying Global Settings.	355
	Making a File System.	357
	Starting and Stopping SNFS	358
	Unmounting or Mounting a File System	359
	Creating a File System Server.	359
	Adding a File System Client	362
	Configuring a Stripe Group	362
	Adding an Affinity.	363
	Creating a Disk-to-Disk Policy Class	366
	Modifying a Disk-to-Disk Policy Class	366
	Manual Disk-to-Disk Relocation	367
	Enabling Stub File Support	367
	Managing Storage Disks with Deduplication Enabled	368
	Adding a Dedup Sdisk.	368
	Modifying a Dedup Sdisk	369
	Deleting a Dedup Sdisk.	369
	Obtaining Dedup Sdisk Information	369
	Obtaining Distributed LAN Client Information.	370
	The proxy Command.	370
	The proxy long Command.	371
	The proxy who Command.	371

Using the Resource Allocation Feature	372
Checking the File System	372
Adding a Stripe Group Without Moving	372
Adding and Moving a Data Stripe Group	373
Moving a Metadata/Journal Stripe Group	374
Specifying an Alternate Retrieval Location.	376

Appendix D

RAS Messages	377
Media and Drive RAS Messages	377
SNFS RAS Messages.	393
Other RAS Messages	405



Figures

Figure 1	StorNext Login Window	8
Figure 2	StorNext Home Page.	9
Figure 3	File System Monitor	11
Figure 4	Library Monitor.	12
Figure 5	Storage Disk Monitor	13
Figure 6	Service - System Status screen	18
Figure 7	Service - Admin Alerts Screen	19
Figure 8	Application Link	21
Figure 9	SNFS Home Page	22
Figure 10	SNSM Menu	25
Figure 11	StorNext Configuration Wizard.	28
Figure 12	Config Menu Options	29
Figure 13	StorNext File System Login Window	33
Figure 14	SNFS Home Page	34
Figure 15	File System Monitor	35
Figure 16	File System Menus	36
Figure 17	Start/Stop SNFS Screen	39
Figure 18	Complete Start/Stop SNFS Screen	40

Figure 19	SNFS Configuration Wizard Screen	41
Figure 20	Enter License Introduction Screen.	45
Figure 21	Quantum License Agreement	46
Figure 22	Enter License Strings Screen	47
Figure 23	Complete Enter License Screen (Temporary)	48
Figure 24	License Reminder	48
Figure 25	Complete Enter License Screen (Permanent)	50
Figure 26	License String Example	51
Figure 27	User Access Control Screen.	52
Figure 28	Modify User Screen.	53
Figure 29	Add New User Screen	54
Figure 30	Modify User Screen.	56
Figure 31	Start/Stop StorNext Screen	58
Figure 32	Complete Start/Stop StorNext Task Screen	59
Figure 33	Select Log Screen	60
Figure 34	Select File Screen	61
Figure 35	Log File Example.	62
Figure 36	Schedule Events Screen 1	64
Figure 37	Schedule Events Screen 2	65
Figure 38	Schedule Events Screen 3	66
Figure 39	Configure E-mail Notification Screen	71
Figure 40	Configure SMTP E-mail Screen	72
Figure 41	Configure E-mail Addresses Screen.	73
Figure 42	Complete E-mail Configuration Screen	74
Figure 43	Cancel Request Screen	75
Figure 44	Alternate Retrieval Location Screen	77
Figure 45	Backup StorNext Screen	81
Figure 46	Complete Backup Task Screen	82
Figure 47	Backup Policy Screen	83

Figure 48 File System - Introduction Screen 87

Figure 49 Add New File System Screen 88

Figure 50 Directory Browser Window 89

Figure 51 Select Directory Window 89

Figure 52 Disk Settings Screen 90

Figure 53 Customize Stripe Group Screen 92

Figure 54 Label Help Screen 93

Figure 55 Complete File System Task Screen 95

Figure 56 Configure File System Screen 96

Figure 57 Modify File System Window 98

Figure 58 Label Disk Device Screen 101

Figure 59 Labeling Warning 102

Figure 60 Label Name Window 102

Figure 61 Data Loss Warning 103

Figure 62 Unlabel Warning 103

Figure 63 Modify Global Settings Screen 105

Figure 64 Make File System Screen 110

Figure 65 Start or Stop File System Screen 111

Figure 66 Mount or Unmount File System Screen 112

Figure 67 Manage Disks Screen 114

Figure 68 Add Disk Screen 114

Figure 69 Configure Stripe Group Screen 118

Figure 70 Add Stripe Group Screen 119

Figure 71 Modify Stripe Group Screen 122

Figure 72 Add Affinity Introduction Screen 126

Figure 73 Add Affinity Screen 127

Figure 74 Select Directory Screen 128

Figure 75 Assign Affinity Screen 129

Figure 76 Exclusive Stripe Group Screen 130

Figure 77	Complete Add Affinity Task Screen	131
Figure 78	Add, Modify, or Delete Affinities Screen	132
Figure 79	Add Affinity Screen	132
Figure 80	Set Affinity Screen	134
Figure 81	Set Affinity Screen	134
Figure 82	Modify Affinity Screen	135
Figure 83	Metadata Dump of File System Screen	138
Figure 84	Store Files Screen	139
Figure 85	Optional Store Parameters Screen	140
Figure 86	Change File Version Screen	141
Figure 87	Recover Files Screen	142
Figure 88	StorNext Recoverable Files Screen	143
Figure 89	Recover Directory Screen	144
Figure 90	Retrieve Files Screen	145
Figure 91	Retrieve Directory Screen	146
Figure 92	Free Disk Blocks Screen	147
Figure 93	Move Files to New Media Screen	148
Figure 94	Modify File Attributes Screen	149
Figure 95	Check File System Screen	153
Figure 96	Check File System Status Screen	154
Figure 97	File System Expansion Introduction Screen	156
Figure 98	Check File System Reminder	156
Figure 99	Data Stripe Group Selection Screen	157
Figure 100	Metadump Warning	158
Figure 101	New Data Stripe Group Screen	159
Figure 102	Complete File System Task Screen	161
Figure 103	Move Stripe Group Screen	163
Figure 104	Move Stripe Group Introduction Screen	164
Figure 105	Move Options Screen	165

Figure 106 Data Stripe Group Move Screen 166

Figure 107 No Writable Stripe Groups Warning 166

Figure 108 Destination Stripe Group Too Small Warning 167

Figure 109 Metadata Stripe Group Message 167

Figure 110 Metadata/Journal Move Screen 168

Figure 111 Complete Stripe Group Move Screen 169

Figure 112 Process Initiated Status Screen 170

Figure 113 Stripe Group Move Status Screen 171

Figure 114 Library Introduction Screen 175

Figure 115 Library Type Screen. 176

Figure 116 Library Name Screen. 177

Figure 117 Media Types Screen 178

Figure 118 SCSI Device Screen 179

Figure 119 Complete Add Library Task Screen 180

Figure 120 ACSLS Library Name Screen 181

Figure 121 DAS Configuration Screen 182

Figure 122 DAS Library Name Screen 183

Figure 123 DAS Media Types Screen. 184

Figure 124 DAS 2 Library Name Screen 185

Figure 125 Vault Library Name Screen 186

Figure 126 Configure Library Screen. 187

Figure 127 Modify SCSI Library Screen 188

Figure 128 Audit Library Screen 190

Figure 129 Change Library State Screen. 191

Figure 130 Tape Drive Introduction Screen. 194

Figure 131 Associated Library Screen 195

Figure 132 Hardware Devices Screen 196

Figure 133 Complete Add Drive Task Screen 197

Figure 134 Match Devices with Slots Screen. 198

Figure 135	Tape Drive Mapping Help Screen	199
Figure 136	Configure Drives Screen	200
Figure 137	Modify Drive Screen	200
Figure 138	Delete Warning Window	201
Figure 139	Change Drive State Screen	202
Figure 140	Clean Drive Screen	203
Figure 141	Configure Drive Pools Screen	204
Figure 142	Add New Drive Pool Screen	205
Figure 143	Warning Message Window.	205
Figure 144	Modify Drive Pool Screen	206
Figure 145	Restart Message Window	207
Figure 146	Delete Warning Message	207
Figure 147	Manage Disk Space Screen	208
Figure 148	Change Watermark Parameters Screen.	210
Figure 149	Add Media - Introduction Screen	212
Figure 150	Associated Library Screen	213
Figure 151	Associated Library Screen 2	214
Figure 152	Select Mailbox Screen.	215
Figure 153	Complete Add Media Task Screen.	216
Figure 154	Select Media Type Screen	217
Figure 155	Add Media IDs Screen.	218
Figure 156	Create New Media ID Screen	219
Figure 157	Remove or Move Media Screen	220
Figure 158	Select Media Screen	221
Figure 159	StorNext Media Browser Screen	222
Figure 160	Complete Remove/Move Media Task Screen.	223
Figure 161	Library Operator Interface Screen	224
Figure 162	LOI Eject Screen	224
Figure 163	Select Destination Library Screen	226

Figure 164 Library Operator Interface Screen 227

Figure 165 LOI Enter Screen 227

Figure 166 Manually Move Media Screen 229

Figure 167 Media Browser Screen 230

Figure 168 Mount Media Screen 231

Figure 169 Dismount Media Screen 232

Figure 170 Remove Media From SNSM Screen 233

Figure 171 Assign Policy Screen 234

Figure 172 Transcribe Media Screen 235

Figure 173 Change Media Attributes Screen 236

Figure 174 Reclassify Media Screen 237

Figure 175 Media Class Browser Window 238

Figure 176 Clean Media Screen 239

Figure 177 Cancel Eject Process Screen 240

Figure 178 Configure Storage Disk Screen 243

Figure 179 Add Storage Disk - Introduction Screen 244

Figure 180 Add Storage Disk Screen 245

Figure 181 Complete Storage Disk Screen 246

Figure 182 Modify Storage Disk Screen 247

Figure 183 Delete Warning Message 249

Figure 184 Change Storage Disk State Screen 250

Figure 185 Clean Storage Disk Screen 251

Figure 186 Storage Policy Introduction Screen 256

Figure 187 Policy Class and Directory Screen 257

Figure 188 Directory Browser Window 258

Figure 189 Relocation Policy Selection Screen 259

Figure 190 Store, Truncate, and Relocate Times Screen 260

Figure 191 Number of File Copies and Media Type Screen 261

Figure 192 Complete Storage Policy Task Screen 262

Figure 193	Manage Policy Classes Screen	263
Figure 194	Add Policy Class Screen	263
Figure 195	Modify Parameters Screen	264
Figure 196	Manage Policy Class Relationships Screen	268
Figure 197	Add Relationship Screen	268
Figure 198	Directory Browser Window	269
Figure 199	Remove Relationships Screen	271
Figure 200	Backup Information Report	275
Figure 201	Drive States Report Screen	277
Figure 202	Drive State Information Report	278
Figure 203	Files Report Screen	279
Figure 204	StorNext File Browser Screen	280
Figure 205	File Information Report	281
Figure 206	Libraries Report Screen	282
Figure 207	Library Information Report	283
Figure 208	Library Space Used Report	284
Figure 209	Media Report Screen	286
Figure 210	Media Browser screen	287
Figure 211	Media Information Report	288
Figure 212	Detailed Media Information Report	288
Figure 213	Policy Classes Report Screen	291
Figure 214	Policy Class Information Report	292
Figure 215	Directory/Policy Class Relationships Report Screen	293
Figure 216	Relation Information Report	294
Figure 217	Request Report Screen	295
Figure 218	Scheduler Report Screen	296
Figure 219	Scheduler Information Report	297
Figure 220	Storage Disk Report Screen	299
Figure 221	Storage Disk Information Report	300

Figure 222 Affinities Report Screen 301

Figure 223 Directory Browser Screen 302

Figure 224 Directory Affinity Report. 302

Figure 225 File System Report Screen. 304

Figure 226 File System Statistics Report 304

Figure 227 Stripe Groups Report Screen. 306

Figure 228 Stripe Group Statistics Report. 307

Figure 229 File System Client Report Screen. 309

Figure 230 File System Client Report 310

Figure 231 File System Distributed LAN Client Statistics Report
Screen. 311

Figure 232 File System LAN Client Report. 312

Figure 233 Health Check Tests Screen 314

Figure 234 Health Check History Screen 315

Figure 235 Health Check Results Screen 316

Figure 236 Capture System State Screen 318

Figure 237 Download Capture File Screen 319

Figure 238 Service - System Status Screen 320

Figure 239 RAS Ticket Details Screen 322

Figure 240 Recommended Actions Screen 324

Figure 241 No Media Found RAS 378

Figure 242 Possible Drive/Media Mount Discrepancy RAS 379

Figure 243 Tape Drive Alerts RAS part 1 380

Figure 244 Tape Drive Alerts RAS part 2 381

Figure 245 Tape Drive Alerts RAS part 3 382

Figure 246 Drive Reported Drive Error RAS. 383

Figure 247 Cleaning of Drive Failed RAS. 384

Figure 248 Wrong Firmware Level/Invalid Drive Type RAS 385

Figure 249 Drive Removed RAS 385

Figure 250	Tape Drive -Configuration Failed RAS	386
Figure 251	Tape Drive - Reported Media Error RAS	387
Figure 252	Cleaning Media Expired RAS.	388
Figure 253	No Cleaning Media Available RAS.	388
Figure 254	Media Suspect Threshold Count Exceeded RAS	389
Figure 255	Media Format Failure RAS.	390
Figure 256	Invalid Media Label Detected RAS.	391
Figure 257	Media Not Found RAS	392
Figure 258	Duplicate Physical Media Found RAS	392
Figure 259	Storage Disk Taken Offline RAS	393
Figure 260	Configuration Not Supported RAS	394
Figure 261	Label Validation Failure RAS	394
Figure 262	Connection Rejected RAS	395
Figure 263	File System Failover RAS	395
Figure 264	I/O Error RAS.	396
Figure 265	Journaling Error Detected RAS	396
Figure 266	SNFS License Required RAS.	397
Figure 267	SNFS License Failure RAS	397
Figure 268	LUN Mapping Changed RAS.	398
Figure 269	Communication Failure RAS	398
Figure 270	Metadata Inconsistency Detected RAS	399
Figure 271	Bad File System Metadata Dump RAS.	399
Figure 272	Metadata Dump Failure RAS.	400
Figure 273	File System or Metadata Capacity Warning RAS	400
Figure 274	File Processing Failure RAS	401
Figure 275	Missing LUNs RAS.	401
Figure 276	Disk Space Allocation Failure RAS.	402
Figure 277	System Resource Failure RAS	402
Figure 278	Affinity Configuration Violations RAS.	403

Figure 279 Quota Limit or Fragmentation Warnings RAS 404

Figure 280 Shutdown Error RAS 404

Figure 281 Initialization Failure RAS 405

Figure 282 Checksum Error RAS 406

Figure 283 Troubleshooting the StorNext Software RAS 407

Figure 284 Software Resource Violations RAS 408

Figure 285 Closing Service Tickets RAS 409

Figure 286 Analyzing Service Tickets RAS 410

Figure 287 Viewing Service Tickets RAS 411

Figure 288 Vault Failure RAS 411

Figure 289 Robotics - Not Ready RAS 412

Figure 290 Robotics - Move Failure RAS 413

Figure 291 Robotics - Wrong Firmware Level/Invalid Library Type
RAS 414

Figure 292 Backup Failed RAS 415

Figure 293 Backup Errors RAS 416

Figure 294 Configuration Violations RAS 417

Figure 295 Invalid Configuration RAS part 1 418

Figure 296 Invalid Configuration RAS part 2 419

Figure 297 Downloading a System State Capture RAS 419

Figure 298 Capturing a System State RAS 420



Chapter 1 Introduction

StorNext is data management software that enables customers to complete projects faster and confidently store more data at a lower cost. Used in the world's most demanding environments, StorNext is the standard for high performance shared workflow operations and multitier archives. StorNext consists of two components: **StorNext File System (SNFS)**, a high performance data sharing software, and **StorNext Storage Manager (SNSM)**, the intelligent, policy-based data mover.

About StorNext File System

StorNext File System streamlines processes and facilitates faster job completion by enabling multiple business applications to work from a single, consolidated data set. Using SNFS, applications running on different operating systems (Windows, Linux, UNIX, HP-UX, AIX, and Mac OS X) can simultaneously access and modify files on a common, high-speed SAN storage pool. This centralized storage solution eliminates slow LAN-based file transfers between workstations and dramatically reduces delays caused by single-client failures. With SNFS, any server can access files and pick up processing requirements of a failed system to continue operations.

About StorNext Storage Manager

StorNext Storage Manager enhances the StorNext solution by reducing the cost of long term data retention, without sacrificing accessibility. SNSM sits on top of SNFS and utilizes intelligent data

movers to transparently locate data on multiple tiers of storage. This enables customers to store more files at a lower cost, without having to reconfigure applications to retrieve data from disparate locations. Instead, applications continue to access files normally and SNSM automatically handles data access – regardless of where the file resides. As data movement occurs, SNSM also performs a variety of data protection services to guarantee that data is safeguarded both on site and off site.

About Distributed LAN Clients

StorNext supports *distributed LAN clients*. Unlike a traditional StorNext SAN client, a distributed LAN client does not connect directly to StorNext via fibre channel or iSCSI, but rather across a LAN through a gateway system called a *distributed LAN server*. The distributed LAN server is itself a directly connected StorNext client, but it processes requests from distributed LAN clients in addition to running applications.

Any number of distributed LAN clients can connect to multiple distributed LAN servers. StorNext File System supports Distributed LAN client environments in excess of 1000 clients, and should support deployments as large as 5000 clients. File system aggregate throughput is not adversely impacted.

Besides the obvious cost-savings benefit of using distributed LAN clients, there will be performance improvements as well.

Distributed LAN clients must be licensed in the same way as StorNext SAN clients. When you request your permanent StorNext license, you will need to specify the number of distributed LAN clients you plan to use. Naturally, you can always purchase additional distributed LAN client licenses as your needs expand. For more information about StorNext licensing, see [Entering the StorNext License](#) on page 43.

StorNext provides distributed LAN client information via the status monitors on the StorNext home page, SNSM home page, and SNFS home page. More detailed information is available through the Clients Report and the Distributed LAN Client Performance Report. For more information about StorNext reports, see [StorNext Reports](#) on page 273.

Before you can fully use distributed LAN clients, you must first configure a distributed LAN server and distributed LAN clients as described in the *StorNext Installation Guide*.

Purpose of This Guide

This guide is intended to assist StorNext users perform day-to-day tasks with the software. This guide also describes how to generate reports. Quantum recommends using the graphical user interface to accomplish tasks, but an appendix provides alternative procedures for users who wish to perform those tasks via the command line interface.

How This Guide is Organized

This guide contains the following chapters:

- [Chapter 1, Introduction](#)
- [Chapter 2, StorNext GUI Overview](#)
- [Chapter 3, SNFS-Only GUI Overview](#)
- [Chapter 4, Common StorNext Tasks](#)
- [Chapter 5, Backing Up StorNext](#)
- [Chapter 6, Managing the File System](#)
- [Chapter 7, Managing Libraries](#)
- [Chapter 8, Managing Drives and Disks](#)
- [Chapter 9, Managing Media](#)
- [Chapter 10, Managing Storage Disks](#)
- [Chapter 11, Data Migration Management](#)
- [Chapter 12, StorNext Reports](#)
- [Chapter 13, Service Management](#)
- [Chapter 14, Customer Assistance](#)
- [Appendix A, Operating Guidelines](#)
- [Appendix B, HA Failover](#)
- [Appendix C, Using The Command Line Interface](#)

- [Appendix D, RAS Messages](#)

Notes, Cautions, and Warnings

The following table describes important information about Notes, Cautions, and Warnings used throughout this guide.


Description	Definition	Consequences
Note:	Indicates important information that helps you make better use of the software.	No hazardous or damaging consequences.
Caution:	Advises you to take or avoid a specified action.	Failure to take or avoid this action could result in loss of data.
Warning:	Advises you to take or avoid a specified action.	Failure to take or avoid this action could result in physical harm to the user or hardware.

Document Conventions

This guide uses the following document conventions to help you recognize different types of information.

Conventions	Examples
For all UNIX-based commands, the # prompt is implied, although it is not shown.	TSM_control stop is the same as # TSM_control stop

Conventions	Examples
For all UNIX-based commands, words in <i>italic</i> are variables and should be replaced with user-defined values.	cvaffinity <filename> where <filename> is a variable and should be replaced with a user-defined value.



Chapter 2

StorNext GUI Overview

This section describes how to access and navigate through the StorNext GUI, which includes both SNFS and SNSM. If you purchased SNFS only, refer to [SNFS-Only GUI Overview](#).

This chapter includes the following topics:

- [Accessing the StorNext GUI](#)
- [The StorNext Home Page](#)
- [The SNFS and SNSM Home Pages](#)
- [The Configuration Wizard](#)

Note: StorNext supports internationalization for the name space of the file system. This support is fully UTF-8 compliant. It is up to the individual client to set the proper UTF-8 locale.

Accessing the StorNext GUI

The StorNext GUI is browser-based and can be remotely accessed from any machine with access to the StorNext server.

Use this procedure to access the StorNext GUI.

1 Open a Web browser.

Note: StorNext-supported browsers are:

- Internet Explorer 5.5, 6 and 7
- Netscape 7.x
- Mozilla 1.0 and later
- FireFox 1.5 and later or 2.0 and later


To ensure proper browser operation, all monitors must be set to display at a minimum resolution of 1073 x 768. If you use a popup blocker, be sure to enable pop-ups in order to ensure that the StorNext Configuration Wizard functions properly.

2 In the browser's **Address** field, type the full address of the machine and its port number, and then press **Enter**. For example: `http://<machine name>:<port number>`. Use the name of the machine and port number you copied when you installed the StorNext software.

Note: Typically, the port number is 81. If port 81 is in use, use the next unused port number. (I.e., 82, 83, etc.)

After you enter the machine name and port number, the following window appears:

Figure 1 StorNext Login Window



3 In the **User ID** field, type admin and press **TAB**.


4 In the **Password** field, type admin and click **OK**.

The initial StorNext GUI appears. A message asks you to start non-running components.

5 Click **OK** to start the StorNext components. The StorNext GUI launches.

The following illustration shows the main components of the StorNext Home Page.

Figure 2 StorNext Home Page



The StorNext Home Page

On the home page you will find the following:

- **Link to the Quantum Home Page** - Click the link for information about Quantum
- [System Monitors](#)
- [StorNext Home Page Drop-Down Menus](#)
- [StorNext System Status](#)
- [StorNext Admin Alerts](#)
- [StorNext Server Status](#)
- [Home and Help Links](#)
- [Application Link](#)

System Monitors

The StorNext Home Page displays three system monitors that are used to monitor the state of the StorNext system:

- [The File System Monitor](#)
- [The Library Monitor](#)
- [The Storage Disk Monitor](#)

Use these monitors to view current statistics of managed or unmanaged file systems and configured libraries and/or drives, including file system, library, and drive information.

The File System Monitor and Library Monitor have a **Refresh** button that allows you to manually update (refresh) the information shown in the monitor. You can also select a rate from the **Refresh Rate** list to automatically refresh at the selected interval:

- No Refresh
- 30 seconds
- 1 minute
- 2 minutes
- 5 minutes

- 10 minutes

The File System Monitor

The File System Monitor enables you to view statistics on each configured file system. When you open a browser to access StorNext, the File System Monitor appears at the top of the browser window.

Figure 3 File System Monitor

State	File System	Total Space (GB)	Used Space (GB)	Free Space (GB)	Used Inodes	# Store Candidates	# Trunc Candidates	# SAN Clients	# LAN Clients	Status
✓	snfs1	189.66	4.31	185.35	35	N/A	N/A	1	0	
✗	snfs2	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	

The File System Monitor provides the following information:

- **Refresh:** Click this button to manually refresh the File System Monitor.
- **Refresh Rate:** Set the File System Monitor to automatically refresh with this drop down menu. Options range from **No Refresh** to every **10 Minutes**.
- **State:** A green checkbox indicates the file system is mounted and active, a red "x" indicates the file system is not active or not mounted
- **File System:** The name of the file system. (You might see more than one file system being monitored, depending upon how your StorNext system is configured). This pane is scrollable to accommodate numerous file systems.
- **Total Space (GB):** Total space (in GB) for the file system
- **Used Space (GB):** Currently used space (in GB) for the file system
- **Free Space (GB):** Amount of free space (in GB) for the file system
- **Used Inodes:** The number of inodes currently in use on the file system
- **#Store Candidates:** Number of files selected for storage to secondary media
- **#Trunc Candidates:** Number of files that have been stored and meet the criteria to become a truncation candidate

Note: The #Store Candidates and # Trunc Candidates fields show “N/A” for a non-managed file system. For a managed file system, if these fields show “refresh,” click the Refresh button to retrieve current information for these fields.

- **# SAN Clients:** The number of StorNext SAN clients (connected via fibre channel or iSCSI) for which you are licensed
- **# LAN Clients:** The number of StorNext distributed LAN clients for which you are licensed. For more information about distributed LAN clients, see [About Distributed LAN Clients](#) on page 2.
- **Status:** The status shows the system status (usage) in percent in addition to low and high watermark settings. The low watermark specifies the level of used disk space that is acceptable to end overflow processing. The high watermark specifies the level of used disk space that initiates overflow processing. In the above example, the system usage is less than 10% with 75% set as the low watermark and 85% set as the high watermark. (High and low watermarks do not apply to non-managed file systems.)

Note: Overflow processing occurs when the system processes beyond the set watermark limitations.

The Library Monitor

The Library Monitor enables you to view library and drive information on each library. When you open a browser to access StorNext, the Library Monitor appears below the File System Monitor.

Figure 4 Library Monitor

State	Library Name	Library Type	Number of Drives	Capacity	Fill Level
✓	scsi_archive1	SCSI	2	60	0

State	Drive Name	Status	Mounted Media	Compression
✓	scsi_archive1_dr1	Free		ON
✓	scsi_archive1_dr2	Free		ON

The Library Monitor provides the following information:

- **Refresh:** Click this button to manually refresh the Library Monitor

- **Refresh Rate:** Set the Library Monitor refresh rate with this drop down menu. Options range from **No Refresh** to every **10 Minutes**.
- **State:** A green checkbox indicates the library is online, a red "x" indicates the library is offline
- **Library Name:** Displays the name of the library
- **Library Type:** Indicates the library type: SCSI, Vault, or Network
- **Number of Drives:** Specifies the number of drives that are configured for the library
- **Capacity:** Total number of media slots in the library
- **Fill Level:** Number of media in the archive
- **Drive Name:** Name (or label) of configured drive (or drives) for the library
- **Status:** Indicates if the drive is free, in use, failed, delayed, or cleaning.
- **Mounted Media:** Media ID of the media currently mounted in the drive
- **Compression:** Indicates whether compression is enabled

The Storage Disk Monitor

The Storage Disk monitor allows you to view configured storage disk information. When you open a browser to access StorNext, the Storage Disk Monitor appears below the Library Monitor.

Figure 5 Storage Disk Monitor

Storage Disk Monitor							
State	Storage Disk Name	Deduplication	Space (GB)	Used Space (GB)	Copy #	# Files	% Savings
<input checked="" type="checkbox"/>	Storage_Disk_1	Disabled	39.1	16.23	1	170	
<input checked="" type="checkbox"/>	Storage_Disk_2	Enabled	184.28	5.38	2	170	0.18

The Storage Disk Monitor provides the following information:

- **State:** A green checkbox indicates the storage disk is online, and a red "x" indicates the storage disk is offline. A yellow triangle icon indicates the storage disk is in an online-pending state. The online-pending state is for storage disks that have deduplication enabled. A deduplication-enabled storage disk is put into an online-pending state during blockpool verification because this process could take a

long time to complete, depending on the blockpool size. Storage disks in the online-pending state are not used for store operations.

- **Storage Disk Name:** Name of the storage disk
- **Deduplication:** Indicates if deduplication is enabled for the storage disk. (This field appears only on 32 bit and 64 bit Linux systems.)
- **Space (GB):** Total amount of space (in gigabytes) on the storage disk. This amount includes space reserved by StorNext.
- **Used Space (GB):** Space (in gigabytes) used on the storage disk
- **Copy #:** The copy number that may be stored on that media
- **# Files:** Total number of files on the storage disk
- **% Savings:** If deduplication is enabled, this field indicates the amount of savings achieved through deduplication. (This field appears only on 32 bit and 64 bit Linux systems.)

StorNext Home Page Drop-Down Menus

The following drop-down menu options located in the grey bar near the top of the home page provide access to StorNext configuration, administration, and reporting options:

- [Config Menu Options](#)
- [Admin Menu Options](#)
- [Reports Menu Options](#)
- [Service Menu Options](#)
- [Help Menu Options](#)



Config Menu Options

The following Config menu options launch the Configuration Wizard and let you perform individual Configuration Wizard tasks:

- **Config Wizard:** Launches the configuration wizard
- **Enter License:** Enter StorNext license information

- **Add File System:** Add a file system to your environment
- **Add Library:** Add a library or vault
- **Add Tape Drive:** Add tape drive
- **Add Media:** Add media
- **Add Storage Disk:** Add a storage disk
- **Add Storage Policy:** Add a storage policy and a directory relation point to a file system
- **E-Mail Notification:** Configure e-mail notifications for System Status Tickets, Backup information, and Policy Class alerts
- **Add Affinity:** Add an affinity to a file system

Note: For more information about the StorNext Configuration Wizard, see [Chapter 2, Data Migration Management](#).

Admin Menu Options

The following Admin menu options enable you to control StorNext's day-to-day operations:

- **Access StorNext Logs:** Access logs of StorNext operations
- **User Access Control:** Control user access to StorNext tasks
- **Download Client Software:** Download SNFS client software. (This procedure is described in the *StorNext Installation Guide*.)
- **Library Operator Interface:** Enter or eject media from the Library Operator Interface
- **Remove/Move Media:** Remove media from a library or move media from one library to another
- **Run Backup:** Run a backup of StorNext software
- **Schedule Events:** Schedule file system events including Clean Info, Clean Versions, Full Backup, Partial Backup, and Rebuild Policy
- **Start/Stop StorNext:** Start or stop the StorNext components

Reports Menu Options

The following Reports menu options allow you to view StorNext reports:

- **SNFS:** View the following SNFS reports:
 - **Affinities:** View the existing affinities for a selected directory in the file system.
 - **File Systems:** View file system statistics including active clients, space, size, disks, and stripe groups.
 - **Stripe Groups:** View statistics for the stripe group, such as space, affinities, and current statuses.
 - **Clients:** View statistics for StorNext clients, including the number of SAN clients and distributed LAN clients, and client performance.
 - **LAN Client Performance:** View information about distributed LAN clients and servers, including read and write speed.
- **Backups:** View a report of StorNext backups
- **Drives:** View tape drive information
- **Files:** View file information
- **Libraries:** View libraries information
- **Library Space:** View information about used library space
- **Media:** View media information
- **Policy Classes:** View policy class information
- **Relations:** View directory/policy class relationship information
- **Requests:** View request information
- **Scheduler:** View scheduler information
- **Storage Disk:** View storage disk information

Note: Detailed descriptions of the Report menu options are located in [Chapter 2, StorNext Reports](#).

Service Menu Options

The following Service menu options help you monitor and capture system status information:

- **Health Check:** Perform one or more health checks on StorNext and view recent health check results

- **State Capture:** Obtain and preserve detailed information about the current StorNext system state
- **System Status:** View tickets indicating faults reported by the StorNext system
- **Admin Alerts:** View informational messages about system activities.

Help Menu Options

The following Help menu options provide access to StorNext documentation, Quantum contact information, and detailed information about the version of StorNext you are using:


- **Online Help:** Provides a listing of StorNext online help topics that can be viewed in a separate browser window. (You can access this same topics list by clicking the Help icon in the upper right corner of the screen.)
- **Documentation:** Provides access to the StorNext documentation set. (Clicking **Help** in the upper right corner of the screen displays this same page.)
- **Support:** Allows you to access Quantum and Technical Support information.
- **About:** Provides detailed information about your version of StorNext, its components, and the system on which it is running. Also lists StorNext patents.

StorNext System Status

The StorNext System Status button located at the bottom center of the StorNext Home Page appears only when there are open service tickets. This same information is always available by selecting **System Status** from the **Service** menu.

Click the **System Status** button to view a list of tickets indicating faults reported by the system. The **Service - System Status** screen appears.

Figure 6 Service - System Status screen



You can display tickets according to Open or Closed status, or view all tickets. If there are multiple screens of tickets, you can display the Previous or Next screen by clicking those buttons. Alternatively, you can go directly to a page by entering that number (such as 3 or 4) and then clicking the Page button.

On the Service - System Status screen you can do any of the following:

- View a specific ticket by selecting the desired ticket and then clicking the **Details** button
- Refresh (update) the list by clicking the **Refresh** button
- Close all tickets by clicking the **Close All** button

StorNext Admin Alerts

The Admin Alert button appears at the bottom of the screen only when an admin alert has been generated. Admin alerts are informational messages about system activities you might want to be aware of, but are not necessarily an error condition. (Admin alerts do not require any action from StorNext users.)


There are different types of admin alerts. Here are some conditions that could produce an admin alert:

- TSM Health Checks disk space warning

- TSM Intrusive Health Checks when drives are mounted
- MSM media console errors
- MSM drive dismount request when drive is already dismounted
- MSM media audit failures

To view admin alerts, select **Admin Alerts** from the StorNext home page's **Service** menu. (Alternatively, you can click the Admin Alert button at the bottom of the screen.) The **Service - Admin Alerts** screen appears.

Figure 7 Service - Admin Alerts Screen



You can display the Previous or Next screen of alerts by clicking those buttons. Alternatively, you can go directly to a page by entering that number (such as 3 or 4) and then clicking the Page button.

On the Service - Admin Alerts screen you can do any of the following:

- View a specific alert by scrolling to the right of the screen (if the alert is longer than can be displayed on one screen)
- Refresh (update) the list by clicking the **Refresh** button
- Delete a selected alert by clicking the **Delete** button
- Delete all alerts by clicking the **Delete All** button

StorNext Server Status

The StorNext Server Status button is located at the bottom right of the StorNext Home Page. This button displays one of three statuses for the file systems residing on the server:

- **Active:** All configured file systems are active
- **Warning:** One or more of the configured file systems have not been mounted or started
- **Stopped:** The storage manager or the file system manager is stopped

Clicking the StorNext Server Status button displays the Start/Stop screen, which enables you to either start or stop StorNext. For more information about starting or stopping StorNext, see [Starting and Stopping StorNext Components](#) on page 25.

Home and Help Links


In the upper right corner of the screen are **Home** and **Help** links. Clicking **Home** returns you immediately to the StorNext home page regardless of your current location.

Clicking **Help** displays a list of current StorNext documentation in pdf format. You can open a particular guide or document by clicking its link.

Application Link

The StorNext application link is located at the bottom left of the StorNext Home Page. Click this link to view information about StorNext.

Figure 8 Application Link



The SNFS and SNSM Home Pages

The Navigation Pane on the left side of the screen contains links to the SNFS and SNSM home pages along with a Home button to return to the main StorNext Home Page.

- [The SNFS Home Page](#) - In SNFS mode, you perform file system-specific tasks. The SNFS Home Page contains several drop-down menus: Config, Admin, and Reports.
- [The SNSM Home Page](#) - In SNSM mode, you perform administrative storage manager tasks. The SNSM Home Page contains several drop-down menus: File, Media, Admin, and Reports.


The SNFS Home Page

The SNFS Home Page contains the following components:

- Drop-down menus that enable you to perform file system-specific administration tasks

- A file system monitor
- System and server status indicators

Figure 9 SNFS Home Page



Note: The figure shows the SNFS Home Page as it appears when it is part of the full StorNext application. The interface for the SNFS standalone application looks slightly different, but has many of the same major features.

SNFS Home Page Drop-Down Menus

Use the drop-down menus on the SNFS home page as you would from the StorNext Home Page when performing file system-specific administration tasks.

- [The SNFS Config Menu](#)
- [The SNFS Admin Menu](#)
- [The SNFS Reports Menu](#)
- [The SNFS Help Menu](#)

Note: You must unmount and stop the file system before using most of the options listed in this section. See [Chapter 2, Common StorNext Tasks](#) for more information.

The SNFS Config Menu

The following options on the SNFS Config menu allow you to add, modify and delete file system configuration files:

- **File Systems:** Add, modify or delete a file system
- **Globals:** Modify global settings
- **Affinities:** Add, modify, or delete an affinity from file systems
- **Disks:** Add or delete a disk to a file system
- **Stripe Groups:** Add, modify, or delete a stripe group from a file system

Caution: Deleting a disk or stripe group requires a remake of the file system and destroys all data on the file system when the task is run.

The SNFS Admin Menu

The following options on the SNFS Admin menu help you manage day-to-day operations of the active file systems:

- **Make File System:** Make a file system

Caution: Making a file system or making/labeling disk devices destroys all data on the disk on which the task is run.

- **Start/Stop File System:** Start or stop a file system
- **Mount/Unmount:** Mount or unmount a file system
- **Label Disk Devices:** Label disk drives
- **Set Affinities:** Create a directory in the file system to the affinity
- **Metadata Dump:** Generate a Metadata dump of a file system

- **Check File System:** Run a check on your file system before expanding the file system or migrating a stripe group.
- **Expand File System:** Add new stripe groups to your file system.
- **Move Stripe Group:** Move data or metadata from one source stripe group to one or more destination stripe groups.

The SNFS Reports Menu

The options on the SNFS Reports menu allow you to view all SNFS reports. These are the same reports described in [Reports Menu Options](#) on page 19.

The SNFS Help Menu

The options on the SNFS Help menu provide access to StorNext documentation, Quantum contact information, and detailed information about the version of StorNext you are using. The SNFS Help menu options are the same as the ones described in [Help Menu Options](#) on page 21.

The SNSM Home Page


The SNSM Home Page contains a Library Monitor and these drop-down menu options:

- [The SNSM File Menu](#)
- [The SNSM Media Menu](#)
- [The SNSM Admin Menu](#)
- [The SNSM Reports Menu](#)
- [The SNSM Help Menu](#)

SNSM Home Page Drop-Down Menus

Use the drop-down menus on this home page as you would from the StorNext Home Page when performing Storage Manager-specific administration tasks.

Figure 10 SNSM Menu



The SNSM File Menu

The following SNSM File menu options allow you to manage file data and directories:

- **Store:** Store files to a storage medium
- **Version:** Show the version(s) of files stored on storage medium
- **Recover File:** Recover deleted files
- **Recover Directory:** Recursively recover deleted directories
- **Retrieve File:** Retrieve truncated files from a storage medium
- **Retrieve Directory:** Recursively retrieve truncated directories from a storage medium
- **Free Disk Blocks:** Truncate files
- **Move:** Move files from one media to another
- **Attributes:** Change file attributes

The SNSM Media Menu

The following SNSM Media menu options help you manage media and libraries:

- **Library:** Perform media movement tasks within a library (Manual Move, Mount, and Dismount)
- **Assign Policy:** Add media types to a policy class
- **Remove:** Remove media from StorNext
- **Assign Policy:** Associate blank media with a policy class
- **Transcribe:** Transcribe (copy) media
- **Attributes:** Change the media's state or attributes
- **Reclassify:** Reclassify a media to a new media class
- **Clean:** Clean a media by policy class, file system, or media identifier

The SNSM Admin Menu

The following SNSM Admin menu options allow you to perform administration tasks:

- **Library:** Perform library tasks such as Config Library, Audit Library, Library State, and Cancel Eject
- **Drive:** Perform drive tasks such as Config Drive, Change Drive State, and Clean Drive
- **Storage Disk:** Perform storage disks tasks such as Config Storage Disk, Change Storage Disk State, and Clean Storage Disk
- **Disk Space:** Perform an immediate file system storage or truncation policy
- **Policy Class:** Add, modify, or delete a policy class
- **Backup:** Configure backup procedure parameters
- **Relation:** Add or remove directory relation points to a policy class
- **Water Mark Parameter:** Set water mark parameters (for more information about watermarks, see **Status**)
- **Config Drive Pool:** Add, modify, or delete drive pools
- **Cancel Request:** Cancel requested operations
- **Alternate Retrieval Location:** Retrieve a copy of a truncated file from a specified foreign machine when copies of the file are not accessible from the local StorNext Storage Manager machine.

The SNSM Reports Menu

The options on the SNSM Reports menu allow you to view all SNFS reports. These are the same reports described in [Reports Menu Options](#) on page 19.

The SNSM Help Menu

The options on the SNSM Help menu provide access to StorNext documentation, Quantum contact information, and detailed information about the version of StorNext you are using. The SNSM Help menu options are the same as the ones described in [Help Menu Options](#) on page 21.


The Configuration Wizard

StorNext includes a Configuration Wizard that guides you through the process of setting up your StorNext system. The wizard includes tasks you would typically perform when you are first configuring your system.

The Configuration Wizard appears automatically when you launch StorNext for the first time. If you do not finish performing all the tasks, the wizard reappears whenever you return to the StorNext home page so you can resume completing tasks where you left off. For example, if you complete tasks 1 through 5, the next time the StorNext wizard appears you will be ready to complete task 6.

If you have not completed all the wizard tasks and do not want the wizard to appear the next time you access the StorNext home page, select the **Don't Show CW Again** option.

Figure 11 StorNext
Configuration Wizard



You can display the Configuration Wizard at any time by selecting **Configuration Wizard** from the StorNext home page's **Config** menu. If you have completed all of the tasks, each task will be marked as Complete. If you have not completed all tasks, the ones you finished will be marked Complete and the wizard will be ready for you to begin the next uncompleted task.

You can perform any of the Configuration Wizard's tasks separately rather than through the wizard. Each of these tasks is selectable from the StorNext home page's **Config** menu.


The initial Configuration Wizard screen has a **Reset** button that allows you to start at the beginning of the wizard.

Figure 12 Config Menu Options



Each of the Configuration Wizard’s tasks are described in detail elsewhere in this guide:

- Step 1: Enter License (See [Entering the StorNext License](#) on page 11)
- Step 2: Add File System (See [Adding a File System](#) on page 12)
- Step 3: Add Library (See [Adding a Library](#) on page 12)
- Step 4: Add Tape Drive (See [Adding a Tape Drive](#) on page 12)
- Step 5: Add Media (See [Adding Media to a Configured Library](#) on page 11)
- Step 6: Add Storage Disks (See [Adding a Storage Disk](#) on page 13)
- Step 7: Add Storage Policy (See [Adding a Storage Policy](#) on page 13)
- Step 8: Email Notification (See [Setting Up E-mail Notification](#) on page 38)



Chapter 3

SNFS-Only GUI Overview

The SNFS GUI is browser-based and can be accessed remotely from any machine with access to the local network and the SNFS server. This chapter describes how to access and navigate through the SNFS-only GUI. If you purchased StorNext, refer to [StorNext GUI Overview](#) on page 7.

This chapter includes the following topics:

- [Accessing the SNFS GUI](#)
- [The SNFS Home Page](#)
- [The SNFS Configuration Wizard](#)

Accessing the SNFS GUI

Use the following procedure to access the SNFS GUI.

- 1 Open a Web browser.

Note: Supported browsers are:

- Internet Explorer 5.5, 6 and 7
- Netscape 7.x
- Mozilla 1.0 and later
- FireFox 1.5 and later, or 2.0 and later


To ensure proper browser operation, all monitors must be set to display at a minimum resolution of 1074 x 768.

- 2 In the browser's **Address** field, type the full address of the machine and its port number, and then press **Enter**. For example: `http://<machine name>:<port number>`. Use the name of the machine and port number you copied when you installed the StorNext software.

Note: Typically, the port number is 81. If port 81 is in use, use the next unused port number. (I.e., 82, 83, etc.)

After you enter the machine name and port number, the following screen appears.

Figure 1 StorNext File System Login Window



- 3 In the **User ID** field, type admin and press **TAB**.
- 4 In the **Password** field, type admin and click **OK**.

Note: For information on changing your password or setting up additional users, refer to the *StorNext Installation Guide*.


The initial StorNext GUI appears.

The SNFS Home Page

The SNFS home page contains the following components:

- A link to the Quantum Home Page
- [The File System Monitor](#)
- [Drop-down Menus and Options](#)
- [Shortcut Menu Options](#)
- [Home and Help Links](#)
- [StorNext Server Status](#)

Figure 2 SNFS Home Page



The File System Monitor

The File System Monitor enables you to view statistics on each configured file system. When you open a browser to access StorNext, the File System Monitor appears at the top of the browser window.

The File System Monitor has a **Refresh** button that allows you to manually update (refresh) the information shown in the monitor. You can also select a rate from the **Refresh Rate** list to automatically refresh at the selected interval:

- No Refresh
- 30 seconds
- 1 minute
- 2 minutes
- 5 minutes
- 10 minutes

Figure 3 File System Monitor

State	File System	Total Space (GB)	Used Space (GB)	Free Space (GB)	Used Inodes	# Store Candidates	# Trunc Candidates	# SAN Clients	# LAN Clients	Status
<input checked="" type="checkbox"/>	snfs1	189.66	4.31	185.35	35	N/A	N/A	1	0	■
<input checked="" type="checkbox"/>	snfs2	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	■

The File System Monitor provides the following information:

- **Refresh:** Click this button to manually refresh the File System Monitor.
- **Refresh Rate:** Set the File System Monitor to automatically refresh with this drop down menu. Options range from **No Refresh** to every **10 Minutes**.
- **State:** A green checkbox indicates the file system is mounted and started, a red “x” indicates the file system is not active or not mounted
- **File System:** The name of the file system (you may see more than one file system being monitored, depending upon how your StorNext system is configured). This pane is scrollable to accommodate numerous file systems.
- **Total Space (GB):** Total space (in GB) for the file system
- **Used Space (GB):** Currently used space (in GB) for the file system
- **Free Space (GB):** Amount of free space (in GB) for the file system

- **Used Inodes:** The number of inodes currently in use on the file system
- **#Store Candidates:** Number of files selected for storage to secondary media. This field does not apply to SNFS-only configurations.
- **#Trunc Candidates:** Number of files that have been stored and meet the criteria to become a truncation candidate. This field does not apply to SNFS-only configurations.
- **# Clients:** The number of StorNext clients (connected via fibre channel or iSCSI) for which you are licensed
- **# Distributed LAN Clients:** The number of StorNext distributed LAN clients for which you are licensed. For more information about distributed LAN clients, see [About Distributed LAN Clients](#) on page 2.
- **Status:** The status shows the system status (usage) in percentage.

Drop-down Menus and Options

The following drop-down menus are located in the grey bar towards the top of the home page and are used to access SNFS configuration, administration, and reporting options:

- [The Config Menu](#)
- [The Admin Menu](#)
- [The Reports Menu](#)
- [The Help Menu](#)

Figure 4 File System Menus



The Config Menu

The following Config menu options enable you to add and modify file systems:

- **File Systems:** Add or delete a file system

- **Globals:** Modify global settings
- **Disks:** Add or delete a disk to a file system
- **Stripe Groups:** Add, modify, or delete a stripe group from a file system
- **Affinities:** Add, modify, or delete an affinity from file systems

Note: Detailed descriptions of the Config menu options are located in [Chapter 6, Managing the File System](#).

The Admin Menu

The following Admin menu options enable you to control day-to-day operations of active file systems:

- **Make File System:** Make a file system
- **Start/Stop File System:** Start or stop a file system
- **Mount/Unmount:** Mount or unmount a file system
- **Label Disk Devices:** Label disk drives
- **Set Affinities:** Create a relation point in the file system to the affinity

Caution: Making a file system or making/labeling disk devices destroys all data on the disk on which the task is run.

Note: Detailed descriptions of the Admin menu options are located in [Chapter 6, Managing the File System](#).

The Reports Menu

The following Reports menu options let you view file system reports:

- **Affinities:** View the Affinities report
- **File Systems:** View the File Systems report
- **Stripe Groups:** View the Stripe Groups report

Note: Detailed descriptions of the Report menu options are located in [Chapter 12, StorNext Reports](#).

The Help Menu

The Help menu provides you with access to StorNext reference material such as this user's guide, the installation guide, and other useful documents.

Shortcut Menu Options

The SNFS Shortcut Menu on the left side of the SNFS home page contains options that let you perform the following Configuration and Administrative tasks:

- **Config**
 - **Configuration Wizard:** Launch the Configuration Wizard. For more information about the Configuration Wizard, see [The SNFS Configuration Wizard](#) on page 40.
 - **Add Affinity:** Add an affinity and directory to a file system. For more information about adding an affinity, see [Adding an Affinity](#) on page 125.
 - **Add File System:** Add a file system. For more information about adding a file system, see [Adding a File System](#) on page 86.
 - **Enter License:** Enter your StorNext license information. For more information about entering a license, see [Entering the StorNext License](#) on page 43.
 - **User Access Control:** Add, modify, or delete the level of user access. For more information about user access control, see [Controlling User Access](#) on page 51.
- **Admin**
 - **Access StorNext Logs:** View logs for SNFS, the Server System, and the StorNext Web Server. For more information about accessing logs, see [Accessing StorNext Logs](#) on page 59.
 - **Download Client Software:** Download the SNFS client software for a specific operating system. For more information about downloading client software, see the *StorNext Installation Guide*.

- **Start/Stop StorNext:** Start or Stop StorNext Software. For more information about starting and stopping StorNext, see [StorNext Server Status](#) on page 38.

Home and Help Links

In the upper right corner of the screen are **Home** and **Help** links. Clicking **Home** returns you immediately to the StorNext home page regardless of your current location.

Clicking **Help** displays a list of current StorNext documentation in pdf format. You can open a particular guide or document by clicking its link.


StorNext Server Status

The StorNext Server Status button is located at the bottom right of the SNFS Home Page. This button displays one of three statuses for the file systems residing on the server:

- **Active:** All configured file systems are active
- **Warning:** One or more of the configured file systems have not been mounted or started
- **Stopped:** The storage manager or the file system manager is stopped


When you click the Server Status button, the **Start/Stop SNFS** screen appears.

Figure 5 Start/Stop SNFS
Screen



- 1 On the **Start/Stop SNFS** screen, select **Start** or **Stop** to manually start or stop the StorNext File System.
- 2 Select **Enable** or **Disable** to enable or disable the feature that automatically starts SNFS upon reboot. Your selection will be reflected the next time you reboot.
- 3 Do one of the following:
 - Click **Cancel** to exit the **Start/Stop SNFS** screen without saving
 - Click **Next** to save your changes and proceed. The **Complete Start/Stop SNFS** screen appears.

Figure 6 Complete Start/Stop
SNFS Screen



- 4 On the **Complete Start/Stop SNFS** screen, do one of the following:
 - Click **Cancel** to exit the **Complete Start/Stop SNFS** screen without saving
 - Click **Back** to return to the **Start/Stop SNFS** screen
 - Click **Next** to proceed. A status window appears
- 5 If you clicked **Next**, click **Close** when the status window displays **Success**.


The SNFS Configuration Wizard

SNFS includes a Configuration Wizard that guides you through the steps necessary to obtain a StorNext license and create a file system. You can complete these steps at any time, and the wizard will resume at the remaining step if you do not finish completing both tasks. If desired, you can select the option **Don't Show CW Again** to prevent the wizard from automatically appearing the next time you launch StorNext.


Alternatively, you can complete the two tasks without launching the Configuration Wizard by clicking the shortcuts **Enter License** or **Add File System**.

Access the Configuration Wizard by clicking the Configuration Wizard shortcut on the home page. The first screen of the **Configuration Wizard** appears.

Figure 7 SNFS Configuration Wizard Screen



For more information about completing the two Configuration Wizard tasks, see [Entering the StorNext License](#) on page 43, and [Adding a File System](#) on page 86.



Chapter 4

Common StorNext Tasks

This chapter provides instructions on performing the following StorNext and SNFS tasks:

- [Entering the StorNext License](#)
- [Controlling User Access](#)
- [Starting and Stopping StorNext Components](#)
- [Accessing StorNext Logs](#)
- [Scheduling StorNext Events](#)
- [Setting Up E-mail Notification](#)
- [Cancelling SNSM Requests](#)
- [Setting up Alternate Retrieval Locations](#)

Entering the StorNext License

Use the Enter License wizard to enter a permanent license string. You can also proceed using the 30-day temporary license that comes with StorNext. You must have a permanent or temporary license to configure or use StorNext.

Note: If you use the temporary license, be sure to obtain a permanent license from Quantum before the 30-day temporary license expires.

To obtain a permanent license, you must contact the Quantum Technical Assistance Center at licenses@Quantum.com and give them the following information:

- The serial number from your product CD or box.
- The number of StorNext SAN clients and distributed LAN clients you want to support.
- The StorNext server identification number. You can find this number on the Configuration Wizard's **Enter License String** screen.

Alternatively, you can obtain a license by going to www.Quantum.com/swlicense and providing the required information.

After the Quantum Technical Assistance Center receives the above information, a representative will send you a license string. Enter this license screen on the **Enter License String** screen to use StorNext with your permanent license.

If you use the temporary license, allow sufficient time for the Quantum Technical Assistance Center to receive your information and send your license string before the 30-day limit expires.


There are two ways to access the Enter License wizard:

- Through the StorNext Configuration wizard
- By selecting **Enter License** from the **Config** menu on the StorNext home page

The only functional difference is that selecting the Enter License wizard from the StorNext Configuration Wizard enables you to generate a 30-day temporary license. This option is not available when you access the Enter License wizard, so if you want to generate a 30-day temporary license you *must* access the Enter License wizard from the StorNext Configuration wizard.


- 1 From the StorNext home page, choose **Enter License** from the **Config** menu. (Alternatively, click **Step 1, Enter License**, in the StorNext Configuration wizard.) The **Enter License - Introduction** screen appears.

Figure 1 Enter License
Introduction Screen




- 2 Click **Next** to continue. The Quantum license agreement appears. You must accept the license agreement in order to continue with the licensing process.

Figure 2 Quantum License Agreement



- 3 Read the license agreement and then click **Accept**. The **Enter License Strings** screen appears. This screen summarizes the information you must send to licenses@Quantum.com in order to receive the license string you enter on this screen.


Figure 3 Enter License Strings Screen



Note: The illustration shows the Enter License Strings screen that appears when you access the Enter License wizard from the StorNext Configuration wizard. The information below the “Distributed LAN Client License String” field does not appear when you access the Enter License wizard from the StorNext home page Config menu.


- 4 If you want to proceed using a temporary license for SAN clients or distributed LAN clients, select the option **Generate 30 day temporary license**. To enter a permanent license, proceed to [Step 8](#) on page 17.
- 5 After selecting **Generate 30 day temporary license**, click **Next** to continue. The **Complete Enter License** screen appears.

Figure 4 Complete Enter License Screen (Temporary)



- 6 On the **Complete Enter License** screen, click **Next** to complete the task, or **Back** to make changes. When you click **Next**, a message reminds you to contact the Quantum Technical Assistance Center within 30 days to receive your permanent license string.

Figure 5 License Reminder




- 7 Click **OK** to close the message box.

- 8 In order to receive your permanent StorNext license you must email to licenses@quantum.com the following information, as listed on the **Enter License Strings** screen ([Figure 22](#) on page 15):
- The StorNext serial number from the StorNext box or CD.
 - The total number of StorNext SAN clients (if any) you want to license for your system. This is the total number of client machines that are connected to a StorNext server through a fibre channel or iSCSI interface.
 - The StorNext server ID displayed on the screen. (In the figure, the ID is **1617GA24E1**.)
 - The total number of distributed LAN clients (if any) you want to license for your system. This is the total number of distributed LAN clients connected to StorNext via a distributed LAN server. For more information about distributed LAN clients, see [About Distributed LAN Clients](#) on page 2.
- 9 Quantum will email you a license string for SAN clients. If you are using distributed LAN clients, they will also send you a separate license string for distributed LAN client usage. When you receive the license strings, copy and paste them into the **License String** field on the **Enter License Strings** screen ([Figure 22](#) on page 15).

Click **Next** to continue. The **Complete Enter License** screen appears.

Figure 6 Complete Enter License Screen (Permanent)



- 10 Review the information on the screen, and then click **Next** to continue, or **Back** to make changes.
- 11 After the status screen informs you that the operation was completed successfully, click **Finish**.

Entering a License String in the .dat File

Alternatively, when you receive your license string via email, you can copy and paste the license string into the `/usr/adic/DSM/config/license.dat` file instead of using the **Enter License Strings** screen. Updating this file enables StorNext to automatically detect the license string when the Configuration Wizard runs.

Here is an example of a StorNext license file with the license string entered (below License Authorization String). This is an example only. Do not enter the license screen shown.

Figure 7 License String
Example

```
# System:          server1
# Identifier:      359061D6AF
# Number Clients: 18
# Expiration Date: None
# Serial Number:  SN09999
# License:
# CAAAS/AAAAE/AJC4J/T6FSA/F7RA9/7XPXA/ZGC2C/9GAGN/JVJ7N/AML8X/JA
#
# License Authorization String:
server 1 359061D6AF 18 CAAASAAAEEAJC4JT6FSAMFRA97XPXAZGC2C9GAGNJVJ7NAML8XJA server1 Quantum

# System:          server1
# Identifier:      359061D6AF
# Number Clients: 32
# Expiration Date: None
# Serial Number:  SN09999
# License:
# AABAA/FJADA/ASF2F/BZ4ZS/CNUDS/JMMSW/PWZG5/4J8FU/FAPKH8TPMN/AA
#
# License Authorization String:
proxy 1 359061D6AF 32 AABAAFJADAASF2FBZ4ZSCNUDSJMMSWPWZG54J8FUFAPKH8TPMNAA server1 Quantum
```

Note: When using this method to obtain your permanent license, you should copy and paste all of the text into the license.dat file, including the portions that are commented out (i.e., preceded by the pound sign, "#"). In particular, be sure to include the serial number.

If you have Distributed LAN clients, you must also copy and paste the text for those clients. These license authorization strings begin with "proxy" as shown in the preceding example.

Controlling User Access

This section describes the following User Access Control options:

- [Changing the Admin Password](#)
- [Adding a New User](#)
- [Modifying an Existing User](#)
- [Deleting an Existing User](#)

Changing the Admin Password

Figure 8 User Access Control Screen


Use this procedure to change the admin password.

- 1 From the StorNext Home Page choose **User Access Control** from the **Admin** menu. The **User Access Control** screen appears.



- 2 From the **User Access Control** screen, select **Admin** in the User List and click **Modify**. The **Modify User "admin"** screen appears.

Figure 9 Modify User Screen



- 3 Enter your new password.
- 4 Confirm the new password by entering it again, and then click **OK**.
- 5 Click **OK** when the Status Screen displays **Success**.

Adding a New User

Use this procedure to add a new StorNext user.

- 1 From the **User Access Control** screen, click **Add**. The **Add New User** Screen appears.

Figure 10 Add New User
Screen

Quantum. Add New User

Enter User Name:

Enter Password:

Re-Enter Password:

Admin Defaults Operator Defaults General User Defaults

StorNext File System Advanced Functions Select All Deselect All

<input type="checkbox"/> Admin: Set Affinities	<input type="checkbox"/> Admin: Check File System	<input type="checkbox"/> Admin: Label Disk Devices
<input type="checkbox"/> Admin: Make File System	<input type="checkbox"/> Admin: Metadata Dump	<input type="checkbox"/> Admin: Mount/Unmount
<input type="checkbox"/> Admin: Start/Stop File Systems	<input type="checkbox"/> Config: Affinities	<input type="checkbox"/> Config: Disks
<input type="checkbox"/> Config: File System	<input type="checkbox"/> Config: Globals	<input type="checkbox"/> Config: Stripe Groups
<input type="checkbox"/> File System Expansion	<input type="checkbox"/> Admin: File System Migration	

StorNext Storage Manager Advanced Functions Select All Deselect All

<input type="checkbox"/> Admin: Audit Library	<input type="checkbox"/> Admin: Backups	<input type="checkbox"/> Admin: Cancel Eject	<input type="checkbox"/> Admin: Cancel Request
<input type="checkbox"/> Admin: Clean Drive	<input type="checkbox"/> Admin: Config Library	<input type="checkbox"/> Admin: Config Drive	<input type="checkbox"/> Admin: Config Drive Pool
<input type="checkbox"/> Admin: Disk Space	<input type="checkbox"/> Admin: Drive State	<input type="checkbox"/> Admin: Foreign File Access	<input type="checkbox"/> Admin: Water Mark Parameters
<input type="checkbox"/> Admin: Policy Classes	<input type="checkbox"/> Admin: Relations	<input type="checkbox"/> Admin: Storage Disk Clean	<input type="checkbox"/> Admin: Storage Disk Config
<input type="checkbox"/> Admin: Storage Disk State	<input type="checkbox"/> Admin: Library State	<input type="checkbox"/> File: Attributes	<input type="checkbox"/> File: Directory Recover
<input type="checkbox"/> File: Directory Retrieve	<input type="checkbox"/> File: Free Disk Blocks	<input type="checkbox"/> File: Move	<input type="checkbox"/> File: Recover
<input type="checkbox"/> File: Retrieve	<input type="checkbox"/> File: Store	<input type="checkbox"/> File: Version	<input type="checkbox"/> Media: Add
<input type="checkbox"/> Media: Attributes	<input type="checkbox"/> Media: Clean	<input type="checkbox"/> Media: Transcribe	<input type="checkbox"/> Media: Dismount
<input type="checkbox"/> Media: Manual Move	<input type="checkbox"/> Media: Mount	<input type="checkbox"/> Media: Move Blank	<input type="checkbox"/> Media: Reclassify
<input type="checkbox"/> Media: Remove			

StorNext Home Functions Select All Deselect All

<input type="checkbox"/> Access StorNext Logs	<input type="checkbox"/> Add Affinity	<input type="checkbox"/> Add Library	<input type="checkbox"/> Add File System
<input type="checkbox"/> Add Media	<input type="checkbox"/> Add Storage Disk	<input type="checkbox"/> Add Storage Policy	<input type="checkbox"/> Add Tape Drive
<input type="checkbox"/> Schedule Events	<input type="checkbox"/> Library Operator Interface	<input type="checkbox"/> Backup	<input type="checkbox"/> Capture State
<input type="checkbox"/> Download Client Software	<input type="checkbox"/> E-Mail Notification	<input type="checkbox"/> Enter License	<input type="checkbox"/> Health Check
<input type="checkbox"/> Remove/Move Media	<input type="checkbox"/> Start/Stop Components	<input type="checkbox"/> System Status	

StorNext Reports Select All Deselect All

<input type="checkbox"/> Affinities	<input type="checkbox"/> Libraries	<input type="checkbox"/> Backups	<input type="checkbox"/> Clients
<input type="checkbox"/> Drives	<input type="checkbox"/> File Systems	<input type="checkbox"/> Files	<input type="checkbox"/> Library Space
<input type="checkbox"/> Media	<input type="checkbox"/> Media Classes	<input type="checkbox"/> Policy Classes	<input type="checkbox"/> Proxy client performance
<input type="checkbox"/> Relations	<input type="checkbox"/> Requests	<input type="checkbox"/> Scheduler	<input type="checkbox"/> Storage Disk
<input type="checkbox"/> Stripe Groups			

OK Cancel

- 2 In the **Enter User Name** field, type the name the new user will enter at the **User ID** field when he or she logs on to StorNext.
- 3 In the **Enter Password** field, type the password the new user will enter when logging on to StorNext.

- 4 In the **Re-enter Password** field, retype the password you entered at the previous field.
- 5 Select the type of access the user will have:
 - **Admin Defaults:** Enables access to the entire StorNext system including SNFS Advanced Functions, SNSM Advanced Functions, StorNext Home Functions, and StorNext Reports
 - **Operator Defaults:** Enables access to most of the StorNext Home Functions and StorNext Reports
 - **General User Defaults:** Enables access to most of the StorNext Reports

Each of the above selections auto-populates the screen to correspond with your selection, but you can customize access by clicking on specific items for the user.

- 6 To simplify assigning access permissions, you can click **Select All** or **Deselect All** for each category. For example, to grant permission to most StorNext reports, click **Select All** and then deselect the reports for which you do not want to grant permission.
- 7 When you are satisfied with the permissions you have assigned, click **OK**.
- 8 Click **OK** when the Status screen displays **Success**. The **User Access Control** screen shows the new user you just added.
- 9 Click **Cancel** to close the window.

Modifying an Existing User

Use this procedure to modify an existing user's StorNext access.

- 1 From the **User Access Control** screen, select a name (other than "admin") from the User List and click **Modify**. The **Modify User** screen appears with the user's name displayed in the header.

Figure 11 Modify User Screen

Quantum **Modify User "sino"**

[Modify Password](#)

[Confirm Password](#)

StorNext File System Advanced Functions Select All Deselect All

<input type="checkbox"/> Admin: Set Affinities	<input type="checkbox"/> Admin: Check File System	<input type="checkbox"/> Admin: Label Disk Devices
<input type="checkbox"/> Admin: Make File System	<input type="checkbox"/> Admin: Metadata Dump	<input type="checkbox"/> Admin: Mount/Unmount
<input type="checkbox"/> Admin: Start/Stop File Systems	<input type="checkbox"/> Config: Affinities	<input type="checkbox"/> Config: Disks
<input type="checkbox"/> Config: File System	<input type="checkbox"/> Config: Globals	<input type="checkbox"/> Config: Stripe Groups
<input type="checkbox"/> File System Expansion	<input type="checkbox"/> Admin: File System Migration	

StorNext Storage Manager Advanced Functions Select All Deselect All

<input checked="" type="checkbox"/> Admin: Audit Library	<input checked="" type="checkbox"/> Admin: Backups	<input checked="" type="checkbox"/> Admin: Cancel Eject	<input checked="" type="checkbox"/> Admin: Cancel Request
<input checked="" type="checkbox"/> Admin: Clean Drive	<input checked="" type="checkbox"/> Admin: Config Library	<input checked="" type="checkbox"/> Admin: Config Drive	<input checked="" type="checkbox"/> Admin: Config Drive Pool
<input checked="" type="checkbox"/> Admin: Disk Space	<input checked="" type="checkbox"/> Admin: Drive State	<input checked="" type="checkbox"/> Admin: Foreign File Access	<input checked="" type="checkbox"/> Admin: Water Mark Parameters
<input checked="" type="checkbox"/> Admin: Policy Classes	<input checked="" type="checkbox"/> Admin: Relations	<input checked="" type="checkbox"/> Admin: Storage Disk Clean	<input checked="" type="checkbox"/> Admin: Storage Disk Config
<input checked="" type="checkbox"/> Admin: Storage Disk State	<input checked="" type="checkbox"/> Admin: Library State	<input checked="" type="checkbox"/> File: Attributes	<input checked="" type="checkbox"/> File: Directory Recover
<input checked="" type="checkbox"/> File: Directory Retrieve	<input checked="" type="checkbox"/> File: Free Disk Blocks	<input checked="" type="checkbox"/> File: Move	<input checked="" type="checkbox"/> File: Recover
<input checked="" type="checkbox"/> File: Retrieve	<input checked="" type="checkbox"/> File: Store	<input checked="" type="checkbox"/> File: Version	<input checked="" type="checkbox"/> Media: Add
<input checked="" type="checkbox"/> Media: Attributes	<input checked="" type="checkbox"/> Media: Clean	<input checked="" type="checkbox"/> Media: Transcribe	<input checked="" type="checkbox"/> Media: Dismount
<input checked="" type="checkbox"/> Media: Manual Move	<input checked="" type="checkbox"/> Media: Mount	<input checked="" type="checkbox"/> Media: Move Blank	<input checked="" type="checkbox"/> Media: Reclassify
<input checked="" type="checkbox"/> Media: Remove			

StorNext Home Functions Select All Deselect All

<input checked="" type="checkbox"/> Access StorNext Logs	<input checked="" type="checkbox"/> Add Affinity	<input checked="" type="checkbox"/> Add Library	<input checked="" type="checkbox"/> Add File System
<input checked="" type="checkbox"/> Add Media	<input checked="" type="checkbox"/> Add Storage Disk	<input checked="" type="checkbox"/> Add Storage Policy	<input checked="" type="checkbox"/> Add Tape Drive
<input checked="" type="checkbox"/> Schedule Events	<input checked="" type="checkbox"/> Library Operator Interface	<input checked="" type="checkbox"/> Backup	<input checked="" type="checkbox"/> Capture State
<input checked="" type="checkbox"/> Download Client Software	<input checked="" type="checkbox"/> E-Mail Notification	<input checked="" type="checkbox"/> Enter License	<input checked="" type="checkbox"/> Health Check
<input checked="" type="checkbox"/> Remove/Move Media	<input checked="" type="checkbox"/> Start/Stop Components	<input checked="" type="checkbox"/> System Status	

StorNext Reports Select All Deselect All

<input checked="" type="checkbox"/> Affinities	<input checked="" type="checkbox"/> Libraries	<input checked="" type="checkbox"/> Backups	<input checked="" type="checkbox"/> Clients
<input checked="" type="checkbox"/> Drives	<input checked="" type="checkbox"/> File Systems	<input checked="" type="checkbox"/> Files	<input checked="" type="checkbox"/> Library Space
<input checked="" type="checkbox"/> Media	<input checked="" type="checkbox"/> Media Classes	<input checked="" type="checkbox"/> Policy Classes	<input checked="" type="checkbox"/> Proxy client performance
<input checked="" type="checkbox"/> Relations	<input checked="" type="checkbox"/> Requests	<input checked="" type="checkbox"/> Scheduler	<input checked="" type="checkbox"/> Storage Disk
<input checked="" type="checkbox"/> Stripe Groups			

- 2 Change the user's password or modify permissions as described in [Adding a New User](#), and then click **OK**.
- 3 Click **Close** when the Status screen displays **Success**.

Note: If there is only one administrator, you can modify only the password. A sole administrator has full access permissions, and you cannot modify these permissions.

Deleting an Existing User

Use this procedure to delete an existing StorNext user.

- 1 From the **User Access Control** screen, select a name from the User List and click **Delete**.
- 2 When asked to confirm that you want to delete the user, click **OK** to proceed or **Cancel** to abort.
- 3 If you click **OK**, click **Close** when the Status screen displays **Success**.

Note: If there is only one administrator, you will not be allowed to delete the administrator.

Starting and Stopping StorNext Components


There are two ways to start or stop the StorNext applications, SNFS and SNSM:

- By using the **Admin** menu's Start/Stop StorNext option
- By clicking the **Server Status** button located at the lower right corner of the StorNext home page

The following procedure describes how to start/stop StorNext components using either of these methods.

- 1 From the StorNext Home Page select **Start/Stop StorNext** from the **Admin** menu. Alternatively, click the Server Status button in the lower right corner of the screen. (This button displays the server's current status, such as Active.) The **Start/Stop StorNext** screen appears.

Figure 12 Start/Stop StorNext Screen




- 2 Select either the **Start** or **Stop** option.
- 3 Select the items you want to start or stop:
 - All Components
 - StorNext File System
 - StorNext Storage Manager

Note: The current status (**Started**, **Stopped**, or **Warning**) is shown next to each StorNext application. A Warning status indicates one of your configured file systems is either not mounted or has not been started.

- 4 Select either **Enable** or **Disable** to enable or disable the feature that automatically starts StorNext at boot time. (Your selection on this screen will be reflected during the next reboot.)
- 5 Click **Next**. The **Complete Start/Stop StorNext Task** screen appears

Figure 13 Complete Start/Stop StorNext Task Screen



- 6 Click **Next**. A status screen appears.
- 7 Click **Finish** when the status screen shows **Success**. The StorNext Server Status in the lower right corner of the screen now displays **Stopped**, **Warning** or **Active**, depending on your action.

Accessing StorNext Logs

You can access and view any of the following types of logs:


- **SNFS Logs:** Logs about each configured file system
- **StorNext Database Logs:** Logs that track changes to the internal database
- **SNSM - File Manager Logs:** Logs that track storage errors, etc. of the Storage Manager

- **SNSM - Library Manager Logs:** Logs that track library events and status
- **Server System Logs:** Logs that record system messages
- **StorNext Web Server Logs:** Various logs related to the web server

Use the following procedure to access the StorNext log files. The process is the same regardless of the type of log you are viewing.


- 1 From the StorNext home page, select **Access StorNext Logs** from the **Admin** menu. The **Select Log** screen appears.

Figure 14 Select Log Screen




- 2 Select the type of log you wish to view, and then click **Next**. The **Select File** screen appears.

Figure 15 Select File Screen



- 3 Select the log file you wish to view, and then click one of the following:
 - **Delete:** Delete the selected log file
 - **Mail:** E-mail the selected log file to Quantum support
 - **Show:** Show the selected log file in a separate window
- 4 If you clicked **Show**, the selected log file appears in a separate window. Click one of the following to navigate around the displayed log file:
 - **Next Page** to view the log's next page
 - **Previous Page** to view the previous page
 - **Top** to move to the beginning of the log
 - **End** to move to the end of the log
 - **Cancel** to close the window

Figure 16 Log File Example



5 When you are finished viewing logs, click **Cancel** to close the **Select File** window.

Scheduling StorNext Events

StorNext events are tasks that are scheduled to run automatically based on a specified schedule. The following events can be scheduled:

- **Clean Info:** This scheduled background operation removes from StorNext knowledge of media.
- **Clean Versions:** This scheduled event cleans old, inactive versions of files.
- **Full Backup:** By default, a full backup is run once a week to back up the entire database, configuration files, and the file system metadata dump file.

- **Health Check:** By default, health checks are set up to run every day of the week, starting at 7:00 a.m.
- **Partial Backup:** By default, a partial backup is run on all other days of the week that the full backup is not run. This backup includes database journals, configuration files, and file system journal files.
- **Rebuild Policy:** This scheduled event rebuilds the internal candidate lists (for storing, truncation, and relocation) by scanning the file system for files that need to be stored.

Note: The Scheduler does not dynamically update when dates and times are changed significantly from the current setting. You must reboot the system for the Scheduler to pick up the changes.


Each of these events initially has a default schedule, but you can configure the schedules to suit your system needs.

Viewing a Schedule

The procedure for viewing an event's existing schedule is the same regardless of the event type.

- 1 From the StorNext Home Page, select **Schedule Events** from the Admin menu. The **Schedule Events** screen appears.

Figure 17 Schedule Events
Screen 1




2 Select an event type:

- Clean Info
- Clean Versions
- Full Backup
- Health Check
- Partial Backup
- Rebuild Policy

3 Click **Configure**. The second **Schedule Events** screen displays the selected event type and any existing schedules.

Figure 18 Schedule Events
Screen 2



- 4 Click **Close** when you are finished viewing the schedule. (You can also click **Back** to return to the previous screen.)

Adding a New Schedule

Use the following procedure to schedule StorNext events. The procedure for adding a new schedule for an event is the same regardless of the event type.

- 1 From the StorNext Home Page, select **Schedule Events** from the **Admin** menu. The first **Schedule Events** screen appears ([Figure 36](#)).
- 2 Select the type of event you want to schedule:
 - Clean Info
 - Clean Versions
 - Full Backup
 - Health Check
 - Partial Backup

- Rebuild Policy
- 3 Click **Configure**. The second **Schedule Events** screen displays the selected event type and any existing schedules ([Figure 37](#)).
 - 4 Click **Add**. The third **Schedule Events** screen appears.

Figure 19 Schedule Events
Screen 3

Schedule Details - Microsoft Internet Explorer

Quantum. Schedule Event

Scheduled Process Name: Full Backup

Name: f_backup_default

Days of Week: --Select Days--
Sunday
Monday
Tuesday
Wednesday

Day of Month: --Select Days--
1
2
3
4

Run Time: 10 :00 PM

Start Window: 3 hours

Back Apply Cancel

- 5 In the **Name** field, enter a name for the new schedule.
- 6 Select either **Days of Week** or **Days of Month**.
 - If you selected **Days of Week**, select the days on which you want the event to run. To select multiple days, hold down the **CTRL** key when you select subsequent days.
 - If you selected **Days of Month**, select the calendar dates on which you want the event to run. To select multiple dates, hold down the **CTRL** key when you select subsequent dates.
- 7 At the **Run Time** field, specify the time of day you want the event to begin.

- 8 At the **Start Window** field, specify the window in which you want the StorNext Scheduler to start the event. The Scheduler attempts to begin the event within the specified Start Window time (e.g, 30 minutes). If the event cannot begin at that time, the Scheduler tries again during the next cycle.
- 9 Do one of the following:
 - Click **Apply** to save your entry and create the new schedule. When the Status screen displays Success, click **Close**.
 - Click **Cancel** to exit the screen without saving your entries. The **Schedule Events** window closes.
 - Click **Back** to return to the previous screen. (When you click **Back** you lose anything you entered on the Add screen.)
- 10 On the first **Schedule Events** screen, click one of the following:
 - **Back**: Go back to the previous screen
 - **Add**: Add a new schedule
 - **Modify**: Change an existing schedule
 - **Delete**: Delete an existing schedule
 - **Reset**: Reset the schedule to the default settings
 - **Close**: Close the window

Modifying an Existing Schedule

The procedure for modifying an existing schedule for an event is the same regardless of the event type.

- 1 From the **StorNext Home Page**, select **Schedule Events** from the **Admin** menu. The first **Schedule Events** screen appears ([Figure 36](#)).
- 2 Select an event type:
 - Clean Info
 - Clean Versions
 - Full Backup
 - Health Check
 - Partial Backup
 - Rebuild Policy

- 3 Click **Configure**. The second **Schedule Events** screen displays the selected event type and any existing schedules ([Figure 37](#)).
- 4 Select the schedule you want to change, and then click **Modify**. The third **Schedule Events** screen appears ([Figure 38](#)).
- 5 In the **Name** field, view or change the name for the new schedule.
- 6 Select either **Days of the Week** or **Days of the Month**.
 - If you selected **Days of the Week**, view or change the days on which you want the event to run. To select multiple days, hold down the **CTRL** key when you select subsequent days.
 - If you selected **Days of the Month**, view or change the calendar dates on which you want the event to run. To select multiple dates, hold down the **CTRL** key when you select subsequent dates.
- 7 At the **Run Time** field, view or change the time of day you want the event to begin.
- 8 At the **Start Window** field, view or change the window in which you want the StorNext Scheduler to start the event. The Scheduler attempts to begin the event within the specified Start Window time (e.g, 30 minutes). If the event cannot begin at that time, the Scheduler tries again during the next cycle.
- 9 Do one of the following:
 - Click **Apply** to save your changes and modify the new schedule. When the Status screen displays **Success**, click **Close**.
 - Click **Cancel** to exit the screen without saving your changes. The **Schedule Events** window closes.
 - Click **Back** to return to the previous screen. (When you click **Back** you lose any changes you made on the **Modify** screen.)

Deleting an Existing Schedule

The procedure for deleting an existing schedule for an event is the same regardless of the event type. For every event type except Partial Backup you must have at least one schedule, so you will not be allowed to delete a solitary schedule.

- 1 From the **StorNext Home Page**, select **Schedule Events** from the **Admin** menu. The first **Schedule Events** screen appears ([Figure 36](#)).

- 2 Select an event type:
 - Clean Info
 - Clean Versions
 - Full Backup
 - Health Check
 - Partial Backup
 - Rebuild Policy
- 3 Click **Configure**. The second **Schedule Events** screen displays the selected event type and any existing schedules ([Figure 37](#)).
- 4 Select the schedule you want to change, and then click **Delete**.

Caution: After you click Delete you will NOT be asked to confirm that you want to delete the schedule. DO NOT click Delete unless you are absolutely certain you want to delete the schedule.

- 5 When the Status screen displays Success, click **Close**.
- 6 Do one of the following:
 - Click **Cancel** to exit the screen. The **Schedule Events** window closes.
 - Click **Back** to return to the previous screen.

Resetting a Schedule

The Reset function deletes ALL existing schedules for an event type and creates one schedule that uses default values. The procedure for resetting schedules for an event is the same regardless of the event type.

- 1 From the **StorNext Home Page**, select **Schedule Events** from the **Admin** menu. The first **Schedule Events** screen appears ([Figure 36](#)).
- 2 Select an event type:
 - Clean Info
 - Clean Versions
 - Full Backup

- Partial Backup
 - Rebuild Policy
- 3 Click **Configure**. The second **Schedule Events** screen displays the selected event type and any existing schedules ([Figure 37](#)).
 - 4 Click **Reset**.
 - 5 When prompted, confirm that you want to delete all existing schedules and create a single schedule that uses default values: click **Yes** to proceed or **No** to abort.
 - 6 Do one of the following:
 - Click **Cancel** to exit the screen. The **Schedule Events** window closes.
 - Click **Back** to return to the previous screen.

Setting Up E-mail Notification

The E-mail Notification feature allows you to specify parties who should be contacted when system alerts or admin alerts occur. You can specify e-mail recipients, alert levels, and information about your e-mail configuration.


Email notification is also an important part of the StorNext backup process. When you select the **Backup** option on the **Configure Email Address** screen (see [Figure 41](#) on page 41,) key information about a completed backup is emailed to the address you specify. This email contains the following important information:

- The required media for restoring from a complete set
- Names of configured storage disks or deduplication storage disks
- Any SNFS configuration files for file systems that are not data migration-enabled are appended to the email

Note: Before configuring e-mail notification, make sure your SMTP server is configured.


- 1 From the StorNext home page, choose **E-mail Notification** from the **Admin** menu. The **Configure E-mail Notification Introduction** screen appears.

Figure 20 Configure E-mail Notification Screen



- 2 On the **Configure E-mail Notification Introduction** screen, review your current e-mail settings (if any) and then click **Next** to continue. The **Configure SMTP E-mail** screen appears.

Figure 21 Configure SMTP E-mail Screen



- 3 On the **Configure SMTP E-mail** screen, enter the fields related to your e-mail system configuration:
 - **SMTP Server:** Enter the identification for the server that stores and processes your e-mail account information. This might be a valid server name or an IP address.
 - **Authentication:** If your e-mail provider requires a password upon sign on, select the **Password** option. Otherwise, select **None**.
 - **Account:** Enter a valid e-mail account for outgoing e-mail messages.
 - **Password:** Enter the e-mail account's sign-on password, if required.
 - **Sender Address:** Enter the e-mail address for the entity sending alert messages to recipients.
 - **Send Test E-mail to:** Enter an e-mail address to which you can send test messages in order to confirm successful configuration.

- 4 Click **Next** to continue. The **Configure E-mail Addresses** screen appears.

Figure 22 Configure E-mail Addresses Screen

Configure E-mail Addresses
Service Tickets require an alert level. E-mail for Service Tickets will be sent for alerts at the specified alert level and higher. Policy class E-mails require a policy class.
If Notify Quantum Technical Assistance Center on Service Ticket is checked, an E-mail will be sent to QUANTUM when a Service Ticket is generated.

E-mail:

Admin Alerts
 Backups
 Service Tickets: --Select Alert Level--
 Policy Class: --Select Policy Class--


Type	Properties	E-Mail Addresses
Service Ticket	high	abc@quantum.com
Admin Alerts		abc@quantum.com

Notify Quantum Technical Assistance Center on Service Ticket

- 5 On the **Configure E-mail Addresses** screen, add e-mail recipients by entering the following fields:
 - **E-mail:** Enter the e-mail address of the person who should receive e-mail alerts.
 - **Admin Alerts:** Select this option to receive e-mail whenever an admin alert is generated. For more information about Admin Alerts, see [StorNext Admin Alerts](#) on page 22.
 - **Backups:** Select this option to receive e-mail after a backup has occurred on your system.
 - **Service Tickets:** Select this option to receive e-mail when a service ticket for your system is generated. Notifications for service tickets will be sent for events at the specified alert level and higher. You must specify an alert level.

- **Policy Class:** Select this option to receive e-mail about policy class. You must specify a policy class.
 - **Notify Quantum on Service Ticket:** Select this option to automatically send the Quantum Technical Assistance Center a message when a service ticket is generated. (The Quantum Technical Assistance Center is not notified about admin alerts if you selected that option.)
- 6 Click **Add** to add to the list of e-mail recipients the e-mail recipient whose information you just entered. Or, select a previously added e-mail recipient from the list and click **Delete** to remove that recipient.
 - 7 If necessary, add additional e-mail recipients by repeating steps 4 and 5. Click **Next** to continue. The **Complete E-mail Configuration** screen appears.

Figure 23 Complete E-mail Configuration Screen



- 8 Review your selections and do one of the following:


- Click **Back** to change information you entered, or add or remove another e-mail recipient.
 - Click **Next** to continue.
- 9 When the Status screen informs you that your e-mail notifications were successfully processed, click **Close**.

Cancelling SNSM Requests

If your system includes StorNext Storage Manager, the Cancel Request option (accessible from the SNSM home page's Admin menu) allows you to cancel a pending mount request that has not yet been executed by the system. On the **Cancel Request** screen you can view and select one or more requests, which are listed according to request ID and request type.

- 1 From the SNSM home page, choose **Cancel Request** from the **Admin** menu. The **Cancel Request** screen appears.

Figure 24 Cancel Request Screen



- 2 Select from the **Select Request ID** list the pending requests you want to cancel. If desired, you can click **Select All** to select all requests.
- 3 Click **Apply** to continue.
- 4 When the Status screen informs you that the selected requests have been successfully cancelled, click **OK**. The requests you selected are now removed from the list on the Cancel Request screen.

Setting up Alternate Retrieval Locations

In situations where file retrieval fails because the normal file copies cannot be retrieved from the machine on which StorNext Storage Manager resides, this feature enables you to retrieve a copy of the truncated file from a different machine.

For example, if StorNext creates two copies of each file, when retrieving a truncated file StorNext tries to retrieve Copy One and then Copy Two. If neither of these copies can be retrieved and this feature is not enabled, the retrieval fails. However, if this feature is enabled for the file system, after retrieving Copy Two fails Storage Manger tries to retrieve the file from the alternate machine you specified during feature setup.

This feature applies only to managed file systems that have at least one configured policy class.

For this feature to work correctly, it is your responsibility to make sure all files you might want to retrieve are copied to the alternate machine. Otherwise retrieval will fail when StorNext attempts to retrieve the file from the alternate location and cannot find the file.

Caution: If a file on the original machine is renamed after it was stored by StorNext, an alternate retrieve will fail for that file.

If a file on the original machine is not the latest version (for example, if the fsrestore command was used to restore a previous version,) whichever file has the correct pathname will be retrieved from the alternate machine.

Setting up this feature consists of the following steps:

- Enter the IP address of the alternate machine you want to use for retrieval. (This alternate machine may or may not contain StorNext. StorNext is not required on an alternate machine.)
 - Enable the feature for the file system(s) desired
 - Specify the path on the remote machine for the file system to access
- 1 From the SNSM home page, choose **Alternate Retrieval Location** from the **Admin** menu. The **Alternate Retrieval Location** screen appears.

Figure 25 Alternate Retrieval Location Screen

The screenshot shows the 'Alternate Retrieval Location' configuration page in the StorNext Admin console. The page has a navigation menu on the left with 'Home', 'SNFS', and 'SNSM' options. The main content area is titled 'Alternate Retrieval Location' and contains the following elements:

- Instructions:** 'Use Alternate Retrieval Location to specify a remote retrieval location to use in situations where files on tape or a storage disk cannot be accessed.'
- List of instructions:**
 - Enter the IP address or host name for the remote node.
 - For each managed file system with a configured policy class, enter the remote node pathname the file system should access
 - Alternate Retrieval Location will be used only on file systems that are enabled
- Input field:** 'Remote Node IP/Host Name' with an empty text box.
- Table:** A table with columns 'Enable', 'File System Name', 'Local Path', and 'Remote Path'. The first row shows 'snfs:1' as the File System Name and '/stornext/snfs:1' as the Local Path. The 'Enable' checkbox is unchecked, and the 'Remote Path' field is empty.
- Buttons:** An 'Apply' button with a green checkmark icon.


The bottom of the screen shows the 'StorNext' logo and the user 'kazar' with an 'Active' status indicator.

- 2 At the **Remote Node IP** field, enter the IP address or host name of the remote machine from which you want to retrieve truncated files in the event that normal retrieval fails.
- 3 Select the **Enable** field to enable the feature for the file system.
- 4 Locate the managed file system you want to use, and enter at the **Remote Path** field the pathname on the remote machine for the file system.
- 5 Click **Apply** to continue.
- 6 When the Status screen informs you that the feature has been successfully enabled, click **OK**. The feature is now enabled for the file system(s) you specified.

Modifying or Disabling This Feature

As your needs or configuration changes, you can modify the remote machine information or disable the feature on one or more file systems.

- 1 From the SNSM home page, choose **Alternate Retrieval Locations** from the **Admin** menu. The **Alternate Retrieval Locations** screen appears.
- 2 As needed, change the information in the **Remote Node IP** and **Remote Path** fields.
- 3 To disable the feature on a particular file system, locate the desired file system name and remove the checkmark from its **Enable** field.
- 4 Click **Apply** to continue.
- 5 When the Status screen informs you that your changes have been successfully completed, click **OK**.



Chapter 5

Backing Up StorNext

This chapter provides instructions on backing up the StorNext software.

This chapter covers these topics:

- [Types of StorNext Software Backups](#)
- [Performing a StorNext Software Backup](#)
- [Managing the Backup Policy](#)

Types of StorNext Software Backups

There are two types of StorNext software backup: Full and Partial. By default, a full backup runs once a week, and partial backups are scheduled to run on all other days of the week that the full backup is not run. If you would like to change the backup schedule, see [Managing the Backup Policy](#) on page 14.

A Full Backup includes:

- The StorNext database
- Configuration files
- File system metadata dump file (after journal files are applied)

A Partial Backup includes:

- StorNext database journals
- Configuration files
- File system journal files

Note: These backups DO NOT back up user data. This procedure backs up only StorNext-related system files.

Setting up Email Notification For Backup

Before you run a StorNext backup, be sure to set up email notification as described in [Setting Up E-mail Notification](#) on page 38.

Email notification is an important part of the StorNext backup process. When you select the **Backup** option on the **Configure Email Address** screen (see [Figure 41](#) on page 41), key information about a completed backup is emailed to the address you specify. This email contains the following important information:


- The required media for restoring from a complete set
- Names of configured storage disks or deduplication storage disks
- Any SNFS configuration files for file systems that are not data migration-enabled are appended to the email

Performing a StorNext Software Backup

Use the following procedure to manually run a StorNext backup.


- 1 From the StorNext home page select **Run Backup** from the **Admin** menu. The **Backup StorNext** screen appears.

Figure 1 Backup StorNext
Screen



- 2 Select the type of backup you want run (**Full** or **Partial**), and then click **Next**. The **Complete Backup Task** screen appears.

Figure 2 Complete Backup
Task Screen




- 3 Click **Next** to start the backup.
- 4 Click **Finish** when the Status screen displays success.

Managing the Backup Policy

The Backups option on the SNSM home page's Admin menu allows you to manage your backup policy by setting the media type for each copy. You can also specify the maximum number of file versions to maintain, and also indicate where backups reside on your managed file system.

- 1 From the SNSM home page, choose **Backups** from the **Admin** menu. The **Backup Policy** screen appears.

Figure 3 Backup Policy Screen




2 Enter the following information:

- **File Copy 1 - 4:** For each copy number (1 - 4,) select the media type. Only the media types applicable to your configuration are selectable.
- **Max Backup Sets:** If desired, specify the maximum number of backup sets you want to create.
- **Drivepool to Use:** Select the location where backups should reside on your managed file system.
- **Checksum Generation:** Select this option if you want to generate a checksum.
- **Checksum Validation:** Select this option to enable checksum validation.
- **Managed File System to use for Backup:** Select the name of the file system to use for the backup.

3 Click **Apply** to save and apply the backup policy.

4 After the Status screen informs you that the operation was performed successfully, click **OK**.



Chapter 6

Managing the File System

This chapter describes how to manage file systems by performing these major configuration and administration tasks:

- [Working With File Systems](#)
- [Managing File System Operations](#)
- [Working With Disks](#)
- [Working With Stripe Groups](#)
- [Working With Affinities](#)
- [Using the SNSM File System Functions](#)
- [Understanding Resource Allocation](#)
- [Performing File System Expansion](#)
- [Performing Stripe Group Movement](#)

Working With File Systems

This section includes the following file-system related tasks:

- [Adding a File System](#)
- [Creating a File System From SNFS](#)

- [Modifying a File System](#)
- [Deleting a File System](#)


Note: This section describes how to create, modify, expand, and delete a file system using the GUI. For instructions on using the CLI to accomplish the same tasks, see [Using The Command Line Interface](#).

Adding a File System

The following procedure describes how to create an empty file system. The number of file systems you can add is limited only by the number of disks available for configuration.


- 1 From the StorNext home page, choose **Add File System** from the **Config** menu. The **File System - Introduction** screen displays both configured file systems and no. of disks available for configuration.

Figure 1 File System -
Introduction Screen



2 Click **Next** to continue. The **Add New File System** screen appears.

Figure 2 Add New File System
Screen



- 3 On the **Add New File System** screen, click **Browse** to enter a mount point (directory) for the file system. This allows you to navigate to an existing directory or create a new one. The **Directory Browser** screen appears.

Figure 3 Directory Browser Window



- 4 Select from the **Select Directory** list an existing directory in which you want to create the file system.

Figure 4 Select Directory Window



- 5 To create a new directory underneath the selected directory, click **Create Directory** and enter the new directory name. Click **OK** to continue. The new directory is shown in the Directory Browser's **Current Directory** Field.
- 6 Click **OK** to accept the new directory. The new directory is shown on the **Add New File System** screen.
- 7 If desired, select the **Enable Data Migration** option. Select this option if you want this file system to be managed with automatic data movement between the primary disk storage and secondary storage (either disk or tape). If you do not enable this option, this file system remains unmanaged and does not move data to the tape library.


Caution: Do not select the Enable Data Migration option if the file system will be used as a storage disk.

- 8 If desired, select the **Enable Distributed Lan Server on this machine** option if you want the machine on which the file system is located to act as a Distributed Lan server. (This option appears only for Linux machines.)

Note: When you enable the Distributed LAN Server, all configured IPV4 addresses will be used on the machine for the Distributed LAN Server Ethernet traffic.

- 9 Click **Next** to proceed from the **Add New File System** screen. The **Disk Settings** screen appears.

Figure 5 Disk Settings Screen



- 10 On the **Disk Settings** screen, type valid values and click **Next**.

- **Block size** field: The block size in bytes for the file system. The block size is the minimum unit of data accessed from physical devices. The default value of 16384 bytes is the recommended setting for best overall efficiency.


Note: If the file system you are adding will be used for deduplication-enabled storage disks, you must accept the default value of 16384 bytes.

- **Stripe group** field: The number of stripe groups for the file system. Selecting a value greater than 1 enables customization to optimize performance and use StorNext features such as dedicated Metadata, Journal and User Data stripe groups.

Note: Quantum recommends that metadata and journals be on a stripe group separate from data. (In some cases, metadata and journals should be on separate stripe groups as well.)

- 11 When a message reminds you to select more than one stripe group if you want an alternate configuration other than the one that provides journal, metadata, and user data, click **OK** to continue. The **Customize Stripe Group** screen appears.

Figure 6 Customize Stripe Group Screen



12 Enter values for the **Customize Stripe Group** screen.

- **Name Field:** The name of the stripe group.
- **Select disks list:** The disks available to assign to the stripe group. You must select at least one disk for each stripe group.
- **Label Type:** If you plan to create LUNs larger than 2TB, you must specify the EFI label type when configuring a file system.

VTOC labels were used for all operating systems in previous StorNext and Xsan releases, and are still required for the SGI IRIX operating system, Solaris releases prior to Solaris 10 Update 2, and LUNs less than 1TB.

EFI labels are required if you plan to create LUNs that are larger than 2TB. (For Solaris, EFI labels are also required for LUNs with a raw capacity greater than 1TB.) EFI labels will not work with the IRIX operating system.

The correct value is automatically selected when you reach the **Customize Stripe Groups** screen, so you can accept the default value unless you have a reason to change the label type.

For more information about 2TB LUN requirements, see the *StorNext Installation Guide*.

- **Label Help:** Click this link to display guidelines for determining whether to select VTOC or EFI labels. The matrix looks like this:

Figure 7 Label Help Screen

Quantum. Platform Support for Large LUNs

This is a list of operating systems that have limitations for labeling disk devices

Operating System	<1TB LUN	1-2TB LUN	>2TB LUN
Restricted support of >2TB LUNs			
AIX	VTOC, EFI	VTOC, EFI	EFI, sVTOC (Note 1)
Support for first 2TB of >2TB LUNs			
Apple Xsan 1.2	VTOC	VTOC	sVTOC
Apple Xsan 1.3 on OS X 10.3 Panther	VTOC, EFI	VTOC, EFI	sVTOC, EFI
HP-UX	VTOC, EFI	VTOC, EFI	sVTOC, EFI
Windows XP 32 bit	VTOC, EFI	VTOC, EFI	sVTOC, EFI
No support of >2TB LUNs			
IRIX	VTOC	VTOC	
Linux 2.4	VTOC, EFI	VTOC, EFI	
Solaris 9 w/Big LUN Patch	VTOC, EFI	EFI	
Solaris 10 vanilla	VTOC, EFI	EFI	
Solaris 9 vanilla	VTOC		

Note 1: AIX appears to be limited to LUNs 2.2TB or smaller

Key	
VTOC	VTOC label
EFI	EFI label
sVTOC	'short' VTOC label on >2TB LUN

DISCLAIMER: While every effort has been made to ensure the accuracy of this information, it is subject to change and should be verified with each particular system vendor.

- **Stripe breadth** drop-down menu: The stripe breadth for the file system. The stripe breadth is the number of kilobytes (KB) that is read from or written to each disk in the stripe. For a typical StorNext installation, 64KB is the recommended setting.


Note: If the file system you are adding will be used for deduplication-enabled storage disks, you must accept the default value of 64 kilobytes.

- **Metadata, Journal, and User Data** checkboxes: Enable one or more of these options (check the boxes) to create a location for metadata, journaling or user data.
 - To enable metadata to be placed on the stripe group, select the **Metadata** checkbox.
 - To enable journaling to be placed on the stripe group, select the **Journal** checkbox.
 - To enable user data to be placed on the stripe group, select the **User Data** checkbox.

Note: Quantum recommends that your user data be on a different stripe group than your metadata and journal data.

- 13 If you selected multiple stripe groups on the **Customize Stripe Group** screen, repeat [Step 12](#) on page 92 for each stripe group.
- 14 On the **Customize Stripe Group** screen, type valid values and click **Next**. The **Complete File System Task** screen appears.

Figure 8 Complete File System Task Screen



- 15 Review your selections. Click **Next** to complete the task or **Back** to make changes.
- 16 After the status screen informs you that the file system was successfully added, click **Next**.

Creating a File System From SNFS

An alternative way to create a file system is to do so from the SNFS home page. However, this method is a more advanced task, and requires steps that are simplified when you create a file system through the Configuration Wizard (or by choosing **Add File System** from the StorNext home page's **Admin** menu). For this reason, the preferred method of creating a file system is the one described in [Adding a File System](#) on page 86.

The following steps are required to create a file system from the SNFS home page:

1. Add a file system.

2. Modify the file system's global settings.
3. Add disks to the file system.
4. Add stripe groups for the file system.
5. Make the file system.
6. Start the new file system.
7. Mount the new file system.
8. Add affinities to the new file system.

Each of the above steps is explained in the following procedure.

- 1 From the SNFS home page, choose **File System** from the **Config** menu. The **Configure File System** screen appears.

Figure 9 Configure File System Screen



- 2 Click **Add** to continue. The **Add File System** screen appears.
- 3 Type a name and mount point for the new file system and click **OK**. A status screen appears.
- 4 When the status screen indicates that the file system has been added, click **Close**.

- 5 Configure global settings for the file system as explained in [Making Global Changes](#) on page 104.
- 6 Add disks to the file system as explained in [Adding Disks](#) on page 113.
- 7 Configure stripe groups for the file system as explained in [Adding a Stripe Group](#) on page 117.

Note: Make sure that each stripe group is associated with no more than one affinity when using disk-to-disk migration.

- 8 Make the file system as described in [Making a File System](#) on page 110.
- 9 Start the file system as described in [Starting and Stopping the File System](#) on page 111.
- 10 Mount the file system as described in [Mounting or Unmounting a File System](#) on page 112.
- 11 Add affinities to the file system as described in [Modifying an Affinity](#) on page 135.

Note: When using disk-to-disk relocation you can define a maximum of two affinities per file system.

Modifying a File System

This section describes how to modify an existing file system's configuration. Changes to an existing file system include adding or modifying a stripe group, adding disks, adding affinities, and performance tuning.

Note: This procedure assumes the file system exists with at least two stripe groups.


- 1 Unmount the file system as described in [Mounting or Unmounting a File System](#) on page 112.
- 2 Stop the file system as described in [Starting and Stopping the File System](#) on page 111.
- 3 Make the appropriate changes to the file system. For more information, refer to:

- [Making Global Changes](#) on page 104
 - [Modifying an Affinity](#) on page 135
 - [Adding Disks](#) on page 113
 - [Adding a Stripe Group](#) on page 117
 - [Modifying a Stripe Group](#) on page 121
- 4 Start the file system as described in [Starting and Stopping the File System](#) on page 111.
 - 5 Mount the file system as described in [Mounting or Unmounting a File System](#) on page 112.

Alternatively, you can modify the file system by using the Modify button on the Configure File System screen to change the file system name or mount point, or to enable the machine as a Distributed LAN server.

- 1 Make sure the file system you want to modify is unmounted as described in [Mounting or Unmounting a File System](#) on page 112.
- 2 Select from the File Systems list the file system to modify, and then click **Modify**. The **Modify File System** window appears.

Figure 10 Modify File System Window



- 3 Make the desired changes to the file system:
 - **Name:** The name of the file system (display only; cannot be changed)
 - **Mount Point:** the file system's mount point location (display only; cannot be changed)

- **Enable Distributed Lan Server on this machine:** Select this option to let the machine on which the file system is located to act as a Distributed LAN server. (This option appears only for Linux machines.) If you no longer wish the machine to be a Distributed LAN server, deselect this option.

Note: When you use this option to configure a machine as a Distributed LAN server, all available IPV4 addresses are used.

Deleting a File System

Use this procedure to delete a file system.

Note: If you want to delete an unmanaged file system that contains storage disks or deduplication-enabled storage disks, you will not be allowed to delete the file system until you first delete all of those sdisks or dedup sdisks.

- 1 From the SNFS home page, choose **File System** from the **Config** menu. The **Configure File System** screen appears.
- 2 Select the file system you want to delete, and click **Delete**.

If you are deleting a managed file system, you are warned that deleting this file system will delete all associated data and you will not be able to recover it. You are also warned that deleting this managed file system restarts the StorNext Storage Manager.

If you are deleting a non-managed system, you are warned that all files and directories in the file system will be permanently deleted. Click **OK** to Continue.

- 3 Click **OK** to close the warning message window.
- 4 After the Status screen informs you that the file system was successfully deleted, click **Close**.

Managing File System Operations

This section describes tasks that can be performed on an individual basis when you administer your file system. Except where noted, tasks can be completed in any order, and do not require that other tasks be successfully completed. This section includes these topics:

- [Disk Device Labeling](#)
- [Making Global Changes](#)
- [Working with the fsnameservers File](#)
- [Making or Unmaking a File System](#)
- [Starting and Stopping the File System](#)
- [Mounting or Unmounting a File System](#)

Disk Device Labeling

Each drive used by SNFS must be labeled. (A new drive must be labeled only one time.) You can label a drive from any StorNext server or client that has a fibre channel (FC) connection to the drive.

Caution: Labeling a disk device will result in a complete loss of data on that disk device.


Selecting a Device to Label, Unlabel, or Probe

Use the following procedure to select a disk device for labeling, unlabeled, or probing.

- 1 From the SNFS home page, choose **Label Disk Devices** from the **Admin** menu. The **Label Disk Device** screen appears.

Note: A device whose name is not selectable in the Disk Devices list is currently in use by a configured file system.

Figure 11 Label Disk Device Screen



From this screen, you can perform these tasks:

- **Select Label Type** - Specify VTOC or EFI label types for your disks. (The default is EFI labels.) EFI labels are required if you plan to create LUNs that are larger than 2TB. (For Solaris, EFI labels are also required for LUNs with a raw capacity greater than 1TB.) EFI labels will not work with the IRIX operating system. If you need help determining label type, click the **Label Help** link.

VTOC labels were used for all operating systems in previous StorNext and Xsan releases, and are still required for the SGI IRIX operating system, Solaris releases prior to Solaris 10 Update 2, and LUNs less than 1TB.

- **Label** - Label any unused devices or relabel any device
- **Unlabel** - Unlabel any device in use by the file system
- **Probe** - Pulse a disk's activity light

Caution: If you unlabel a device, all data on that device will be lost. Additionally, the unlabeled device will no longer be used by a StorNext file system until it is relabeled.

- 2 From the **Disk Devices** list, select the device you want to label, unlabel, or probe, and then click **Label**.

Labeling a Device

Use the following procedure to label a disk device.

- 1 Specify the label type by choosing VTOC or EFI at the Select Label Type field.
- 2 Click **Label**. A message window warns you that the device might have a file system on it, and that labeling the device will destroy any data.

Figure 12 Labeling Warning



- 3 Verify that the disk you are labeling is empty, and then click **OK** to close the message window. The **Label Name** window appears.

Figure 13 Label Name Window



- 4 Enter a name for the device, or accept the displayed default name and then click **OK**. A message window warns you all data will be lost if you label this device.

Figure 14 Data Loss Warning



- 5 Click **OK** to continue and label the device. The **Label Disk Devices Status** window appears.
- 6 Click **Close** when the status displays **Success**. The **Label Disk Device** screen now includes the label you created.

Caution: Quantum recommends that you reboot nodes after you label or relabel a disk. Rebooting ensures that the system recognizes the labeled or relabeled disk.

Unlabeling a Device

Use the following procedure to unlabel a disk device.

- 1 After selecting from the Disk Devices box the device you want to unlabel, click **Unlabel**. A message warns you that the device will be unusable after unlabeling.

Figure 15 Unlabel Warning



- 2 Click **OK** to close the warning window and proceed, or click **Cancel** to abort the unlabeled process. If you click **OK**, a status window appears.
- 3 Click **Close** when the status displays **Success**. The **Label Disk Device** screen now shows the device with no label name.

Note: If you decide later to make an unlabeled device usable by the StorNext File System, you must first relabel the device. The relabeling process is identical to labeling initially, as described in [Labeling a Device](#) on page 102.

Probing a Device

Use the following procedure to probe a disk device.

- 1 After selecting from the Disk Devices box the device you want to probe, click **Probe**. The **Probe Disk Device Status** window appears.
- 2 Click **Close** when the status displays **Success**.

Note: The probe should activate the light on the disk or RAID.

Making Global Changes

The global section of the file system configuration file contains general parameters that control system performance, components related to the file system's resource consumption, and whether features are enabled or disabled.

For most of these parameters, restarting the File System Manager (FSM) causes the modified parameters to take effect. However, the **File System Block Size** and **Windows Security** parameters require that the file system be remade before they take effect. Remaking the file system results in data loss, so you should carefully plan the initial configuration of these two parameters in order to reduce the number of file system remakes. If a parameter change requires a file system remake, the system notifies the administrator in the system log.


The global section also contains several parameters that can dramatically improve or degrade system performance, so you should exercise caution when modifying performance parameters.

Before making any changes to the file system’s configuration, carefully review the `cvfs_config(4)` man pages or the “CVFS Configuration File” help file.

The following task describes how to modify global configuration settings. These settings affect all stripe groups in the file system.

- 1 From the SNFS home page, choose **Globals** from the **Config** menu. The **Modify Global Settings** screen appears.

Figure 16 Modify Global Settings Screen



If you are configuring global settings as part of the procedure for making a new file system, the list at the top of the screen shows your current step in the process. Previous and next steps are hyperlinked, so you can click any step to move to that task.

- 2 In the **File Systems** drop-down menu, select a file system.
- 3 In the remaining fields, choose or enter values and options for the file system you selected.

- **File System Block Size:** This value defines the granularity of the file system's allocation size. The default setting is 16,384. The block size must be specified in powers of 2.
- **Inode Cache Size:** This value defines the number of inodes that can be cached in the SNFS server. The default setting for the cache size is 16, and the minimum is 8.
- **Maximum Log Size:** This value defines the maximum number of bytes (size) to which a SNFS Server log file can grow. When the log file reaches the specified size, it is rolled and a new log is started. In this situation, the two log files could use twice the maximum log size space specified in this field. The range is from 1 to 256 megabytes.
- **Maximum Number of Logs:** This value determines the number of rolled logs kept. Choices range from 1 to 10.
- **Journal Size:** This value controls the size of the file system journal. The range is 1 to 256 megabytes.
- **Thread Pool Size:** This value defines the number of client pool threads to be activated and used by the SNFS server. This setting affects system performance. There should be at least two threads per client. Increasing the number of threads will improve file system response time in operations that affect allocation and metadata functions. The range is from 16 to 1024 threads.
- **Migrating File System:** Enable this option (check the box) if the data on the file system should be migrated to tertiary storage. Migration cannot be disabled once it is enabled.
- **Global Super User:** Enable this option (check the box) to allow a user with super-user privileges to assert these privileges on the file system.

Note: If the **Global Super User** option is enabled, super users have global access rights on the file system. This selection is the same as the `maproot=0` directive in the Network File System (NFS).

If the **Global Super User** option is not enabled, super users can modify only files they can access, like any other users.

- **Windows Security:** Enable this option (check the box) to allow all Windows clients to use Microsoft Windows native security. If this

option is enabled, then a Windows administrator may set Windows security on all files and directories for any user or group defined on the local Windows client or Windows domain.

- **Use Physical Memory Only:** When this option selected, the file system will use only physical memory, not swapped or paged.
- **Quotas:** Select this option to enable enforcing quotas for users and groups.

Note: Quotas are based on actual usage, and are not enforced based on space allocated.

- **Inode Expansion Blocks:** Use these fields to configure the minimum (floor), maximum (ceiling), and increment of the block allocation size for a dynamically expanding file.
 - **Minimum:** This value specifies the minimum number of blocks allocated when a file requires additional space.
 - **Maximum:** This value specifies the maximum number of blocks allocated on subsequent expansion. The range is from 1 to 32768 blocks.
 - **Increment:** When the allocation space is exhausted, this value specifies the increment added to the last allocation size, up to the maximum number of file system blocks.
 - **Max Connections:** Specify the maximum number of simultaneous connections for the file system.
 - **Reserved Space:** This option enables delayed allocations on clients. Reserved space is a performance feature that allows clients to perform buffered writes on a file without first obtaining real allocations from the metadata controller. The allocations are later performed when the data is flushed to disk in the background by a daemon performing a periodic sync.

If the Reserved Space option is not enabled, slightly more disk space can be used at the expense of buffer cache performance, which could be adversely affected and cause fragmentation.
- **Debug Log Settings:** Settings to turn on debug functions for the file system server. The log information may be useful if a problem occurs. A Quantum Technical Assistance Center representative may ask for certain debug options to be activated to analyze a file system or hardware problem.

- **Disable Debugging:** Disables detailed file system debug tracing. When debug tracing is enabled, file system performance could be significantly reduced.
 - **LDAP Configuration**
 - **UNIX File Creation Mode on Windows** - Mode bits for UNIX files
 - **UNIX Directory Creation Mode on Windows** - Mode bits for UNIX directories
 - **UNIX ID Fabrication on Windows** - Allows you to enable or disable using fabricated IDs on a per-file system basis. If enabled, Windows user IDs are mapped using fabricated IDs.

The default value for enabling fabrication is based on the type of StorNext server you are using. On Windows the default is No.
 - **UNIX Nobody UID on Windows** - UNIX user ID to use if no other mapping can be found
 - **UNIX Nobody GID on Windows** - UNIX group ID to use if no other mapping can be found
- 4 Click **Apply**. The **Modify Global Setting Status** screen appears.
 - 5 After the status screen indicates that the global settings were successfully modified, click **Close**.

Working with the fsnameservers File

The SNFS `fsnameservers` file specifies machines serving as File System Name Server coordinator(s) to the `fsmpm` daemon. The File System Name Server coordinator is a critical component of the StorNext File System Services (FSS). A principal function of the coordinator is to manage failover voting in a high-availability configuration. Therefore, it is critical to select highly reliable systems as coordinators. Redundancy is provided by listing multiple machine entries in the `fsnameservers` file, one entry per line. The first machine listed is the primary coordinator and any subsequent machines listed serve as backup coordinators. To create redundancy, it is recommended that you list two machines. Typically, the selected systems are also configured for FSM services, but this is not a requirement.

If the `fsnameservers` file does not exist, then the file system operates as a local file system, requiring both a client and a server. The file system

will not communicate with any other StorNext File System product on the network, thus eliminating the sharing of the FSS over the SAN.

Caution: It is extremely important that all copies of `/usr/cvfs/config/fnameserver` file in a SAN be identical. A stale configuration on a system that is not in use can cause election problems if `fsmppm` processes are running with mismatched `fnameservers`.

It is also critical to verify that complete network connectivity exists between all systems running SNFS (client or server). This is important because all StorNext systems participate in the failover process.

After you type the IP addresses of two reliable machines on your network in the `fnameservers` file, copy the `fnameservers` file to every machine running SNFS. After changing the `fnameservers` file, always restart (stop and start) all file system services running on the SNFS.

Redundant NICs are supported in the `fnameservers` file; simply list both NIC addresses in the file.

Making or Unmaking a File System

Use the following procedures to make or re-make a file system.

Caution: This task destroys all existing data for the selected file system and creates a new configuration. Making or re-making a file system results in a complete loss of user data.

After remaking a managed file system, you must delete or recreate the relation points.

Conditions to Make or Re-Make a File System

The following are reasons to make or re-make a file system.

- Creating a new file system
- Removing a stripe group from the file system
- Removing a disk from a stripe group
- Changing a stripe group's stripe breadth


- Changing the sector count of a disk

Making a File System

Use this procedure to make a file system.

- 1 From the SNFS home page, choose **Make File Systems** from the **Admin** menu. The **Make File System** screen appears.

Figure 17 Make File System Screen



- 2 In the **File Systems** list, select a new or existing file system and click **Apply**. A confirmation screen warns you that all relation points on the system will be removed during this process, and prompts if you still want to make the selected file system.
- 3 Click **OK**. A second confirmation screen warns that all data will be lost on the selected file system.
- 4 Click **OK**.
- 5 After the status screen informs you that the task has completed successfully, click **Close**.


Starting and Stopping the File System

This task enables you to start or stop a file system. You need to stop a file system if you want to take it out of service to make configuration changes or perform other administrative tasks. You need to start a file system to place it back into service.

- 1 From the SNFS home page, choose **Start/Stop File System** from the **Admin** menu. The **Start or Stop File System** screen appears.

Note: Unselectable (greyed out) file systems in the **Start or Stop File System** screen require a metadata dump in order to be started. (This does not apply for SNFS-only systems.)

Figure 18 Start or Stop File System Screen



- 2 Do one of the following:
 - Select a file system from the **Inactive File Systems** list and click **Start**.
 - The **Start Stop File System Status** screen appears.
 - Select a file system from the **Active File Systems** list and click **Stop**.
 - The **Start Stop File System Status** screen appears.
- 3 After the status screen informs you that task has successfully completed, click **Close**.

Mounting or Unmounting a File System

This task enables you to mount or unmount a file system. You must unmount a file system if you want to take it out of service to make configuration changes or perform other administrative tasks. You must mount a file system to place it back into service.

Note: Before you unmount the file system you must stop the SNSM component as described in [Starting and Stopping StorNext Components](#) on page 57.

Note: The StorNext GUI performs NFS export/unexport on mount/dismount.

- 1 From the SNFS home page, and choose **Mount/Unmount** from the **Admin** menu. The **Mount or Unmount File System** screen appears.

Figure 19 Mount or Unmount File System Screen



- 2 Do one of the following:
 - Select a file system from the **Unmounted File Systems** list and click **Mount**. The **Mount File System Status** screen appears.
 - Select a file system from the **Mounted File Systems** list and click **Unmount**. The **Unmount File System Status** screen appears.

- 3 After the status screen informs you that the task has successfully completed, click **Close**.

Working With Disks

A disk:

- can be used as an individual element or as one of many disks in a Redundant Array of Inexpensive Disks (RAID).
- can contain metadata information, journaling, and/or data.
- can constitute a single stripe group itself, or can be one node of a multi-disk stripe group.
- can be a local hard disk located inside a server (if only a metadata disk), or most commonly, used in a RAID visible to all machines in the Storage Area Network (SAN) over FC.

The procedures in this section describe how to manage disks by adding, deleting, and defragmenting disks. Each disk is assigned to a disk type that specifies the number of sectors on the disk.

- [Adding Disks](#)
- [Deleting Disks](#)
- [Defragmenting a Disk](#)


Adding Disks

Use this procedure to add a disk to a selected file system.

Note: A disk must have a label before you can add it. For information about labeling a disk, see [Labeling a Device](#) on page 102.


- 1 If the file system is mounted, unmount the file system as described in [Mounting or Unmounting a File System](#) on page 112.
- 2 If the file system is started, stop the file system as described in [Starting and Stopping the File System](#) on page 111.
- 3 From the SNFS home page, choose **Disks** from the **Config** menu. The **Manage Disks** screen appears.

Figure 20 Manage Disks
Screen



- 4 From the **File Systems** drop-down menu, select the file system to which you want to add the disk, and then click **Add**. The **Add Disk** screen appears.

Figure 21 Add Disk
Screen



- 5 Select from the **Available Disks** list the disk you want to add to the file system. You can select multiple disks by pressing the CTRL key and clicking the disk name. Click **OK** to continue. The **Add Disk Status** screen appears.

- 6 After the status screen indicates that the disks have been added, click **Close**. The **Manage Disks** screen ([Figure 20](#) on page 114) appears, and the disks list includes the disks you just added.
- 7 Start the file system as described in [Making a File System](#) on page 110.
- 8 Mount the file system as described in [Mounting or Unmounting a File System](#) on page 112.

Caution: When you add a new disk or stripe group to your SAN, often an OS-dependent operation must be run to make the added device recognizable by a host. Some of these utilities can disrupt access to existing disks, causing access hangs or failures. To avoid this, stop all file system operations on the affected host *before* rescanning for the new device.

Deleting Disks

Use this procedure to delete a disk from a selected file system.

Caution: Deleting a disk used in a stripe group results in a complete loss of user data and requires re-making the file system. For information on re-making the file system, refer to [Making or Unmaking a File System](#) on page 109.

- 1 From the SNFS home page, choose **Disks** from the **Config** menu. The **Manage Disks** screen ([Figure 20](#)) appears.
- 2 Select from the **File Systems** drop-down menu the file system that contains the disk you want to delete.
- 3 Select from the **Disks** list the disk you want to delete.
- 4 Click **Delete**. A message asks you to confirm that you want to delete the disk.
- 5 Click **OK** to proceed with the deletion. The **Delete Disk** status screen appears.
- 6 After the status screen indicates that the disk has been deleted, click **Close**.

Defragmenting a Disk

This procedure describes how to defragment a disk using the `snfsdefrag` utility to relocate SNFS file data into a single, pre-allocated extent. Reducing the number of extents in a file improves system performance by minimizing disk head movement when I/O occurs. You can defragment a single file or multiple files, perform a recursive defragment on files in a directory, or defragment the entire file system.

1 Log onto a metadata controller as root. The metadata controller is the computer where the StorNext server (not the client) is installed.

2 Do one of the following:

- To defragment a single file or a list of files, type:

```
/usr/cvfs/bin/snfsdefrag <filename> [filename ....]
```

- To perform a recursive defragment on files in a directory, type:

```
/usr/cvfs/bin/snfsdefrag -r <directory_name>
```

If the `directory_name` value is equal to the mount point, then the entire StorNext file system is defragmented.

Note: For more information about using the `snfsdefrag` command, refer to the `snfsdefrag(1)` man pages.

Working With Stripe Groups

A stripe group is a logical disk volume in a file system that consists of one or more LUNs used to store metadata information, journaling information, and user data. Stripe groups are used to create logical volumes that can span multiple controllers on an array and even multiple arrays for even greater performance. Stripe groups are composed of LUNs of the same size and disk class based on your cost and performance requirements. (Fibre channel provides the highest performance and duty cycle.)

A stripe group contains definitions about read and write permissions, real time I/O constraints, a stripe breadth definition, multi-pathing methodology, and an affinity association. A file system can contain multiple stripe groups.

Stripe groups are bound together to create a StorNext file system. When data is written into the file system, two critical things happen.

First, data is separated from metadata. Metadata operations are typically small and random, and they require a lot of head movement on disks. Data however, tends to be written in large sequential patterns with less head movement. By separating data and metadata, thrashing is minimized and performance is maximized.

The second critical thing that can happen when data is written to the file system is file steering. Stripe groups can be different sizes and categories of disk, so you could have a file system with one stripe group of SATA disk intended for proxy files or temporary storage (i.e., less critical data on lower duty cycle, slower disk). You might have another larger fibre channel stripe group for storing raw content (i.e., high value data, on higher duty cycle, higher performance disk).

Getting data to specific stripe groups is accomplished using affinities, a mapping that ties a directory in the file system to a specific stripe group. When you write files to a directory, StorNext uses affinities to transparently write those files to the desired disk type. This means you don't have to have two or more separate file systems (e.g., one for critical data and another for temporary data). Instead, you have a single namespace for simplified management.

Working with a stripe group consists of these tasks:


- [Adding a Stripe Group](#)
- [Modifying a Stripe Group](#)
- [Deleting a Stripe Group](#)

Adding a Stripe Group

Use this procedure to add a stripe group.


- 1 From the SNFS home page, choose **Stripe Groups** from the **Config** menu. The **Configure Stripe Groups** screen appears.

Figure 22 Configure Stripe Group Screen



- 2 From the **File Systems** drop-down menu, select the file system to which you want to add a stripe group, and then click **Add**. The **Add Stripe Group** screen appears.

Figure 23 Add Stripe Group
Screen



3 Enter the fields on the **Add Stripe Group** screen.

- **Name:** The name of the new stripe group.
- **Breadth:** The breadth size is the number of file system blocks the system reads/writes before moving to the next disk in the stripe group. This value should be optimized for the major application that uses the disks. The selectable breadth values range from the size of the file system block size to a maximum of 8192KB. The default setting is 64KB.
- **Available Disks:** Disks available to the stripe group. To populate this list, you must first add disks to the file system by doing the following:

- a From the SNFS home page, choose **Disks** from the **Config** menu.
 - b Select the file system in which the disks reside, and then click **Add**.
 - c Add one or more disks.
- **Disks in Stripe Group:** Disks associated with the new stripe group.
 - **Available Affinities:** Affinities associated with existing stripe groups that, if selected, would also apply to this stripe group.
 - **Affinities in Stripe Group:** Affinities associated with the new stripe group.
 - **Metadata, Journal, and Exclusive:** Enable one or more of these options (check the boxes) to create a location for metadata and journaling on the stripe group.
 - To enable placing metadata and journaling but not user data on the new stripe group, select the **Metadata**, **Journal**, and **Exclusive** checkboxes. Do not associate an affinity with a metadata, journal, or an exclusive stripe group.
 - To enable placing only metadata on the new stripe group, select both the **Metadata** and **Exclusive** checkboxes.
 - To enable placing only journaling on the new stripe group, select both the **Journal** and **Exclusive** checkboxes
 - Journaling can be enabled on only one stripe group.
 - If the **Metadata** and/or **Journal** boxes are not checked on any of the stripe groups in this file system, default settings cause metadata and journaling to be placed on the first defined stripe group.
 - **Realtime IO/sec** (optional): The number of disk I/O operations per second that are available to real-time applications. This is an optional setting and can be left blank.
 - **Realtime MB/sec** (optional): The number of megabytes per second that are available to real-time applications. This is an optional setting and can be left blank.
 - **Non-realtime IO/sec** (optional): The number of disk I/O operations per second that are available to non-realtime

applications. If both Non-realtime IO/sec and Non-realtime MB/sec Fields are selected, the system uses the lesser of the two values. This is an optional setting and can be left blank.

- **Non-realtime MB/sec** (optional): The lower threshold (in MB per second) that is reserved for non-realtime applications. If both Non-realtime MB/sec and Non-realtime IO/sec Fields are selected, then the system uses the lesser of the two values. This is an optional setting and can be left blank.

- 4 Click **OK**.
- 5 After the status screen indicates that the new stripe group has been added, click **Close**.


Caution: When you add a new disk or stripe group to your SAN, often an OS-dependent operation must be run to make the added device recognizable by a host. Some of these utilities can disrupt access to existing disks, causing access hangs or failures. To avoid this, stop all file system operations on the affected host *before* rescanning for the new device.

Modifying a Stripe Group

Use this procedure to modify a stripe group configuration.

- 1 From the SNFS home page, choose **Stripe Groups** from the **Config** menu. The **Configure Stripe Groups** screen appears.
- 2 Select from the **File Systems** drop-down menu the file system whose stripe group you want to modify.
- 3 Select from the **Stripe Groups** list the stripe group you want to modify.
- 4 Click **Modify**. The **Modify Stripe Group** screen appears.

Figure 24 Modify Stripe Group
Screen



5 Type valid values for the stripe group.

- **Stripe Group Name:** The selected stripe group.
- **Stripe Breadth:** The breadth size is the number of file system blocks the system reads/writes before moving to the next disk in the stripe group. This value should be optimized for the major application that uses the disks. The selectable breadth values range from the size of the file system block size to a maximum of 8192KB. The default setting is 64KB.
- **Available Disks:** Disks available to the stripe group. To populate this list, you must first add disks to the file system by doing the following:
 - a From the SNFS home page, choose **Disks** from the **Config** menu.

- b Select the file system, and then click **Add**.
- c Add one or more disks.
- **Disks in Stripe Group:** Disks associated with the stripe group.
- **Available Affinities:** Affinities associated with existing stripe groups that, if selected, would also apply to this stripe group.
- **Affinities in Stripe Group:** Affinities associated with the stripe group.
- **Read-only:** Select this option to make the stripe group read only.
- **Status:** The status of the stripe group, either online or offline.
 - Select **Up** to put the stripe group online.
 - Select **Down** to take the stripe group offline. Data stored in the stripe group is unavailable when the status is **Down**.
- **Metadata, Journal, and Exclusive:** Enable one or more of these options (check the boxes) to create a location for metadata and journaling on the stripe group.
 - To enable placing metadata and journaling but not user data on the new stripe group, select the **Metadata**, **Journal**, and **Exclusive** checkboxes. Do not associate an affinity with a metadata, journal, or an exclusive stripe group.
 - To enable placing only metadata on the new stripe group, select both the **Metadata** and **Exclusive** checkboxes.
 - To enable placing only journaling on the new stripe group, select both the **Journal** and **Exclusive** checkboxes.

Note: Journaling can be enabled on only one stripe group.

If the **Metadata** and/or **Journal** checkboxes are not checked on any of the stripe groups in this file system, default settings cause metadata and journaling to be placed on the first defined stripe group.

If the file system was made, the metadata cannot be removed.

- **Realtime IO/sec** (optional): The number of disk I/O operations per second that are available to real-time applications. This is an optional setting and can be left blank.

- **Realtime MB/sec** (optional): The number of megabytes per second that are available to real-time applications. This is an optional setting and can be left blank.
- **Non-realtime IO/sec** (optional): The number of disk I/O operations per second that are available to non-realtime applications. This is an optional setting and can be left blank.
- **Non-realtime MB/sec** (optional): The lower threshold (in megabytes per second) that is reserved for non-realtime applications. This is an optional setting and can be left blank.

6 Click **OK**.

7 After the screen indicates that the stripe group has been modified, click **Close**.

Deleting a Stripe Group

Use this procedure to delete a stripe group.

Caution: Deleting a stripe group causes a complete loss of data and requires re-making the file system. Refer to [Making or Unmaking a File System](#) on page 109.

- 1 From the SNFS home page, choose **Stripe Groups** from the **Config** menu. The **Configure Stripe Group** screen ([Figure 22](#) on page 118) appears.
- 2 Select from the **File Systems** drop-down menu the file system that contains the stripe group you want to delete.
- 3 Select from the **Stripe Groups** list the stripe group you want to delete.
- 4 Click **Delete**. A confirmation screen prompts you to confirm that you want to delete the stripe group.
- 5 Click **OK**. The **Delete Stripe Group Status** screen appears.
- 6 After the status screen indicates that the stripe group has been deleted, click **Close**.

Working With Affinities

An affinity is a label assigned to one or more stripe groups. An affinity enables you to direct data to its associated stripe groups. A file system can have one or more affinities associated with it. These associations, defined in stripe groups, are made with the affinity key to a directory or file.

For example, if you create a directory association with Affinity aff1, all data written to that directory is written to StripeGroup StripeGroup2, and therefore only to disk2.

If you make an association with Affinity aff2 and a separate directory in the file system, all data is directed to StripeGroup StripeGroup3, which contains three disks: disk3, disk4, and disk5. All data is written to these disks when directed to the associated directory with Affinity aff2.

This section includes the following topics:


- [Adding an Affinity](#)
- [Adding an Affinity Through SNFS](#)
- [Setting the Affinity in a Directory](#)
- [Modifying an Affinity](#)
- [Deleting an Affinity](#)
- [File System Configuration Restrictions](#)

Adding an Affinity

Use the following procedure to add an affinity to a configured file system.


- 1 From the StorNext home page, select **Add Affinity** from the **Config** menu. The **Add Affinity Introduction** screen appears, showing a list of existing configured file systems to which you can add an affinity.

Figure 25 Add Affinity
Introduction Screen



2 Click **Next** to continue. The **Add Affinity** screen appears.


Figure 26 Add Affinity Screen



- 3 Select the file system to which you want to add the affinity. Enter a name for the affinity (up to eight characters long) that begins with a letter.

Click **Next** to continue. The **Select Directory** screen appears.


Figure 27 Select Directory
Screen



- 4 Specify the directory for the new affinity. You can enter a pathname directly in the field, or click **Browse** to select from a list of existing directories. If desired, you can also create a new directory after you click **Browse**.


Click **Next** to continue. The **Assign Affinity** screen appears.

Figure 28 Assign Affinity
Screen



- 5 Select from the displayed the list the stripe group to which you want to assign the new affinity. If none of the shown stripe groups are available, you must exit the procedure and create additional stripe groups before proceeding.
Click **Next** to continue. The **Exclusive Stripe Group** screen appears.

Figure 29 Exclusive Stripe Group
Group Screen




- 6 If desired, you can select a stripe group that will be used exclusively for the new affinity's files. To make a stripe group exclusive, select the desired stripe group from the left column (Not Exclusive) and click the upper arrow (>) to move the stripe group to the Exclusive list.

Conversely, you can change a previously configured stripe group from exclusive by selecting it from the Exclusive column and clicking the lower arrow (<) to move it to the Not Exclusive column.

Click **Next** to continue. The **Complete Add Affinity Task** screen appears.

Figure 30 Complete Add Affinity Task Screen




- 7 Click **Next** to continue. After the status screen informs you that the affinity was added successfully, click **Close**.

Adding an Affinity Through SNFS

Alternatively you can use the following procedure to add an affinity to an existing file system through SNFS.


- 1 If the file system is mounted, unmount the file system as described in [Mounting or Unmounting a File System](#) on page 112.
- 2 If the file system is started, stop the file system as described in [Starting and Stopping the File System](#) on page 111.
- 3 From the SNFS home page, choose **Affinities** from the **Config** menu. The **Add, Modify, or Delete Affinities** screen appears.

Figure 31 Add, Modify, or Delete Affinities Screen



- 4 From the **File Systems** drop-down menu, select the file system to which you want to add the affinity, and then click **Add**. The **Add Affinity** screen appears.

Figure 32 Add Affinity Screen



- 5 Select a Stripe Group with which to associate the affinity, and then type a name for the affinity in the **Affinity Name** Field. (Affinity

names cannot be longer than eight characters.) Click **OK** to continue.

The following stripe groups cannot be selected:

- An exclusive stripe group
- A metadata stripe group
- A journal stripe group
- A stripe group on a managed file system with a configured affinity

- 6 After the Status screen informs you that the affinity was successfully added, click **Close**. The **Add, Modify, or Delete Affinities** screen ([Figure 31](#) on page 132) appears, showing in the Affinities list the new affinity you just added.
- 7 Start the file system as described in [Making a File System](#) on page 110.
- 8 Mount the file system as described in [Mounting or Unmounting a File System](#) on page 112

Setting the Affinity in a Directory

When you use SNFS to add an affinity, you must set the affinity in a directory on the selected file system. (If you use the Add Affinity Wizard, this step is done automatically.)


- 1 From the SNFS home page, choose **Set Affinities** from the **Admin** menu. The **Set Affinity** screen appears.

Figure 33 Set Affinity Screen



- 2 Select from the **File Systems** list the file system on which you want to set the new affinity.
- 3 Select from the **Affinities** list the affinity to set on the selected file system.
- 4 Click **Apply** to continue. The **Set Affinity** screen appears.

Figure 34 Set Affinity Screen




- 5 On the **Set Affinity** screen, click **Browse** and select or create the directory on the file system to which you want to set the affinity. Click **OK**.
- 6 When the status screen informs you that the operation was completed successfully, click **OK**.

Modifying an Affinity

Use this procedure to modify an affinity on a configured file system.

- 1 If the file system is mounted, unmount the file system as described in [Mounting or Unmounting a File System](#) on page 112.
- 2 If the file system is started, stop the file system as described in [Starting and Stopping the File System](#) on page 111.
- 3 From the SNFS home page, choose **Affinities** from the **Config** menu. The **Add, Modify, or Delete Affinities** screen ([Figure 31](#) on page 132) appears.
- 4 Select from the **File Systems** drop-down menu the file system you want to modify.
- 5 Select from the **Affinities** list the affinity you want to modify.
- 6 Click **Modify**. The **Modify Affinity** screen appears.

Figure 35 Modify Affinity Screen



- 7 Select the stripe group(s) with which to associate the affinity.
- 8 Click **OK** to continue.
- 9 After the status screen informs you that the affinity has been modified, click **Close**. The **Add, Modify, or Delete Affinities** screen

([Figure 31](#) on page 132) appears, showing the association you just created.

- 10 Start the file system as described in [Making a File System](#) on page 110.
- 11 Mount the file system as described in [Mounting or Unmounting a File System](#) on page 112.

Deleting an Affinity

Use this procedure to delete an affinity on a configured file system.

Caution: Deleting affinities from a file system configuration after the file system has been in use for a while could result in abnormal behavior. Contact the Quantum Technical Assistance Center before deleting affinities from a file system configuration.

- 1 If the file system is mounted, unmount the file system as described in [Mounting or Unmounting a File System](#) on page 112.
- 2 If the file system is started, stop the file system as described in [Starting and Stopping the File System](#) on page 111.
- 3 From the SNFS home page, choose **Affinities** from the **Config** menu. The **Add, Modify, or Delete Affinities** screen ([Figure 31](#) on page 132) appears.
- 4 Select from the **File Systems** drop-down menu the file system that contains the affinity you want to delete.
- 5 Select from the **Affinity** list the affinity you want to delete.
- 6 Click **Delete**. A message asks you to confirm that you want to delete the affinity.
- 7 Click **Yes** to confirm the deletion. The **Delete Affinity Status** screen appears.
- 8 After the status screen indicates that the affinity has been deleted, click **Close**.
- 9 Start the file system as described in [Making a File System](#) on page 110.
- 10 Mount the file system as described in [Mounting or Unmounting a File System](#) on page 112.

File System Configuration Restrictions

When using affinities, the StorNext administrator must adhere to the file system configuration restrictions listed below to ensure that data management policies execute properly. These restrictions apply only to managed file systems, not to non-managed (SNFS-only) installations.

- Taking into account all managed file systems, you cannot use more than two unique affinity names. However, these two affinity names can be re-used across file systems.
- All data stripe groups must have exactly one affinity association, but several stripe groups can be associated to the same affinity. This restriction does not apply to exclusive metadata or journal stripe groups.
- At least one policy class must be created to use the configured affinities.

Caution: You must designate at least one data stripe group in each managed file system as non-exclusive. If you do not make this designation, writing to the area of the file system that is not associated with any of the affinities will result in an out-of-space error.


Deleting affinities from a file system configuration after the file system has been in use for a while could result in abnormal behavior. Contact the Quantum Technical Assistance Center before deleting affinities from a file system configuration.

Performing a Metadata Dump

This procedure replaces any existing metadata dump data for the selected file system, and should be run only if the metadata file has been lost or corrupted.

- 1 From the SNFS home page, choose **Metadata Dump** from the **Admin** menu. The **Metadata Dump of File System** screen appears.

Figure 36 Metadata Dump of File System Screen



- 2 Select the file system on which to perform the metadata dump, and then click **Apply**.
- 3 After the Status Screen informs you that the metadata dump was performed successfully, click **Close**.

Using the SNSM File System Functions

The StorNext Storage Manager's File menu contains file system functions that enable you to accomplish the following tasks:

- [Storing Files](#)
- [Changing a File Version](#)
- [Recovering a File](#)
- [Recovering a Directory](#)
- [Retrieving a File](#)
- [Retrieving a Directory](#)
- [Freeing Disk Blocks](#)

- [Moving Files to New Media](#)
- [Modifying a File's Attributes](#)


These functions are not available if you have only StorNext File System and not StorNext Storage Manager.

Storing Files

Use this function to expedite storing files to media rather than waiting for data to be migrated automatically.

- 1 From the SNSM home page, choose **Store** from the **File** menu. The **Store Files** screen appears.

Figure 37 Store Files Screen



- 2 Do one of the following:

- Enter in the **Enter Individual Filename** field the file you want to store to media. You must enter the file's complete pathname.
- Select from the **File List** one or more filenames to store to media. (To expedite filename selection, you can click the **Select All** button to select all files in the list, and then deselect the files you don't want to include. If you change your mind, you can deselect all selected files by clicking the **Deselect All** button.)

- Click **Browse** to display the **StorNext File Browser** window. On this window locate and select the files you want to store to media, and then click **OK**.
- 3 If desired, click the Properties button to change the default parameters that are applied during file storage. The **Optional Store Parameters** screen appears.

Figure 38 Optional Store Parameters Screen



- 4 On the **Optional Store Parameters** screen, enter values and then click **Apply**.
- **Number of Copies:** Specify the number of copies (1 - 4) to create for each file.
 - **Truncation:** Specify whether file truncation is applied immediately to each file, or by policy.
 - **Drive Pool Name:** Choose the default drive pool from a list of available drive pools. (Drive pools must be previously created.)
 - **Minimum File Size:** Specify the minimum size a file must be in order to qualify for storage.
 - **Select Media Type:** Specify whether files are stored according to the policy's media type, or on storage disk (SDISK or Dedup SDISK).

- 5 On the **Store Files** screen, click **Apply** to continue.
- 6 After the Status screen informs you that the operation was completed successfully, click **OK**.

Changing a File Version

Use this function to find alternate versions of a specified file.

- 1 From the SNSM home page, choose **Version** from the **File** menu. The **Change File Version** screen appears.

Figure 39 Change File Version Screen



- 2 Do one of the following:
 - Enter in the **Enter Individual Filename** field the file whose version you want to change. You must enter the file's complete pathname.
 - Click **Browse** to display the **StorNext File Browser** window. On this window locate and select the files you want to store to media, and then click **OK**.
- 3 Click the **Get Versions** button to display a list of available alternate versions for the file you specified.
- 4 Select the version you want to use, and then click **Apply**.

- 5 After the Status screen informs you that the operation was completed successfully, click **OK**.

Recovering a File

This function allows you to recover a deleted file. Undeleted files are recovered in a truncated state. To return the file back to disk, you must use the Retrieve File function as described in [Retrieving a File](#) on page 145.


- 1 From the SNSM home page, choose **Recover File** from the **File** menu. The **Recover Files** screen appears.

Figure 40 Recover Files Screen



- 2 Enter in the **Enter File Filter** field a file filter to help you locate the deleted file you want. The filter can include wildcard characters (the asterisk *) anywhere in the filter string.
- 3 Click **Browse** to locate and select files that have been deleted. The **StorNext Recoverable Files** screen shows a list of all deleted files that apply to the filter you entered. The date and time the file was deleted is also shown.

Figure 41 StorNext
Recoverable Files Screen



- 4 Select from the list the files you want to undelete. To expedite file selection, you can use the **Select All** or **Deselect All** buttons. If the file list spans more than one screen page, click **Next** to view the next page, or **Back** to return to the previous page. When you are finished selecting files for recovery, click **OK**.
- 5 On the **Recover Files** screen, verify that the files shown are the ones you want to recover. Click the **Select All** button to include all files.
- 6 When you are ready to recover the selected files, click **Apply**.
- 7 After the Status screen informs you that the operation was completed successfully, click **OK**.


Recovering a Directory

Use this function to recover a deleted directory. When you restore a directory, the deleted files and inodes from the directory and its sub-directories are recovered, but the files are in a truncated state. To retrieve the files you must use the Retrieve Directory function as described in [Retrieving a Directory](#) on page 146.

Note: When you use this function, you must first select the file system to which you want to restore the deleted directory. That file system must be started and mounted.

- 1 From the SNSM home page, choose **Recover Directory** from the **File** menu. The **Recover Directory** screen appears.

Figure 42 Recover Directory Screen



- 2 Select from the **Select File System** dropdown list the file system to which you want to restore the deleted directory.
- 3 Select from the **Select Directory** list the directory you want to recover.
- 4 If desired, click the **Details** button to view all files and sub-directories that will be restored when the selected directory is recovered.
- 5 Click **Apply** to recover the selected directory.
- 6 After the Status screen informs you that the operation was completed successfully, click **OK**.

Retrieving a File

This function allows you to retrieve an entire file or a portion of the file from media to disk. If you retrieve a partial file, you must give the file a new name to prevent overwriting the current version.

- 1 From the SNSM home page, choose **Retrieve File** from the **File** menu. The **Retrieve Files** screen appears.

Figure 43 Retrieve Files Screen



2 Do one of the following:

- Enter in the **Enter Individual Filename** field the file you want to retrieve. You must enter the complete pathname.
- Click **Browse** to display the **StorNext File Browser** window. On this window locate and select the files you want to store to media, and then click **OK**. On the **Retrieve Files** screen, verify that the files shown are the ones you want to retrieve.

3 If desired, click the **Properties** button to view the properties for the files you are retrieving.

4 Click **Apply** to retrieve the selected files.

5 After the Status screen informs you that the operation was completed successfully, click **OK**.

Retrieving a Directory

This function allows you to recursively retrieve (copy) a directory's files and sub-directories from media, and then place them back on disk.

1 From the SNSM home page, choose **Retrieve Directory** from the **File** menu. The **Retrieve Directory** screen appears.

Figure 44 Retrieve Directory
Screen



2 Do one of the following:

- Enter in the **Enter Directory** field the directory from which you want to recursively retrieve files and sub-directories. You must enter the complete pathname.
- Click **Browse** to display the **Directory Browser** window. On this window locate and select the directory you want, and then click **OK**.

3 Click **Apply** to retrieve the selected files.

4 After the Status screen informs you that the operation was completed successfully, click **OK**.

Freeing Disk Blocks

Use this function to free disk blocks by removing one or more files from disk. Before you can remove a file from disk, it must first be on media. (The file remains on media after you remove it from disk.)

1 From the SNSM home page, choose **Free Disk Blocks** from the **File** menu. The **Free Disk Blocks** screen appears.

Figure 45 Free Disk Blocks
Screen



2 Do one of the following:

- Enter in the **Enter Individual Filename** field the file you want to remove from disk. You must enter the file's complete pathname.
- Click **Browse** to display the **StorNext File Browser** window. On this window locate and select the files you want to remove from disk, and then click **OK**. On the **Free Disk Blocks** screen, verify that the files shown are the ones you want to remove.

3 Click **Apply** to remove the selected files from disk.

4 After the Status screen informs you that the operation was completed successfully, click **OK**.


Note: If the stub files feature is enabled and the file size is smaller than the stub file size you specified when creating the storage policy, the stub file remains on the disk. In this situation, you might receive a Failure status after you click **Apply**.

Moving Files to New Media

This function enables you to move files from one piece of media to another. When you use this function, files on the original media are deleted. (That is, this is not a copy function that leaves files on the source media and places a copy on the destination media.)

- 1 From the SNSM home page, choose **Move** from the **File** menu. The **Move Files to New Media** screen appears.

Figure 46 Move Files to New Media Screen



- 2 Do one of the following:
 - Enter in the **Enter Individual Filename** field the file you want to move. You must enter the file's complete pathname.
 - Click **Browse** to display the **StorNext File Browser** window. On this window locate and select the files you want to move, and then click **OK**. On the **Move Files to New Media** screen, verify that the files shown are the ones you want to move.
- 3 Specify the **Media ID** and **Media Type** for the destination media. If desired, select the **Move to Blank Media** option. (When you select this option, StorNext searches for blank media on which to move the selected files.)
- 4 Click **Apply** to move the selected files to new media.
- 5 After the Status screen informs you that the operation was completed successfully, click **OK**.

Modifying a File's Attributes

With this function you can change a file's attributes, including the associated policy class attributes and number of file copies to save on media during storage.

- 1 From the SNSM home page, choose **Attributes** from the **File** menu. The **Modify File Attributes** screen appears.

Figure 47 Modify File Attributes Screen

The screenshot shows the 'Modify File Attributes' web interface. At the top, there is a navigation bar with 'File', 'Media', 'Admin', 'Reports', and 'Help'. The main content area is titled 'Modify File Attributes' and includes a subtitle: 'Select files and modify the associated policy class attributes or change the number of file copies to save on media during storage.' Below this, there is a 'Select Individual File' section with a text input field and a 'Browse' button. A 'File List' section contains a table and buttons for 'Select All' and 'Deselect All'. The configuration section includes radio buttons for 'Number of Copies' (ONE, TWO, THREE, FOUR), 'Truncation Options' (After Store, During Cleanup, No Truncation), and 'Stub Size' (By Policy, User Set). At the bottom, there are 'Apply' and 'Reset' buttons. The browser window title is 'Modify File[s] Attributes' and the system tray shows 'spock Active' and 'Local intranet'.

- 2 Do one of the following:
 - Enter in the **Select Individual File** field the file whose attributes you want to modify. You must enter the file's complete pathname.
 - Click **Browse** to display the **StorNext File Browser** window. On this window locate and select the files whose attributes you want to modify, and then click **OK**.
- 3 Specify the number of copies to maintain for each selected file (One, Two, Three, or Four).
- 4 Indicate whether to truncate files after storing, during the file cleanup process, or not at all.
- 5 If the Stub File feature is enabled, indicate how StorNext determines the stub file size:

- **By Policy:** Use the stub file size entered when the storage policy was created
 - **User Set:** Change the stub file size by entering the new size (in kilobytes) at this field
- 6 Click **Apply** to modify attributes for the selected files.
 - 7 After the Status screen informs you that the operation was completed successfully, click **OK**.

Understanding Resource Allocation

StorNext provides two Resource Allocation tools that allow you to make changes to your file system: *File System Expansion*, and *Stripe Group Movement*.

About File System Expansion

StorNext's File System Expansion feature enables you to dynamically add LUNs to a selected file system without interrupting that file system's operation.

The only disruption that occurs during File System Expansion is a short pause of new metadata requests as StorNext updates its internal system and clients to be aware of the new overall capacity and physical disk resources that are used.

File System Expansion is often done in conjunction with the Stripe Group Movement feature. That is, you might want to add new stripe groups knowing you'll want to use those stripe groups for Stripe Group Movement.

StorNext provides a File Expansion Wizard to simplify the process. Quantum recommends using this wizard for File System Expansion, but you can also use the command line interface. (For information about using the CLI, see [Using the Resource Allocation Feature](#).)

Note: After expansion you must perform a metadata dump. The File Expansion Wizard includes an option that will do this for you automatically, but the process can take longer than if you do the metadump manually.

About Stripe Group Movement

Stripe Group Movement moves data files off one or more data stripe groups onto the remaining data stripe groups in a file system, which frees data LUNS so they can be decommissioned or reused. In a similar way, the metadata on a single LUN can be moved to a new LUN. StorNext provides a Movement Wizard to simplify these processes, which is launched when you select Move Stripe Group from the Admin menu on the SNFS home page.

During data stripe-group movement, you indicate one or more source stripe groups from which to move data. StorNext automatically moves all data from the source stripe groups to the remaining stripe groups in the file system. All other data stripe groups are targets, allowing an even distribution of data across remaining disk resources. During movement, the file system is online and read/write operations occur normally, but the source data stripe group(s) are in read-only mode (write disabled).

After all data has been removed from the source stripe group, you must mark the stripe group as “down,” which prevents new data from being written to the source stripe group. At this point the source LUNs are empty and read-only, so there will be no access to them.

Although stripe group and LUN configuration line items must never be deleted from a configuration once its corresponding file system has been created, marking a stripe group as down allows its LUNs to be relabeled and reused. The names of the LUNs in the configuration file can also be changed to free up the old names for reuse.

When moving metadata off one LUN onto a new LUN, the file system must be stopped. The Movement Wizard allows you to select one source and one destination LUN. On completion, the old LUN is relabeled with a suffix of **.old**, and the new LUN is relabeled with the old LUN’s name. The old LUN can then be relabeled and reused.

Quantum recommends using the Movement Wizard for Stripe Group Movement, but you can also use the command line interface. (For

information about using the CLI, see [Using the Resource Allocation Feature.](#))

Expansion and Movement Steps

Here are the steps required for expanding a file system and moving stripe groups:

- 1 Check the file system before you begin. (See [Checking the File System](#) on page 152.)
- 2 Expand the file system. (See [Performing File System Expansion](#) on page 154.)
- 3 Move data stripe groups or metadata/journal stripe groups. (See [Performing Stripe Group Movement](#) on page 161.)
- 4 Mark source stripe groups as down. (See [Modifying a Stripe Group](#) on page 121.)
- 5 Reboot all clients after the expansion.


Checking the File System

Before you perform either File System Expansion or Stripe Group Movement, you must first perform a check on the file system you plan to use for these features. This operation could take a significant amount of time depending on the size of the file system, so plan accordingly.

Use the following procedure to perform a file system check.

- 1 From the SNFS home page, choose **Check File System** from the **Admin** menu. The **Check File System** screen appears.

Figure 48 Check File System Screen



- 2 Select from the **File Systems** list the file system you want to check. Only file systems eligible for File System Expansion or Stripe Group Movement are shown in the list. Next to the file system name, the file system's current state is shown in parentheses: mounted or unmounted, and started or stopped.
- 3 If desired, select the **Check Read Only** option if you want to perform the file system check in read-only mode.

Note: If the file system you select is currently started and mounted, the check will be automatically performed in read-only mode. In read-only mode on a live file system (started and mounted,) you could receive false errors.


- 4 When you are ready to proceed with the check on the selected file system, click **Apply**. The status window informs you that the process was initiated. Close this window by clicking **Close**.

Checks in progress are shown as **In Progress** under the heading **Current Status of Check File System**. If the move does not appear with an In Progress status, click **Refresh** to update the list.)

- 5 Also listed under **Current Status of Check File System** are any previously run file system checks. To clear the list of previously completed checks (marked **Success** or **Failure**,) click **Clear**.

- 6 To view a status summary for any successfully run file system checks, click the **Success** or **Failure** indicator under the **Status** heading.

Figure 49 Check File System Status Screen



Performing File System Expansion

In practical terms, expanding an existing file system means allocating additional storage to that file system. You add additional storage by creating a new stripe group in the file system configuration file, and assigning new disks to the stripe group.

Here is an overview of the steps required to enable File System Expansion:

1. Define new disk LUNs and make them available across the fibre channel SAN to the metadata controller and StorNext clients.
2. Use the StorNext GUI to create a new stripe group from the LUNs.
3. Add the new stripe group to an existing file system.

After you perform these steps, StorNext automatically updates the available file system capacity and presents the expanded file system to the StorNext clients.

Note: During File System Expansion the file system is shut down. After expansion you must perform a new metadata dump runs after the file system is updated. For managed file systems, the File System Expansion wizard includes an option for you to run the metadump automatically, or you can do it separately from the wizard.

As a final part of the File System Expansion process, StorNext restarts the file system. This restart typically takes less than two minutes and is often completed in seconds, but during the restart all new read/write requests are paused while operations in progress continue normally.


Caution: When you add a new disk or stripe group to your SAN, often an OS-dependent operation must be run to make the added device recognizable by a host. Some of these utilities can disrupt access to existing disks, causing access hangs or failures. To avoid this, stop all file system operations on the affected host *before* rescanning for the new device.

Caution: When adding a disk to an existing stripe group while I/O is stopped on the clients, you should first unmount the file system before the MDC does the offline disk expansion. After the expansion is complete, refresh disks before mounting the file system again. If you do not unmount the file system before expansion, your clients could lock up, requiring a reboot.

Follow these steps when you are ready to use the File System Expansion Wizard:

- 1 From the SNFS home page, choose **Expand File System** from the **Admin** menu. The **File System Expansion - Introduction** screen appears, showing a list of file systems available for expansion.

Figure 50 File System
Expansion Introduction Screen



- 2 Select from the list the file system on which you want to perform File System Expansion. Click **Next** to continue.
- 3 A message reminds you that you must check your file system before beginning the expansion process.

Figure 51 Check File System
Reminder



Do one of the following:

- If you have not checked the selected file system, click **Cancel** and then run the file system check as described in [Checking the File System](#) on page 152.
- If you have already checked the selected file system, click **OK** to proceed. The **Data Stripe Group Selection** screen appears.


Figure 52 Data Stripe Group Selection Screen



- 4 The **Data Stripe Group Selection** screen shows all currently configured stripe groups. Specify the number of data stripe groups you want to add to the file system.
- 5 If you want StorNext to automatically perform a metadata dump after file system expansion is finished, select the **Perform metadump** option. (This option appears only for managed file systems.)

If you do not select this option, a message warns you that backups on the file system will fail until you perform a metadump. Click **OK** to proceed.


Figure 53 Metadump Warning



Note: If you proceed without selecting the **Perform metadump** option, you should manually perform a metadata dump after file system expansion is finished to make sure your backups do not fail. For information about performing a metadata dump, see [Performing a Metadata Dump](#) on page 137.

- 6 On the Data Stripe Group Selection screen, click **Next** to continue. The **New Data Stripe Group** screen appears.

Figure 54 New Data Stripe Group Screen



7 Enter values for the **New Stripe Group** screen.

- **Name field:** The name of the stripe group.
- **Select disks list:** The disks available to assign to the stripe group. You must select at least one disk for each stripe group.
- **Label Type:** If you plan to create LUNs larger than 2TB, you must specify the EFI label type when configuring a file system.

VTOC labels were used for all operating systems in previous StorNext and Xsan releases, and are still required for the SGI IRIX operating system, Solaris releases prior to Solaris 10 Update 2, and LUNs less than 1TB.

EFI labels are required if you plan to create LUNs that are larger than 2TB. (For Solaris, EFI labels are also required for LUNs with a raw capacity greater than 1TB.) EFI labels will not work with the IRIX operating system.


The correct value is automatically selected on the **New Stripe Groups** screen, so you can accept the default value unless you have a reason to change the label type.

For more information about 2TB LUN requirements, see the *StorNext Installation Guide*.

- **Label Help:** Click this link to display guidelines for determining whether to select VTOC or EFI labels. (See [Figure 7](#) on page 93.)
 - **Stripe breadth** drop-down menu: The stripe breadth for the file system. The stripe breadth is the number of kilobytes (KB) that is read from or written to each disk in the stripe. For a typical StorNext installation, 64KB is the recommended setting.
 - **Metadata or Data:** Specify whether you plan to use the stripe group for data or for metadata.
- 8 Click **Next** to continue. If you are entering more than one stripe group, your choices are saved and you are ready to make selections for the next stripe group. Repeat step 5 for each stripe group you are adding.

If you are adding only one stripe group, the **Complete File System Task** screen appears after you click **Next**.

Figure 55 Complete File System Task Screen



- 9 Click **Next** to complete the process.
- 10 After the status screen informs you that the expansion was completed successfully, click **Finish** to exit the wizard.

Performing Stripe Group Movement

The time it takes to complete the Stripe Group Movement process depends on the amount of data being moved between source and target stripe groups. When moving a data stripe group, the file system continues to run during the movement process. StorNext does not block any new read/write requests, or block updates to existing files on the source stripe group. All operations (including metadata operations) are handled normally, but no new writes are allowed to the source stripe group, which will be marked read-only.

Although the Movement Wizard focuses primarily on data movement, you can also move metadata stripe groups. You can move a metadata stripe group to a new stripe group of the same or greater capacity as the original metadata stripe group. However, during metadata stripe group movement the file system must be down, and no new read/write operations can occur until all metadata has been transferred and the file system is restarted. The exact amount of downtime is based on the disk size.

Note: When moving a metadata stripe group, the Movement Wizard shuts down the file system. Depending on the number and size of files in your system configuration, metadata movement could take a long time, so plan accordingly.

After data movement is complete, you must mark the source stripe group as “down.”

For movement purposes, StorNext treats metadata and journal stripe groups the same way, so it doesn't matter whether the stripe group you want to move is a metadata stripe group, a journal stripe group, or a combined metadata and journal stripe group. The only caveat is that stripe groups used for metadata/journal move cannot contain data. (StorNext has different mechanisms for moving data stripe groups versus metadata/journal stripe groups.)

Note: During Stripe Group Movement, affinities are preserved when files are moved from one stripe group to another. Consequently, there must be sufficient space on one or more destination stripe groups for any affinities on the source stripe group. (You must add any affinities from the source stripe group to the new stripe group.)

Launching the Movement Wizard

Use the following procedure to perform stripe group movement.

- 1 From the SNFS home page, choose **Move Stripe Group** from the **Admin** menu. The **Move Stripe Group** screen appears, showing the following information for completed moves and moves in progress:
 - **File System:** The name of the file system involved in the move
 - **Type:** Specifies whether the move is for a data stripe group or a metadata/journal stripe group

- **Lun/Stripe Group:** The name of the lun/stripe group involved in the move
- **Start Time:** The date and time the move was launched
- **Percentage Complete:** The current percentage complete status for the move
- **Status:** For current moves the status will be **Progress**. For previously completed moves the status will be either **Success** or **Fail**.

Figure 56 Move Stripe Group Screen

Move Stripe Group

- This page shows the current Stripe Group Move status.
- Click Move to start a new move.

To perform a complete move:


- 1) [Check File System](#)
- 2) [Expand File System](#)
- 3) [Move \(Metadata/Journal\) Data Stripe Groups](#)
- 4) [Mark Source Stripe Groups as Down \(Data Stripe Groups only\)](#)

File System	Type	Lun/Stripe Group	Start Time	Percentage Complete	Status
snfs1	Metadata	spodkmail -> beas01	Feb-27-2007 15:49:50	100%	Success
snfs1	Metadata	spodkmail -> spod005	Feb-27-2007 14:15:08	12%	Fail
snfs1	User Data	StripeGroup3	Feb-27-2007 15:02:29	100.00%	Success

Buttons:

- 2 To update the percentage complete for a move in progress, click **Refresh**.
To remove previously completed moves (with either a **Success** or **Fail** status,) click **Clear**.
- 3 To view details about either a previously completed move or a move in progress, click the link in the **Status** column. The status window for completed moves shows files skipped and files defragmented. The status for moves in progress shows percentage complete.
- 4 Click **Move** to launch the Movement Wizard. The **Move Stripe Group - Introduction** screen appears.

Figure 57 Move Stripe Group Introduction Screen




- 5 Select the file system for which you want to move stripe groups. Click **Next** to continue.
- 6 A message reminds you that you must check your file system before beginning the movement process.

Do one of the following:

- If you have not checked the selected file system, click **Cancel** and then run the file system check as described in [Checking the File System](#) on page 152.
- If you have already checked the selected file system, click **OK** to proceed. The **Move Options** screen appears.

Figure 58 Move Options Screen




- 7 Specify whether to move a data stripe group or a metadata or journal stripe group. (You cannot move a stripe group that contains both user data and metadata/journal data on the same LUN.) Click **Next** to continue.

If you are migrating **ONLY** a metadata/journal stripe group, go to [Step 9](#) on page 167.

If you are migrating a data stripe group, the **Data Stripe Group Move** screen appears.


Figure 59 Data Stripe Group
Move Screen



- 8 Select the source stripe group from which data will be moved. The stripe group you select will be marked as read only, and data from the stripe group will be distributed among one or more available stripe groups. (This feature does not allow you to specify the destination stripe groups.)


If there are no writable destination stripe groups available, you will receive a warning message and not be allowed to continue. If you receive this message you must first create destination stripe groups before you proceed with the move.

Figure 60 No Writable Stripe
Groups Warning



Similarly, if StorNext determines that the available destination stripe groups are not large enough to accommodate the data on the source stripe group, you will receive a warning message and not be allowed to continue using the selected source stripe group.


Figure 61 Destination Stripe Group Too Small Warning



Click **Next** to continue. If you are migrating only data stripe groups, the **Complete Stripe Group Move** screen appears. Go to [Step 12](#) on page 169.


- 9 When you move a metadata/journal stripe group, a message informs you that the file system must be stopped during the move.

Figure 62 Metadata Stripe Group Message



- 10 Click **OK** to continue. The **Metadata/Journal Move** screen appears.

Figure 63 Metadata/Journal
Move Screen




- 11 Select from the displayed lists of available LUNs the source (from) LUN and the destination (to) LUN. Click **Insert** to continue. The LUN pair are displayed beneath the **SOURCE>>DESTINATION** heading.

If you made a mistake during selection, you can clear the LUN pair by clicking **Remove**, and then repeat the selection.

When you are ready to continue with the move, click **Next**. The **Complete Stripe Group Move** screen appears.


Figure 64 Complete Stripe Group Move Screen



Note: The figure shows the **Complete Stripe Group Move** screen for a metadata stripe group move. The screen for a metadata/journal stripe group move looks similar.

- 12 On the **Complete Stripe Group Move** screen, verify the displayed information about the move and then click **Next** to continue. The **Process Initiated** status screen appears.

Figure 65 Process Initiated
Status Screen



- 13 After the status screen informs you that the move was initiated successfully, click **Finish** to exit the wizard.


Completing Metadata Stripe Group Movement

After the move begins, follow these steps to complete the operation.

- 1 Check move status by clicking the move's **Status** link on the **Move Stripe Group** screen ([Figure 56](#) on page 163).

The status screen for a move in progress shows the process ID, current files, and number of files copied. This screen looks similar to this:

Figure 66 Stripe Group Move Status Screen



- 2 When the move status is **Success**, mark the source stripe group as “down” and then restart the file system.
- 3 If the status is **Failed**, click the link to determine the cause of failure. If desired, retry the move.

Reusing LUNs in a Data Stripe Group After a Move


If you want to reuse LUNs in the data stripe group used in a successful move, you must follow these steps:

- 1 Mark the stripe group as “Down” as described in [Modifying a Stripe Group](#) on page 121.
- 2 From the command line interface, rename the LUNs in the downed stripe group. (You cannot accomplish this step through the GUI.)

Reusing LUNS in a Metadata Stripe Group After a Move

LUNs in the metadata stripe group used in a successful move can be reused immediately after the move. There is no need to rename LUNs, because during the move process StorNext automatically appends “.old” to the source’s LUN label name.

Although no additional steps are required, you might want to manually rename LUNs from either the StorNext GUI or the command line interface to provide better identification.



Chapter 7

Managing Libraries

StorNext configures, allocates, and manages all components associated with libraries. Resources that can be used for libraries include a wide variety of storage drives and media. Once configured, StorNext performs automated and manual operations to ensure the libraries resources are operating at efficient levels and in stable states. StorNext also offers flexibility when maintaining or replacing library components by minimizing system downtime.

This chapter describes the following tasks that help you manage your library:

- [Adding a Library](#)
- [Modifying a Library](#)
- [Deleting a Library](#)
- [Rescanning a Library](#)
- [Auditing a Library](#)
- [Changing the Library State](#)

Adding a Library

Use the procedure in this section to add libraries to StorNext. StorNext's Storage Manager (SNSM) component supports three library types:

- **SCSI:** A SCSI or fibre channel-attached library.
- **Network (ACSL or DAS):** A network-attached library. There are two types of network-attached library: **ACSL** or **DAS**.
- **Vault:** A library that stores media moved from a robotic library. A vault library is used to only store media, and cannot be reconfigured after it has been designated as a vault.


When adding a library to StorNext, follow the instructions specific to your library type.

Starting the Add Library Wizard

Use this procedure to launch the Add Library wizard.


- 1 From the StorNext home page, choose **Add Library** from the **Config** menu. The **Library Introduction** screen appears, showing all currently configured libraries.

Figure 1 Library Introduction Screen



2 Click **Next**. The **Library Type** screen appears.

Figure 2 Library Type Screen



- 3 Select the type of library you have: SCSI, Network, or Vault. (If you select Network, choose **ACSL** or **DAS** from the drop-down list.)
- 4 Click **Next** to continue.
- 5 Follow the procedures for your library type
 - [Adding a SCSI Library](#) on page 14
 - [Adding an ACSLS Network Library](#) on page 18
 - [Adding a DAS Network Library](#) on page 20
 - [Adding a Vault Library](#) on page 24

Adding a SCSI Library

After you specify SCSI on the **Library Type** screen and then click **Next**, the **Library Name** screen appears.


Figure 3 Library Name Screen



(If you have no SCSI libraries configured, a message informs you that no SCSI devices were detected. If you receive this message, configure the library and devices before proceeding.)


- 1 On the **Library Name** screen, accept the default library name or type a name, and then click **Next**. The **Media Types** screen appears.

Figure 4 Media Types Screen




- 2 Select a media type from the list, and then click **Next**. The **SCSI Device** screen appears.

Figure 5 SCSI Device Screen



- 3 Select a SCSI device from the list and click **Next**. The **Complete Add Library Task** screen appears.

Figure 6 Complete Add Library Task Screen



- 4 Review your selections. When finished, click **Next** to complete the task or **Back** to make changes.
- 5 After the status screen informs you that the library was successfully added, click **Close**.

Adding an ACSLS Network Library

After you select Network on the **Library Type** screen and then choose **ACSL**S from the drop-down list, the **Library Name** screen appears.


Figure 7 ACSLS Library Name Screen

- 1 Enter the fields on the **Library Name** screen.
 - **Library Name:** The name of the library. This can be any name you choose.
 - **Host Name:** The actual host name or IP address of the ACSLS server.
- 2 Click **Next** to continue. The **Media Types** screen ([Figure 117](#) on page 16) appears.
- 3 Select a media type from the list and click **Next**. The **Complete Add Library Task** screen ([Figure 119](#) on page 18) appears.
- 4 Review your selections. Click **Next** to complete the task, or **Back** to make changes.
- 5 After the status screen notifies you that the library was successfully added, click **Close**.

Adding a DAS Network Library

Figure 8 DAS Configuration Screen

After you select Network on the **Library Introduction** screen and then choose **DAS** from the drop-down list, the **DAS Configuration** screen appears.



1 Do one of the following:

- Select the **Failover** option, or both the **Failover** and **Dual Aisle** options, click **Next**, and proceed to [Adding a DAS Network Library with Failover](#) on page 21.
- Select only the **Dual Aisle** option, click **Next**, and proceed to [Adding a DAS Network Library with or without Dual Aisle Configuration](#) on page 22.

Note: For more information about failover or dual aisle configurations, contact the Quantum Technical Assistance Center. Refer to [Customer Assistance](#) on page 11 for contact information.

Adding a DAS Network Library with Failover

- 1 On the **Library Name** screen, enter valid values and click **Next**.

Figure 9 DAS Library Name Screen

The screenshot shows a web browser window titled "Add Library - Microsoft Internet Explorer". The page content is titled "Library Name". On the left side, there is a vertical banner with the Quantum logo at the top, a line drawing of a Scalar i2000 server tower in the middle, and the StorNext logo at the bottom. The main content area contains the following text and input fields:


- Library Name**
Enter a name for the new library. The name can be any combination of letters and numbers, but it cannot contain spaces.
- Enter the hostname or IP address of the DAS server.
- Enter the client name of this host that has been configured on the library's network.
- Enter the standby server host name that has been configured on the library's network.
- Enter the standby server client name that has been configured on the library's network.

At the bottom of the form, there are three buttons: "Back" (with a left arrow), "Next" (with a right arrow), and "Cancel" (with a red X).

- **Library Name:** The name of the library. This can be any name you choose
- **DAS Server Name:** The name of the DAS server
- **DAS Client Name:** The name of the DAS client configured on the DAS server
- **Standby Server Host Name:** Standby server host name configured on the network
- **Standby Server Client Name:** Standby server client name configured on the network

The **Media Types** screen appears.

Figure 10 DAS Media Types
Screen




- 2 Use the drop-down lists to map the mail boxes (EIF ports) to specific media, and then click **Next**. The **Complete Add Library Task** screen ([Figure 119](#) on page 18) appears.
- 3 Review your selections. Click **Next** to complete the task, or **Back** to make changes.
- 4 After a status screen informs you that the library has been successfully added, click **Close**.

Adding a DAS Network Library with or without Dual Aisle Configuration

- 1 On the **Library Name** screen, enter valid values and click **Next**.

Figure 11 DAS 2 Library Name Screen



- **Library Name:** The name of the library. This can be any name you choose.
- **Host Name:** The actual host name or IP address of the DAS server.
- **Client Name:** The name of the client for the current configuration. The client name is queried by StorNext and automatically displayed.


The DAS **Media Types** screen ([Figure 123](#) on page 22) appears.

- 2 Use the drop-down lists to map the mail boxes (EIF ports) to specific media, and then click **Next**. The **Complete Add Library Task** screen ([Figure 119](#) on page 18) appears.
- 3 Review your selections. Click **Next** to complete the task, or **Back** to make changes.
- 4 After a status screen informs you that the library was successfully added, click **Close**.

Adding a Vault Library

After you select Vault on the **Library Type** screen, the **Library Name** screen appears.

Figure 12 Vault Library Name Screen




- 1 Type a name for the library (or accept the displayed default name), and then click **Next**. The **Complete Add Library Task** screen ([Figure 119](#) on page 18) appears.
- 2 Review your selections. Click **Next** to complete the task, or **Back** to make changes.
- 3 After a status screen informs you that the library was successfully added, click **Close**.

Modifying a Library

This task describes how to modify a library by changing its media type.

- 1 From the SNSM home page, choose **Library > Config Library** from the **Admin** menu. The **Configure Library** screen appears.


Figure 13 Configure Library Screen



- 2 In the **Select Library** list, select the library you want to modify, and then click **Modify**. The **Modify Library** screen appropriate to the selected library type appears, showing information about the library and its associated device.

The following illustration shows the **Modify SCSI Library** screen. The screen for modifying a DAS or ACSLS library looks similar.

Figure 14 Modify SCSI Library
Screen



- 3 According to your needs, move media types from the **Unconfigured** media types list to the **Configured** media types list, or vice versa.

To move media, select the desired media type and then click the right or left arrow to move the media to the opposite list. (Click the right arrow to move the selected media type from the **Configured** list to the **Unconfigured** list, or the left arrow to move the selected media type from the **Unconfigured** list to the **Configured** list.)

- 4 When you are finished moving media to the appropriate list, click **Apply**. The **Modify Library Status** screen appears.
- 5 After the status screen indicates that the library has been modified, click **Close**.

Deleting a Library

Before you delete a library you **must** do the following:

- Remove ALL media from the library as described in [Removing Media](#) on page 20.

- Delete ALL drives associated with the library as described in [Deleting a Tape Drive](#) on page 19.
- 1 From the SNSM home page, choose **Library > Config Library** from the **Admin** menu. The **Configure Library** screen ([Figure 126](#) on page 25) appears.
- 2 Select from the **Select Library** list the library you want to delete, and then click **Delete**. A message window prompts you to confirm that you want to delete the library.
- 3 Click **OK** to proceed. The **Delete Library Status** screen appears.
- 4 After the status screen indicates that the library has been deleted, click **Close**.

Rescanning a Library

Whenever you add drive slots or media slots to your library, you should use the Rescan function to ensure that StorNext recognizes the new devices and can fully utilize them. (This function does not apply to vaults.)

Note: The rescan process disrupts any current activity because it places existing devices temporarily offline. Be sure to use the Rescan function after hours or during times of system inactivity when you are certain no StorNext operations will be disrupted.

- 1 From the SNSM home page, choose **Library > Config Library** from the **Admin** menu. The **Configure Library** screen ([Figure 126](#) on page 25) appears.
- 2 Click **Rescan** to begin the rescan process.

Auditing a Library

Use this procedure to audit a library. An audit is a physical check of each library component to verify its integrity and make sure the database and library are synchronized. Quantum recommends that you audit the library after each restore.

- 1 From the SNSM home page, choose **Library > Audit Library** from the **Admin** menu. The **Audit Library** screen appears.

Figure 15 Audit Library Screen



- 2 Select from the **Select Library** list the library you want to audit.
- 3 If desired, select the **Remap Audit** option to synchronize the StorNext and SNSM databases with the library databases.

Note: Quantum recommends that you select the **Remap Audit** option. If you are using an AML library, the audit could take hours or days to complete.

Audits should be run when there are no other processes running on the library.

- 4 Click **Apply**. The library audit launches, and the **Library Audit** status screen appears.
- 5 After the status screen informs you that the library audit was successful, click **Close**.

Changing the Library State

This function allows you to change a library's *logical* state to online or offline. *Logical* state means the library could be online, but StorNext treats it as offline.

To physically take a SCSI library offline, use the library's front panel. To take a network library physically offline, use the library's controller software.


- 1 From the SNSM home page, choose **Library > Library State** from the **Admin** menu. The **Change Library State** screen appears.

Figure 16 Change Library State Screen



- 2 Select the library whose logical state you want to change. (The library's current state is shown.)

- 3 Select the new state for the library you selected.
- 4 Click **Apply** to change the library's state.
- 5 When the Status screen informs you that the library state was successfully changed, click **OK**.
- 6 If desired, repeat steps 2 - 5 to change the state for additional libraries.



Chapter 8

Managing Drives and Disks

Tape drives provide I/O for a StorNext media library. For a detailed list of supported media for the current StorNext release, refer to the *StorNext Release Notes*.

Managing tape drives and disks consists of the following tasks:

- [Working with Tape Drives](#)
- [Working with Drive Pools](#)
- [Managing Disk Space](#)
- [Changing Watermark Parameters](#)

Working with Tape Drives

Working with tape drives involves the following tasks:

- [Adding a Tape Drive](#)
- [Modifying a Tape Drive](#)
- [Deleting a Tape Drive](#)
- [Changing a Drive State](#)
- [Cleaning a Tape Drive](#)


Adding a Tape Drive

Use this procedure to add tape drives to your libraries. You can add any number of connected tape drives to the StorNext system.

- 1 From the StorNext home page, choose **Add Tape Drive** from the **Config** menu. (Alternatively, from the SNSM home page, you can choose **Drive > Config** from the **Admin** menu, and then click **Add**.)


The **Tape Drive Introduction** screen appears, listing the number of currently configured tape drives and the hardware devices available for configuration.

Figure 1 Tape Drive Introduction Screen



- 2 Click **Next** to add a tape drive. The **Associated Library** screen appears.

Figure 2 Associated Library Screen




- **Configured Libraries** list: Select the configured library with which to associate the tape drives you are adding.
- **Fibre-Channel Attached Drives** checkbox: Check this box if you have fibre channel-attached tape drives. If you check this box, the **Match Devices with Slots** screen appears. Go to [Matching Devices with Slots](#) on page 15.

Note: For SCSI direct-attached tapes that are not fibre channel, it is not necessary to match the device with the correct slot because slot-to-drive matching is automatically performed. If StorNext cannot perform slot matching, the **Match Devices with Slots** screen appears.


- 3 After you select a configured library on the **Associated Library** screen, click **Next**. The **Hardware Devices** screen appears.

Figure 3 Hardware Devices
Screen



- **Hardware Devices** list: Select the hardware devices you want to add as drives. Drive names are automatically generated in this format: <library_dr1>, <library_dr2>
 - **Enable Compression** checkbox: Check this box to allow data compression on the tape drives you are adding.
- 4 Click **Next** to continue. The **Complete Add Drive Task** screen appears.

Figure 4 Complete Add Drive Task Screen




- 5 Review your selections. Click **Next** to complete the task or **Back** to make changes.
- 6 When the Status screen informs you that the tape drive was successfully added, click **Next**. The **Tape Drive Introduction** screen ([Figure 130](#) on page 12) appears.
- 7 Do one of the following:
 - Add more tape drives. Click **Next** to repeat the Adding a Tape Drive procedure ([Step 1](#) on page 12).
 - When the Status Screen displays Success, click **Done**.

Matching Devices with Slots

When you select the **Fibre-Channel Attached Drives** checkbox on the **Associated Library** screen, the **Match Devices with Slots** screen


appears. The system-specific information on this screen varies from configuration to configuration.

Figure 5 Match Devices with Slots Screen



- 1 If you need to see existing drive and slot mappings, click **Show Mapping Help** to display a list of current device mappings. Click **Close** when you are finished viewing the information on the **Tape Drive Mapping Help** screen.

Figure 6 Tape Drive Mapping Help Screen



Note: In some cases StorNext may not be able to determine the device to slot mapping. Limitations include ACSLS and DAS libraries.

- 2 On the **Match Devices with Slots** screen, select from the **Device** list the device you want to configure.
- 3 Select from the **Slot** list the slot with which to match the device.
- 4 Click **Insert** to add the device/slot combination to the Type>>Device>>Slot list.
- 5 Repeat steps 2 through 4 for each device and slot combination.

Note: If you want to enable compression on some devices but not others in the Type>>Device>>Slot list, you must complete this procedure twice: once to enable compression on selected devices, and a second time to specify devices without compression.

Modifying a Tape Drive

Use this procedure to modify configured tape drives.


- 1 From the SNSM home page, choose **Drive > Config** from the **Admin** menu. The **Configure Drives** screen appears.

Figure 7 Configure Drives Screen



- 2 Select a drive from the **Configured Drive List**, and then click **Modify**. The **Modify Drive** screen appears.

Figure 8 Modify Drive Screen



- 3 Modify any of the following information:
 - **Drive Name:** The name of the modified drive.
 - **Dismount Delay:** Change the dismount delay time (in seconds).


- **Select New Drive Path:** The drop-down menu displays available drive paths to which the drive can be changed.
 - **Compression Enabled:** Choose Yes to enable, or No to disable compression.
- 4 After making drive modifications, click **Apply**.
 - 5 After the Status Screen informs you that your modifications were made successfully, click **Close**.

Deleting a Tape Drive

Use this procedure to delete a configured tape drive.

- 1 From the SNSM home page, choose **Drive > Config** from the **Admin** menu. The **Configure Drives** screen appears.
- 2 From the **Configured Drive List**, select the drive you want to delete, and then click **Delete**. A message prompts you to confirm that you want to delete the drive. There is no undo function that will undelete the drive if you change your mind, so be absolutely certain you want to delete the drive before you continue.

Figure 9 Delete Warning Window




- 3 Click **OK** to close the message window and proceed.
- 4 After the Status Screen informs you that the drive has been successfully deleted, click **Close**.

Changing a Drive State

The drive state function allows you to change the logical (database) state to online or offline.

- 1 From the SNSM home page, choose **Drive > Drive State** from the **Admin** menu. The **Change Drive State** screen appears, showing all configured drives and the current state for each drive.

Figure 10 Change Drive State
Screen



- 2 Select the drive whose state you want to change.
- 3 Select the state (Online or Offline) to assign to the selected drive, and then click **Apply**.
- 4 After the Status Screen informs you that the drive's status was successfully changed, click **Close**. The **Change Drive State** screen shows the changed state for the drive you selected.
- 5 If desired, repeat steps 2 - 4 to change the state for additional tape drives.

Cleaning a Tape Drive

Use this procedure to manually clean a tape drive.

Note: To perform this procedure, you must have at least one cleaning tape specifically labeled for cleaning. For example, **CLN_XXX**.

- 1 From the SNSM home page, choose **Drive > Clean Drive** from the **Admin** menu. The **Clean Drive** screen appears.

Figure 11 Clean Drive Screen



- 2 Select from the **Select Drive List** the drive you want to clean, and then click **Apply**. The **Clean Drive** status screen appears.
- 3 After the Status screen notifies you that the drive has been cleaned successfully, click **Close**.

Working with Drive Pools

Drive pools are groups of tape drives allocated for various administrator-defined storage tasks. Drive pools enable you to delimit storage processes based on data type, performance, security, location, or all of these variables. Drive pools can reside in a single tape library or span multiple tape libraries.

Working with drive pools involves the following tasks:

- [Adding a Drive Pool](#)
- [Modifying a Drive Pool](#)
- [Deleting a Drive Pool](#)


Adding a Drive Pool

Use this procedure to add a drive pool.

Note: This procedure requires restarting the Storage Manager component.


- 1 From the SNSM home page, choose **Config Drive Pool** from the **Admin** menu. The **Configure Drive Pools** screen appears.

Figure 12 Configure Drive Pools Screen



- 2 Click **Add** to add a new drive pool. The **Add New Drive Pool** screen appears.

Figure 13 Add New Drive Pool Screen



- 3 Enter a name for the new drive pool in the **Drive Pool Name** field.
- 4 Select from the **Select Drive IDs** list the drives you want to include in the new drive pool, and then click **Apply**. A message window warns you that continuing with the task restarts the Storage Manager.

Figure 14 Warning Message Window



- 5 Click **OK** to close the message window and continue.
- 6 After the Status screen informs you that the drive pool was successfully added, click **Close**.


Modifying a Drive Pool

Use this procedure to modify a drive pool.

Note: This procedure requires restarting the Storage Manager component.

- 1 From the SNSM home page, choose **Config Drive Pool** from the **Admin** menu. The **Configure Drive Pools** screen (Figure 141 on page 22) appears.
- 2 Select from the **Select Drive Pool** list the existing drive pool you want to change, and then click **Modify**. The **Modify Drive Pool** screen appears.

Figure 15 Modify Drive Pool Screen



- 3 Select from the **Associated Drives** list the drive you want to move to the Available Drives list, and then click **Remove Left**. The drive is removed from the **Associated Drives** list and appears in the **Available Drives** list.
- 4 If desired, you can also move drives from the Available Drives list to the Associated Drives list. To do so, select from the **Available Drives** list the drive you want to move to the **Associated Drives** list, and then click **Insert Right**. The drive is removed from the **Available Drives** list and appears in the **Associated Drives** list.

- 5 Click **Apply**. A message window warns you that continuing with the task restarts the Storage Manager.

Figure 16 Restart Message Window



- 6 Click **OK** to close the message window and proceed.
- 7 After the status screen informs you that your drive pool modifications were successful, click **Close**.


Deleting a Drive Pool

Use this procedure to delete a drive pool. You cannot delete a drive pool if there are drives currently associated with it. To disassociate drives from the drive pool, see [Modifying a Drive Pool](#) on page 23.

Caution: At least one drive pool must be configured at all times. Do not delete the default drive pool.

- 1 From the SNSM home page, choose **Config Drive Pool** from the **Admin** menu. The **Configure Drive Pool** screen ([Figure 141](#) on page 22) appears.
- 2 Select from the **Select Drive Pool** list the drive pool you want to delete, and then click **Delete**. A message window asks you to confirm that you want to delete the drive pool. There is no undo function that will undelete the drive pool if you change your mind, so be absolutely certain you want to delete the drive pool before continuing.

Figure 17 Delete Warning Message




- 3 Click **OK** to close the message window and continue.
- 4 After the Status screen informs you that the drive pool was successfully deleted, click **Close**.

Managing Disk Space

This function allows you to apply a storage policy or truncation policy to a file system, and to adjust other parameters that help you manage disk space.

- 1 From the SNSM home page, choose **Disk Space** from the **Admin** menu. The **Manage Disk Space** screen appears.

Figure 18 Manage Disk Space Screen



- 2 Select from the **File Systems** list the file system that contains the disks whose space you want to manage.
 - 3 Enter the following fields:
 - **Storage Policy:** Choose this option to apply a storage policy to the file system.
- OR

- **Truncation Policy:** Choose this option to apply a truncation policy to the file system.
 - **Execute at Highest Priority:** Choose this option to make your storage or truncation policy execute at the highest priority.
 - **Stop at Occupied Disk Space:** Specify the disk-full percentage at which the storage or truncation policy is no longer applied.
 - **MINTIME Before File Migration:** If you chose the Storage Policy option, enter the minimum amount of time (in minutes) before a file is migrated.
 - **MINTIME Before File Truncation:** If you chose the Truncation Policy option, enter the minimum amount of time (in days) before a file is truncated.
 - **Minimum File Size to Truncate:** If you chose the Truncation Policy option, enter the minimum size (in bytes) a file must be before it is eligible for truncation.
- 4 Click **Apply** to continue.
 - 5 After the Status screen informs you that the operation was performed successfully, click **OK**.

Changing Watermark Parameters


Watermarks help you determine disk space thresholds for your file system. These thresholds determine the point at which StorNext applies or stops applying a truncation policy. This function also allows you to specify a minimum usage percentage, which is the percentage of occupied disk space the nightly truncation policy attempts to achieve.

- 1 From the SNSM home page, choose **Watermark Parameters** from the **Admin** menu. The **Change Watermark Parameters** screen appears.

Figure 19 Change Watermark Parameters Screen



- 2 Select from the **Managed File System(s)** dropdown list the file system for which you want to set watermark parameters.
- 3 Enter the following fields:
 - **Low Watermark %**: Enter the percentage of occupied disk space a file system must reach before StorNext stops applying the truncation policy.
 - **High Watermark %**: Enter the percentage of occupied disk space a file system must reach before StorNext applies a truncation policy.
 - **Min Use %**: Enter the target percentage of occupied disk space the nightly truncation policy should attempt to achieve.
- 4 Click **Apply** to continue.
- 5 After the Status screen informs you that the operation was performed successfully, click **OK**.



Chapter 9

Managing Media

In StorNext, data is usually stored on tapes in a media library. For a detailed list of supported media, refer to the *StorNext Release Notes*.

Managing tape media consists of the following tasks:

- [Adding Media to a Configured Library](#)
- [Removing and Moving Media](#)
- [Using the SNSM Media Functions](#)

Caution: If you are sharing a library between applications, refer to the *StorNext Installation Guide*.


Adding Media to a Configured Library

The following task describes how to add media to a configured library.

Caution: Before continuing with this procedure, make sure there is no media in your tape drives.


- 1 From the StorNext home page, choose **Add Media** from the **Config** menu. The **Add Media - Introduction** screen shows a list of current libraries, along with the number of media each library can hold.

Figure 1 Add Media -
Introduction Screen




- 2 Click **Next** to add media to a configured library. The **Associated Library** screen appears.

Figure 2 Associated Library
Screen



- 3 On the **Associated Library** screen, select a configured library to which you want to add media, and then click **Next**. The second **Associated Library** screen appears.


Figure 3 Associated Library
Screen 2



If you selected a media vault on the first **Associated Library** screen, go to [Step 1](#) on page 16.

- 4 On the second **Associated Library** screen, do one of the following:
 - Specify **Mailbox** and click **Next**.
 - Specify **Bulk Load** and click **Next**. Go to [Step 6](#) on page 15
- 5 When you specify **Mailbox** on the **Associated Library** screen, the **Select Mailbox** screen appears.


Figure 4 Select Mailbox Screen



Select a mailbox and click **Next**.

- 6 When the **Complete Add Media Task** screen appears, review your selections.

Figure 5 Complete Add Media Task Screen



Click **Next** to complete the task or click **Back** to make changes.


- 7 After a status screen informs you that media has been successfully added, click **Finish**.

Adding a Vault

Use this procedure to add a vault to the library.


- 1 When you select a media vault on the **Associated Library** screen, the **Select Media Type** screen appears.

Figure 6 Select Media Type
Screen




- 2 On the **Select Media Type** screen, select the type of media you want to add to the vault. The **Add Media IDs** screen appears.

Figure 7 Add Media IDs Screen



- 3 On the **Add Media IDs** screen, do one of the following:
- Enter the name for the new media in the **Media Name** field. Go to [Step 6](#) on page 19.
 - Select one or more previously entered media IDs from the list, and then click **Next**. Go to [Step 7](#) on page 19.
 - Add multiple media labels by clicking the **New Media** button. The **Create New Media ID** screen appears.

Figure 8 Create New Media ID
Screen



- 4 On the **Create New Media ID** screen, perform the following steps:
 - 1 Enter the new media label name.
 - 2 Enter the number of media IDs you want to generate for the label.
 - 3 Enter a starting value for the new media ID.
 - 4 If you want to create additional media IDs, click **Apply** and repeat steps 1 through 3.
 - 5 When you are finished creating media IDs, click **OK** to save your changes and exit.

The **Add Media IDs** screen appears again.

- 6 On the **Add Media IDs** screen, select one or more media IDs from the list, and then click **Next** to continue. The **Complete Add Media Task** screen ([Figure 153](#) on page 16) appears.
- 7 On the **Complete Add Media Task** screen, click **Next** to finish adding the media, or **Back** to make changes.
- 8 After the status screen informs you that the media has been successfully added, click **Finish**.

Caution: StorNext will use and overwrite all available tapes, so make sure only the library contains only media you want StorNext to use.

Removing and Moving Media

These procedures enable you to remove or move media from a library to a different library or vault while preserving the existing data.


- [Removing Media](#)
- [Moving Media](#)

Removing Media

Use this procedure to remove blank media, backup tapes, or cleaning media from the StorNext system.

- 1 From the StorNext home page, choose **Remove/Move Media** from the **Admin** menu. The **Remove or Move Media** screen appears.


Figure 9 Remove or Move Media Screen



- 2 Click **Remove Media**.


- 3 Select from the **Select Library** list the library from which to remove the media. If the library has more than one media type, specify the type of media you want to remove from the selected library. Click **Next** to continue. The **Select Media** screen appears.

Figure 10 Select Media Screen




- 4 To remove individual media, enter the media name in the **Enter Media** field and click **OK**. The **Complete Remove/Move Media Task** screen appears. Go to [Step 10](#) on page 23.
- 5 To remove multiple media, click **Browse**. The **StorNext Media Browser** screen appears.

Figure 11 StorNext Media
Browser Screen




- 6 On the **StorNext Media Browser** screen, select the **Media Class** to remove: Blank, Cleaning, or Backup. The **StorNext Media Browser** screen displays in the **Select Media** field all the media associated with the class you have selected. The **Current Media Class** field shows the media class you just selected.
- 7 If desired, you can use the **Media Filter** field to narrow your search of your selected media type. For example, to display all media with 01 in its name, enter ***01*** (where the asterisks are wild cards representing any characters before or after the specified search criteria).
- 8 Select from the **Select Media** list the specific media you want to remove (any number of media from one to all), and then click **OK** to close the **StorNext Media Browser** screen. The **Select Media** screen reappears, showing the selected media.
- 9 Select from the **Media List** the media you want to remove, and then click **Next**. The **Complete Remove/Move Media Task** screen appears.

Figure 12 Complete Remove/
Move Media Task Screen




- 10 Review your selections and click **Next** to apply them or **Back** to make changes. The **Remove/Move Media** status screen appears.
- 11 After the status screen informs you that the media has been successfully removed, click **Finish**. The **Library Operator Interface (LOI)** screen appears.

Figure 13 Library Operator Interface Screen



- 12 The name of the library on which the media resides is displayed. Click **Eject Media**. The **Library Operator Interface** screen updates and indicates the selected library. Available media operations for each library are shown.

Figure 14 LOI Eject Screen



- **Select All**: Click this button to select all media in the **Select Media** list for removal
- **Deselect All**: Click this button to deselect all media in the **Select Media** list

- **Details:** Click this button to view information about selected media
- **Eject:** Click this button to eject (remove) media from the library
- **Fail:** Click this button to mark selected media as Failed
- **Close:** Click this button to close the current window

Note: All of the media operations described may not be available for each library.

13 For the selected media IDs, click **Eject**.

14 Click **Close**.

Moving Media

Use this procedure to move data, blank media, backup tapes, or cleaning media to another library.


- 1** From the StorNext home page, choose **Remove/Move Media** from the **Admin** menu. The **Remove or Move Media** screen ([Figure 157](#) on page 20) appears.
- 2** Select **Move Media**.
- 3** Select from the **Select Library** list the library from which to remove the media. Select from the **Media Type** list the media type. Click **Next** to continue. The **Select Media** screen ([Figure 158](#) on page 21) appears

Note: If you are using a vault library, you must manually move the media.

- 4** To move individual media, enter a media name in the **Enter Media** field and click **OK**. The **Complete Remove/Move Media Task** screen ([Figure 160](#) on page 23) appears. Go to [Step 7](#) on page 25.
- 5** To move multiple media, click **Browse**. The **StorNext Media Browser** screen ([Figure 159](#) on page 22) appears.
- 6** Select a media class. The **Select Media** list is populated with all available media in the class you selected.
- 7** Select from the **Select Media** list the media you want to move, and then click **OK**. The **Select Media** screen reappears, showing your selection.


- 8 Click **Next** to continue. The **Select Destination Library** appears.

Figure 15 Select Destination Library Screen




- 9 Select the destination library from the list, and then click **Next**. The **Complete Remove/Move Media Task** screen ([Figure 160](#) on page 23) appears.
- 10 Click **Next** to continue.
- 11 After the status screen informs you that the task was completed successfully, click **Finish**. The **Library Operator Interface (LOI)** screen appears with the source library flagged.

Figure 16 Library Operator Interface Screen



- 12 Eject the media from the source library, and then click the highlighted **Eject Media** button. The **LOI Eject** screen ([Figure 162](#) on page 24) appears.
- 13 Select the media to be moved, and then click **Eject**. Click **Close** to continue. The **LOI** screen reappears with the destination library flagged.
- 14 Move the media to the destination library and click the highlighted **Enter Media** button. The **LOI Enter** screen appears.

Figure 17 LOI Enter Screen



- 15 Click **Eject** to continue, and then click **Close**. The **LOI** screen reappears without flags or highlights.
- 16 Click **Close** to finish.

Using the SNSM Media Functions

The StorNext Storage Manager's Media menu contains functions that enable you to accomplish the following tasks:

- [Moving Media Manually](#).
- [Mounting and Dismounting Media](#).
- [Removing Media From the Storage Manager](#).
- [Removing Media From the Storage Manager](#).
- [Assigning Blank Media to a Policy Class](#).
- [Transcribing Media](#).
- [Changing Media Attributes](#).
- [Reclassifying a Media Class Grouping](#).
- [Cleaning Media](#).
- [Cancelling the Eject Media Process](#).

These functions are not available if you have only StorNext File System and not StorNext Storage Manager.


Moving Media Manually

Use this function to flag media you plan to move manually from the source library to a destination library. After you use this function you must manually remove the media from the source library and use the Library Operator Interface (LOI) to enter the media into the destination library.

Caution: The library that contains the media you want to manually move must be taken offline before you continue. If the library is not offline, the move process will fail. For information about taking a library offline, see [Changing the Library State](#) on page 29.


- 1 From the SNSM home page, choose **Library > Manual Move** from the **Media** menu. The **Manually Move Media** screen appears.

Figure 18 Manually Move Media Screen



- 2 Select from the **Source Archive** dropdown list the source library that contains the media you want to move.
- 3 Click **Browse** when you are ready to flag specific media for moving manually. The **Media Browser** screen appears.

Figure 19 Media Browser
Screen



- 4 On the **Media Browser** screen, select from the **Select Media** list the media you plan to move manually. If desired, you can reduce the number of entries in the media list by entering a filter in the **Media Filter** field. This field accepts the wildcard character (asterisk *). You can also select all media by clicking the **Select All** button.
- 5 Click **OK** after you are finished selecting media.
- 6 Select from the **Destination Archive** dropdown list the destination library to which you plan to manually move the media you just selected.
- 7 When you are ready to flag the selected media for moving manually, click **Apply**.
- 8 When the Status screen informs you that the media have been successfully flagged for moving manually, click **Close**.
- 9 Complete the process by manually removing the media from the source library.
- 10 Access the Library Operator Interface from the StorNext home page by choosing **Library Operator Interface** from the **Admin** menu. To enter media into the destination library, select the destination library and then click **Enter Media**.

Mounting and Dismounting Media

Use the procedures in this section to mount and dismount media in a tape drive as needed.

Mounting Media

Use the following procedure to mount media.

- 1 From the SNSM home page, choose **Library > Mount** from the **Media** menu. The **Mount Media** screen appears.

Figure 20 Mount Media
Screen

Quantum StorNext Home Help

File Media Admin Reports Help

Mount Media
First select the archive the drive is in. Then select the media ID you wish to mount. Then select the drive you wish to mount the media in. The robot does this operation in automated libraries. In manual libraries, the Mount Media button on the Library Operator Interface (LOI) is highlighted when you need to perform the mount operation.

Archive
-- Select Archive --

Media ID
-- Select Media --

Filter
Filter

Select Drive ID
-- Select Drive --

Apply Reset

StorNext spock Active

- 2 Select from the **Archive** drop-down menu the archive that contains the media you want to mount. The screen updates to show available media in the archive you selected.
- 3 Select from the **Media ID** list the media you want to mount.
- 4 If desired, use the **Filter** field to narrow your search when looking for a specific media ID. For example, enter ***01** and click the **Filter** button to find all media IDs ending in 01.
- 5 Select from the **Drive ID** drop-down menu the drive in which you wish to mount the media.
- 6 Review your selections and then click **Apply**.

- 7 After the Status screen informs you that the task has completed successfully, click **Close**.

Dismounting Media

Use the following procedure to dismount media.

- 1 From the SNSM home page, choose **Library > Dismount** from the **Media** menu. The **Dismount Media** screen appears.

Figure 21 Dismount Media Screen




- 2 Select either a **Media** or **Drive ID** from the drop-down menu. Depending on which ID you selected, the information in the other field automatically populates.
- 3 Review your selections and click **Apply**.
- 4 After the Status screen informs you that the operation was successful, click **Close**.

Removing Media From the Storage Manager

Use this function to remove media from the StorNext Storage Manager. When you use this function all files are removed from the selected media, and the media is ejected from the library into the EIF (mailbox).

- 1 From the SNSM home page, choose **Remove** from the **Media** menu. The **Remove Media From SNSM** screen appears.

Figure 22 Remove Media From SNSM Screen




- 2 Select from the **Select Media** list the media you want to remove. To reduce the number of media shown, you can enter a search filter in the Enter Filter field. This field accepts the wildcard character (asterisk *).
- 3 Click **Apply** when you are ready to remove media.
- 4 When the Status screen informs you that the media were successfully removed, click OK.

Assigning Blank Media to a Policy Class

This function allows you to assign blank media to a policy class. Blank media must exist in the library before you can move the media to a policy class.

- 1 From the SNSM home page, choose **Assign Policy** from the **Media** menu. The **Assign Policy** screen appears.

Figure 23 Assign Policy Screen



- 2 Select from the **Select Media** list the blank media you want to move.
- 3 Select from the **Select Policy Classes** list the policy class to which you want to move the blank media you selected.
- 4 Click **Apply** when you are ready to move the selected blank media.
- 5 When the Status screen informs you that the blank media were successfully assigned to policy classes, click OK.

Transcribing Media

You can either transcribe the contents of one media type to another media type, or reclaim (defragment) media. During the transcription or reclamation process, SNSM uses two drives to transcribe one media to another media, file by file.

Caution: For SNSM to successfully transcribe one media to another media, two drives must be online. If only one drive is online, the transcription or reclamation process fails.

- 1 From the SNSM home page, choose **Transcribe** from the **Media** menu. The **Transcribe Media** screen appears.

Figure 24 Transcribe Media
Screen



- 2 Select from the **Select Media to Transcribe** list one or more media, and then click **Apply**.
- 3 After the status screen informs you that the media has been transcribed, click **Close**.

If transcription or reclamation starts and all the drives are in use, SNSM prioritizes and queues the job. When two drives become available, the queued job starts and all active files are transcribed. When transcription is complete, SNSM updates the database to reflect the new location of the files.

If the transcription or reclamation process encounters a file that spans multiple media, only the portion of the file that exists on the old media is transcribed.

When the transcription or reclamation process is complete, only deleted files remain on the source media. To remove the source copy of the deleted files, you must clean the media as described in [Cleaning Media](#) on page 38. After the cleaning process is complete and the source copy of the deleted files are removed, the media is available for reuse as blank media.

Changing Media Attributes

This function allows you to change the current state of one or more pieces of media. You can change the state to one of the following:

- **Unsuspect:** Indicates that the media is physically sound and not in a suspect (potentially damaged or unusable) condition.
- **Write Protect:** Write protected media cannot be overwritten or have data added.
- **Write Unprotect:** This media can be overwritten or have data added.
- **Available:** The media is available for writing or reading.
- **Unavailable:** The media is not available for writing or reading.
- **Unmark:** Cancels an “Error” or “Checkout” state and makes the media either accessible or available for checkout, respectively.

- 1 From the SNSM home page, choose **Attributes** from the **Media** menu. The **Change Media Attributes** screen appears.

Figure 25 Change Media Attributes Screen



- 2 Specify the new state you want to apply to the media.
- 3 Select from the **Select Media** list the media to which you want to apply the new state you specified.
- 4 Click **Apply** to change the states of the media you selected.

- 5 When the Status screen informs you that the media statuses were successfully changed, click **OK**.

Reclassifying a Media Class Grouping

A media class grouping is a media management tool that segregates media into classes. Membership in a media class grouping is exclusive, so media can belong to only one class grouping. Media class groupings are restricted to one type of media. Valid media types are Clean, Data, and Migrate.

For example, LTO media cannot belong to a media class grouping that contains AIT media. When media is initially imported into the SNSM, it enters a media class grouping. Import mapping to media class is automatic and driven by media type.


- 1 From the SNSM home page, choose **Reclassify** from the **Media** menu. The **Reclassify Media** screen appears.

Figure 26 Reclassify Media Screen



- 2 Select from the **Source Media Class** list the desired source media class, and then click **Browse** to select the media IDs you want to reclassify. The **Media Class Browser** window appears.

Figure 27 Media Class Browser Window



- 3 Select one or more media from the **Select Media** list, and then click **OK** to return to the **Reclassify Media** screen. The IDs for the selected media are shown in the **Entered Media IDs** list.
- 4 Select from the **Destination Media Class** list a destination for media IDs, and then click **Apply**.

The new media class must be associated with the type of media you are reclassifying. For example, if you select 3590 DATA as your source media class, select 3590 CLEAN as your destination media class.

Caution: Your source and destination media types must be the same.
For example:
Source: F0_3590_DATA
Destination: F0_3590_CLEAN

- 5 After the Status screen informs you that the operation was successful, click **Close**.


Cleaning Media

This function allows you to clean a file system, policy class, or media. Periodic cleaning helps prevent inactive information from growing to an unmanageable size. When you run this function, the StorNext Storage Manager removes inactive files that have not been accessed since the specified endtime. This process does not affect current file versions on the media.

Caution: Inactive file versions cleaned from the media cannot be recovered or used again.

- 1 From the SNSM home page, choose **Clean** from the **Media** menu. The **Clean Media** screen appears.

Figure 28 Clean Media Screen




- 2 Do one of the following:
 - Select from the **File System** list the file system you want to clean.
 - Select from the **Media** list IDs for specific media you want to clean.
 - Select from the Policy Class list specific policy classes you want to clean.
- 3 Specify at the **Endtime** field an ending time for the cleaning process, using the format YYYY:MM:DD:hh:mm:ss. For example, enter **2007:12:06:10:08:00** for an endtime of 10:08 on December 6, 2007.
- 4 Click **Apply** to start the cleaning process.
- 5 When the Status screen informs you that the operation was completed successfully, click **OK**.

Cancelling the Eject Media Process


Use this function to cancel the eject process for selected media. The media remains in the assigned slots.

- 1 From the SNSM home page, choose **Library > Library State** from the **Admin** menu. The **Cancel Eject Process** screen appears.

Figure 29 Cancel Eject Process Screen



- 2 Select from the **Media ID** list the media for which you want to cancel the eject process. If necessary, select a library from the **Library** list to view and select media in that library. To select all media in the Media ID list, click the **Select All** button.
- 3 Click **Apply** to cancel the eject process for the selected media.
- 4 When the Status screen informs you that the operation was successful, click **OK**.



Chapter 10

Managing Storage Disks

StorNext provides the ability to configure *storage disks* that function and operate the same way as physical tape media. Storage disks are external devices on UNIX-based file systems that can be used for long term data storage.

When a storage disk is configured, the StorNext Storage Manager moves data to storage disks for long-term retention in addition to, or instead of tape. This enables users to leverage the specialized third-party functionality of appliances or store small files that might take longer to retrieve from tape. Many users will still use tape for long-term storage and vaulting, but storage disk can be used to create tape-free archives.

Here are a few differences storage disks have over tape media, aside from the obvious cost-saving benefit:

- A storage disk either belongs to no policy class, or belongs to a single policy class
- A storage disk can store file copies only with the same copy ID

Note: Before you create a storage disk, the disks you plan to use must reside in an existing, mounted file system.

After you create a storage disk, observe the following usage recommendations:

- If your file system includes storage disks, avoid using that file system for any data other than storage disk stored data.

- Use complete and physically dedicated file systems (snfs, local, nfs, or other,) for storage disk data, not shared file systems or file systems with linked directories.
- If your file system includes storage disks and you accidentally fill it with unrelated user data (i.e., non-storage disk data,) call the Quantum Technical Assistance Center and ask for a procedure to clean up and transcribe data.

Storage Disk Deduplication

StorNext supports storage disk *deduplication only on non-managed file systems*. Deduplication frees disk space by eliminating redundant data. The deduplication process does not retain duplicate data, so there is only one copy of the data to be stored. (Indexing of all data is retained in case that data is required later.) The main benefit of deduplication is that it reduces storage capacity requirements because only unique data is stored. Without deduplication, offline copies of a file consume as much disk space as the original file.

When you create a new storage disk, you will be given the option of enabling deduplication. StorNext refers to a storage disk with deduplication enabled as a *dedup SDISK*. If your system configuration consists only of storage disks, the same rules that apply to storage disks apply to deduplication-enabled storage disks. For example, in a storage disk-only configuration the first storage disk must always use file copy 1.

You can create up to 4 dedup sdisks. (You can have a total of 16 storage disks, of which 4 can be dedup sdisks.)

You must have a minimum of 2GB of RAM for each dedup sdisk you plan to use.

Note: The 2GB of RAM per dedup sdisk is in *addition* to the memory required for StorNext.

At this time storage disk deduplication is supported only on 32 bit and 64 bit Linux platforms.

Adding a Storage Disk

Storage disks are treated the same as media in the system. Before you configure a new storage disk, the disk you want to use must be in a file system that is already created and mounted.

Note: When you are creating storage disks, exercise caution before enabling the deduplication feature. Once you create a dedup SDISK, you cannot change it to a non-enabled storage disk. Conversely, you cannot convert a non-enabled storage disk to a dedup SDISK. However, you can delete a non-enabled storage disk or dedup SDISK and then recreate the storage disk with deduplication either enabled or disabled.

Use the following procedure to add storage disks.


- 1 From the SNSM home page, choose **Storage Disk > Config** from the **Admin** menu. The **Configure Storage Disk** screen appears.

Figure 1 Configure Storage Disk Screen




- 2 Click **Add**. The **Add Storage Disk - Introduction** screen appears, listing any currently configured storage disks.

Figure 2 Add Storage Disk -
Introduction Screen



3 Click **Next**. The **Add Storage Disk** screen appears.

Figure 3 Add Storage Disk
Screen



4 Enter the fields on the **Add Storage Disk** screen.


- **Storage Disk Name:** Enter a name for the storage disk, or accept the displayed default name.
- **Enable Deduplication:** Select this option if you want to enable deduplication and create a dedup SDISK. Once deduplication is enabled, you cannot disable it after the storage disk has been created. Likewise, if deduplication is NOT enabled at this time, you cannot enable it later.
- **Select a File System Mount Point:** Select a file system mount point for the storage disk. If deduplication is enabled, **the mount point must be on a non-managed StorNext file system.**
- **Click Browse to Select a Directory for Storage Disk Files:** Click the **Browse** button to display the **Directory Browser** window. On this window select or create a directory on which to store storage disk files.

- **Copy # Used for all Policy Classes:** Select the copy number (Copy 1, Copy 2, Copy 3, or Copy 4) for the new storage disk. This copy number will be used by all policy classes.

Note: Only copies of the number you choose can be written to this disk. For example, if you select copy number 1, only files with that copy number can be written to the storage disk.

- 5 Click **Next**. The **Complete Storage Disk** screen appears.

Figure 4 Complete Storage Disk Screen




- 6 Review information for the new storage disk. Click **Next** to complete the task or **Back** to make changes.
- 7 After the Status screen informs you that the storage disk was successfully added, click **Finish**.

Modifying a Storage Disk

Use this procedure to modify a previously configured storage disk.

- 1 From the SNSM home page, choose **Storage Disk > Config** from the **Admin** menu. The **Configure Storage Disk** screen ([Figure 178](#) on page 13) appears.
- 2 Select a storage disk from the **Current Storage Disks** list, and then click **Modify**. The **Modify Storage Disk** screen appears.

Figure 5 Modify Storage Disk Screen



Note: For a blank storage disks (i.e., a storage disk that has not been written to, and one where no file system files reside,) you can modify any of the parameters on the **Modify Storage Disk** screen. If the storage disk has been written to, you can change only the number of streams.

- 3 Modify any of the following information:

- **Mount Point:** The file system mount point for the storage disk. You can change the mount point only on a blank or unused file system. To change the mount point, select an existing mount point from the drop-down list, and then click **Browse**. The directory changes in the **Mount Point** field. (You can modify this parameter only for an unused storage disk.)
- **Directory:** The directory selected for file storage when the storage disk was initially created. (You cannot modify this setting.)
- **# of Streams:** The number of streams (1 - 8) I/O streams that can concurrently write to the disk.
- **Copy:** The copy number (1-4) specified when the storage disk was created. This copy number will be used by all policy classes. (You can modify this parameter only for an unused storage disk.)

Note: The **Deduplication** field indicates whether deduplication is enabled for the storage disk. This field cannot be modified.

- 4 After making storage disk modifications, click **Apply**.
- 5 After the Status Screen informs you that your modifications were made successfully, click **Close**

Deleting a Storage Disk

Use the following procedure to delete a storage disk.

- 1 From the SNSM home page, choose **Storage Disk > Config** from the **Admin** menu. The **Configure Storage Disk** screen ([Figure 178](#) on page 13) appears.
- 2 Select the storage disk you want to delete, and then click **Delete**.

A message will warn you that all data on the disk will be deleted with the disk, and ask you to confirm that you want to continue.

Caution: All data on the storage disk will be removed when you continue, so proceed with extreme caution. There is no “undo” feature that will undelete the storage disk if you change your mind, so be absolutely certain you want to delete the storage disk before you continue.

Figure 6 Delete Warning Message



- 3 Click **OK**.
- 4 After the Status screen informs you that the storage disk was successfully deleted, click **Close**.

Note: Deleting a storage disk does not unmount it. After deleting you must manually unmount the disk. Alternatively, you can use the mounted storage disk to create a new storage disk.

Changing a Storage Disk State

Changing a storage disk’s state means changing the storage disk’s logical state to online or offline.

- 1 From the SNSM home page, choose **Storage Disk > State** from the **Admin** menu. The **Change Storage Disk State** screen appears.

Figure 7 Change Storage Disk State Screen



- 2 Select the storage disk whose state you want to change.
- 3 Select the state (**Online** or **Offline**) to assign to the selected storage disk, and then click **Apply**.

Note: The “Online Pending” state applies only to deduplication-enabled storage disks. This state means a verification process is currently in progress. When verification is complete, the status automatically changes to “Online.” When a storage disk is in the “Online Pending” state, you can retrieve from the storage disk but cannot make modifications. Also, no store or delete operations are performed on the storage disk.

When a storage disk is in an “Online Pending” state, you can change the state only to “Offline.” If you change the state from “Online Pending” to “Online,” the operation will fail.

- 4 After the Status screen informs you that the storage disk’s status was successfully changed, click **Close**. The **Change Storage Disk State** screen shows the changed state for the storage disk you selected.

- 5 If desired, repeat steps 2 - 4 to change the state for additional storage disks.

Cleaning a Storage Disk

The cleaning process scans the storage disk and removes inactive files that have not been accessed since the endtime, and orphaned file copies (i.e., archive copies on the file system disk,). Using this process helps prevent the inactive information from growing to an unmanageable size.

This process does not affect current file versions on the media.

Caution: Inactive file versions cleaned from the media cannot be recovered or used again.


- 1 From the SNSM home page, choose **Storage Disk > Clean** from the **Admin** menu. The **Clean Storage Disk** screen appears.

Figure 8 Clean Storage Disk Screen



- 2 Select from the **Select Storage Disk List** the storage disk you want to clean, and then click **Apply**.
- 3 After the Status screen informs you that the storage disk was successfully cleaned, click **Close**.

Note: When you clean a deduplication-enabled storage disk, blocklets are not immediately freed from the associated blockpool. The unused blocklets will be freed when the weekly clninfo schedule is run. If you want to immediately free blocklets, you can run the `fscclean -b` command from the command line.



Chapter 11

Data Migration Management

This chapter describes how to use StorNext to manage data. This chapter covers these topics:

- [Policy Classes and Relationships](#)
- [Adding a Storage Policy](#)
- [Modifying a Policy Class](#)
- [Deleting a Policy Class](#)
- [Applying a Policy Class](#)

This chapter includes procedures executed from the command line. Before initially executing any StorNext command line programs, you must first source either the `.profile` or the `.cshrc` file to update the user environment with StorNext environment variables.

If you are running `sh`, `ksh`, or `bash`, type: `./usr/adic/.profile`

For all other shells, type: `source /usr/adic/.cshrc`

Note: Files do not migrate by default policy rules if the time is set in the future. Before migrating files, verify that the time settings on all of your client and server machines are synchronized.

Policy Classes and Relationships

A policy class defines how files will be managed in a directory and subdirectories. These are the available policy class settings:

- Number of copies to create
- Media type to use when storing data
- Amount of time to store data after data is modified
- If disk-to-disk relocation is enabled, the amount of time (in days) before relocating a file
- Amount of time before truncating a file after a file is modified

Policy classes can be related to one or more directories. In this situation, all files in that directory and sub-directories are governed by the policy class. The connection between a policy class and a directory is called the relation point.

Here are some examples of policy class usage:

- A directory in which to store backups every night is created. This directory is seldom accessed after the files are copied over. A policy class could be set up to create two tape copies of the files, store one copy of the files to AIT media after residing on disk for 10 minutes, and then truncate the other set of files immediately after storing the other set to tape in order to free up disk space. This policy can be associated with a directory such as: `/sandsm/dsm1/backup`.
- A directory has been created to store all documents that are accessed frequently, and if truncated, need to be retrieved quickly. The policy class in this case could be set up to create a single tape copy, store the files to LTO media 15 minutes after being on disk, and then truncate after 60 days of non-use. This policy can be associated with a directory such as: `/sandsm/dsm1/docs`.

Stub Files

StorNext includes a licensable *Stub File* feature. When this feature is enabled, third-party applications can gather information about a file by reading a portion of the file (called a stub) rather than reading the entire file. When you create a policy class you can enable stub file support and specify the size of the stub file (in kilobytes). When stub file support is enabled, the beginning portion of the file (up to the size you

specified) remains on disk after data blocks are freed during policy management or space management.

Disk-to-Disk Relocation

Disk-to-Disk relocation allows you to move data from one set of disks (disk stripe group) to another without affecting the file name space. You can perform this procedure only if you have the full StorNext suite. (Customers with StorNext File System only cannot perform this procedure.)

Before enabling disk-to-disk relocation, you must perform these preliminary tasks if you have not already done so:

- Create a new managed file system with no more than two affinities, or reconfigure an existing managed file system by adding affinities. (For instructions on creating the file system, see [Managing the File System](#) on page 11.)
- Create a policy class or modify an existing one.
- Create a relation point.
- Determine relocation criteria.

After you have accomplished these preliminary tasks, you can enable disk-to-disk relocation.

Adding a Storage Policy


Use the following procedure to add a new storage policy. When you create the policy, you can enable disk-to-disk relocation, stub file support, or both features.

Note: You cannot create a policy class when a client is NFS mounted to an active StorNext file system. When you create a policy class, the StorNext GUI executes the command `fsaddrrelation`. The `fsaddrrelation` command fails if the specified directory is not empty and the associated file system is busy. When a client is NFS mounted, the NFS daemon holds a reference in the directory. Consequently, the unmount fails, aborting the `fsaddrrelation` process.

For more information, see the man page for `fsaddrrelation`.


- 1 From the StorNext home page, choose **Add Storage Policy** from the **Config** menu. The **Storage Policy Introduction** screen appears, showing any previously configured policy classes.

Figure 1 Storage Policy Introduction Screen



- 2 Click **Next** to continue. The **Policy Class and Directory** screen appears.

Figure 2 Policy Class and Directory Screen



3 Select any of the following options:

- **Enable Disk-to-Disk:** Select this option to activate the Disk-to-Disk Relocation feature. For more information about Disk-to-Disk Relocation, see [Disk-to-Disk Relocation](#) on page 13.
- **Enable Checksum Generation:** If enabled, checksums are generated and retained (in the database) for files stored by the corresponding policy class
- **Enable Checksum Validation:** If enabled, checksums are compared to retained values for the files retrieved by the corresponding policy class
- **Enable Stub Files:** Select this option to enable the Stub Files feature. For more information about Stub Files, see [Stub Files](#) on page 12.

4 Click **Browse**. The **Directory Browser** window appears.


Figure 3 Directory Browser Window



- 5 Select a directory that contains the files you want to migrate, and then click **OK**. The **Policy Class and Directory** screen ([Figure 187](#)) reappears, showing the directory you specified.
- 6 Review your selections and click **Next**. If you enabled the disk-to-disk relocation feature, the **Relocation Policy Selection** screen appears.


If you did not enable disk-to-disk relocation, go to [Step 8](#) on page 17.

Figure 4 Relocation Policy Selection Screen



- 7 On the **Relocation Policy** screen, select from the first drop-down list the affinity from which to move your data. Select from the second drop-down list the affinity to which you want to move the data.
- 8 Click **Next** to continue. The **Store, Truncate, and Relocate Time** screen appears.


Figure 5 Store, Truncate, and Relocate Times Screen



- **Minimum Store Time (Minutes):** The minimum amount of time a file must remain unaccessed before it is considered a candidate for storage
- **Minimum Truncation Time (Days):** The minimum number of days a file must remain unaccessed before it is considered a candidate for truncation
- **Minimum Relocation Time (Days):** The minimum number of days a file must remain unaccessed on the primary affinity before it is considered a candidate for relocation to a secondary affinity. This option does not appear when you select the **Enable Stub Files** option on the **Policy Class and Directory** screen.
- **Stub File Size (Kbytes):** The desired target size to allocate for the file stub. This is the readable portion of the file that remains after truncation when the Stub File feature is enabled. This option appears only when you selected the stub file option.


After entering values, click **Next** to continue. The **Number of File Copies and Media Type** screen appears.

Figure 6 Number of File Copies and Media Type Screen



- 9 Select up to four copies to store for each file, including the primary file (File Copy 1). Select the media type for each copy, and then click **Next**. The **Complete Storage Policy Task** screen appears

Figure 7 Complete Storage Policy Task Screen



- 10 Review the information and click **Next** to continue, or **Back** to edit your choices.
- 11 After the Status screen informs you that the procedure completed successfully, click **Close**.

Adding a Policy Class Through SNSM Without a Relation Point

Use this procedure from the SNSM home page to add a policy class without a relation point.


- 1 From the SNSM home page, choose **Policy Classes** from the **Admin** menu. The **Manage Policy Classes** screen appears.

Figure 8 Manage Policy Classes Screen



2 Click **Add** to continue. The **Add Policy Class** screen appears.

Figure 9 Add Policy Class Screen



- 3 Enter in the **Policy Class Name** field the name of the new policy class.
- 4 Click **Properties** to change the default parameters of the new policy class. The **Modify Parameters** screen appears.

Figure 10 Modify Parameters Screen

Quantum. Modify Parameters for "pc1"

Standard Options

File Copy 1 LTO

File Copy 2 --- Media Types ---

File Copy 3 --- Media Types ---

File Copy 4 --- Media Types ---

File Age Before Migration 10 Minutes File Age Before Truncation 5 Days

Advanced Options

Truncate File Immediately After Store

Max Inactive Versions 10

Drive Pool to Use fs_F0drivepool

Checksum Generation Checksum Validation

Stub Files Stub File Size KB

Minimum Set Store Size MB

Soft Limits 20000 Hard Limits 25000

Auto Store

Name	Run Days	Start Time	Start Window
------	----------	------------	--------------

Add Delete

Enable Disk-To-Disk

Affinity From a1 Affinity To a1

File Age Before Relocation Days

Enable Retrieve to Affinity

Affinity Name a1

OK Cancel

- 5 On the **Modify Parameters** screen, enter both the Standard Options and Advanced Options as desired.

- **File Copy 1, 2, 3, and 4:** The copy number used when storing assigned media. The copy number assigned to the media determines which copy goes to the media. You must specify media for all copies used. You must use at least one copy (Copy 1), and can use up to four copies.

If you have created more than one drivepool, you can accomplish file steering by specifying which drivepool to use for each available file copy. For example, you could send data on File Copy 2 to "fs_drivepool_2," and data on File Copy 3 to "fs_drivepool_3." (Before you can assign a drive pool to a file copy, you must first create the drive pools by using the Configure Drive Pools function.) For more information about configuring drive pools, see [Adding a Drive Pool](#) on page 22.

- **File Age Before Migration:** This value determines the minimum number of minutes a file must reside unmodified on disk before it is considered a candidate for storage on media.
- **File Age Before Truncation:** This value determines the minimum number of days a file must reside on a disk unaccessed before it is considered a candidate for truncation from disk. Truncation removes the disk blocks of a stored file, but not the file itself.
- **Truncate Immediately After Store:** Enable this option (check this box) to truncate files immediately after they are stored.
- **Max Inactive Versions:** The maximum number of inactive versions of a file StorNext keeps track of for recovery purposes.
- **Drive Pool to Use:** Associates the drive pool to use with the policy class. If you specify a drive pool, the drive pool name must be defined before any data operation can occur.
- **Checksum Generation:** If this option is enabled, (the box is checked), checksums are generated and retained in the database for files stored by the corresponding policy class.
- **Checksum Validation:** If this option is enabled (the box is checked), checksums are compared to retained values for the files retrieved by the corresponding policy class. The Checksum feature consumes additional space in the StorNext database whether it is enabled or not. When disabled, this feature consumes approximately 2 bytes per stored file; when enabled, this feature consumes approximately 18 bytes per stored file. The database stores data in files on the host computer, so the

increase in database size translates to a corresponding increase in disk space requirements. The exact amount of space consumed (whether the feature is enabled or disabled) may vary.

- **Stub Files:** Select this option to enable the Stub File feature.
- **Stub File Size:** When the Stub File feature is enabled, specify the target size for the stub file in kilobytes.
- **Minimum Set Store Size (1 to 999 in MB or GB):** This value determines the minimum size (in megabytes or gigabytes) all valid store candidates in the policy class combined must reach before they are stored.
- **Maximum File Store Age (1 to 720 in hours):** This value determines the time after which any valid store candidate in the policy class is stored.
- **Soft Limits:** This value represents the soft limit on the number of media allocated for the policy class.
- **Hard Limits:** This value represents the hard limit on the number of media allocated for the policy class.
- **Auto Store:** Use this option to automatically store files for the current policy class. If this option is disabled (unchecked), Quantum recommends that the files for the policy class be stored by scheduled events.
- **Add or Delete Schedule:** Add or delete a scheduled store event. To create or delete a scheduled event, follow the procedures described in [Scheduling StorNext Events](#) on page 30
- **Enable Disk-To-Disk:** Enables the disk-to-disk relocation functionality. Before you can enable the disk-to-disk functionality on this screen, two affinities (the From and To affinities described below) must be created.
 - **Affinity From:** The primary affinity where a file resides. (E.g, A1)
 - **Affinity To:** The secondary affinity to which the file will relocate. (E.g, A2)
 - **File Age Before Relocation (in days):** This value determines the minimum time in days a file must reside unaccessed on the primary affinity before being relocated to a secondary affinity.

- **Enable Retrieve to Affinity:** Enable this option if you want to restore to an alternative affinity. When enabled, this feature allows you to retrieve truncated files to a specified affinity that might be different from the affinity assigned during ingest. (StorNext supports up to two affinities.)

For example, you could migrate files from A1 (affinity 1) to tape storage, but have the policy pull a file back to A2 (affinity 2) instead of to A1.

- **Affinity Name:** Specify which of the two affinities you want the policy to retrieve files from.
- 6 After you are finished setting parameters for the policy class, click **Apply**.
 - 7 After the Status screen informs you that the policy class was successfully added, click **Close**.

Adding a Relation Point to a Policy Class

After you have added a policy class, you can use this procedure to add a relation point to it.

- 1 From the SNSM home page, choose **Relations** from the **Admin** menu. The **Manage Policy Class Relationships** screen appears.

Figure 11 Manage Policy Class Relationships Screen



- 2 Select the policy class for which you wish to create a relation point, and then click **Add**. The **Add Relationship** screen appears.

Figure 12 Add Relationship Screen



- 3 Click **Browse** to view directories to which you can make the relationship. The **Directory Browser** window appears.

Figure 13 Directory Browser Window



- 4 Select from the list a directory to which to add the relationship directory. A second **Directory Browser** window appears, showing the directory you selected in the **Current Directory** field.
- 5 Click **OK**. The **Add Relationship** screen appears, showing the new directory.
- 6 Click **Apply**.
- 7 After the Status screen informs you that the operation completed successfully, click **Close**.

Modifying a Policy Class

Use the following procedure to modify a policy class.

- 1 From the SNSM home page, choose **Policy Classes** from the **Admin** menu. The **Manage Policy Classes** screen ([Figure 193](#)) appears.
- 2 Select from the **Select Policy Class** list the policy class you want to modify, and then click **Modify**. The **Modify Parameters** screen ([Figure 195](#)) appears.
- 3 On the **Modify Parameters** screen, modify values as desired and then click **OK**.

- 4 After the status screen informs you that the policy class has been successfully modified, click **Close**.

Deleting a Policy Class

Before you can delete a policy class, you must first delete everything associated with that policy class, including media and directories. Also, you must clean the media and return it to a system-blank state.


To remove all files in directories associated with the policy class, follow these steps:

1. Log on as root.
2. Change directories to the location where the relation resides.
3. Remove all files and directories.

Use the following procedure to delete a policy class.

- 1 To remove the relations for the policy class, from the SNSM home page choose **Relations** from the **Admin** menu. The **Manage Policy Class Relationships** ([Figure 196](#)) appears.
- 2 In the list, select a policy class and click **Remove**. The **Remove Relationships for...** screen appears.

Figure 14 Remove Relationships Screen




- 3 In the **Select Directory** list, select a directory and click **Apply**.
- 4 After the Status screen informs you that the relationship has been successfully removed, click **Close**.
- 5 From the SNSM home page, choose **Policy Classes** from the **Admin** menu. The **Manage Policy Classes** screen ([Figure 193](#)) reappears.
- 6 Select from the **Select Policy Class** list the policy class you want to delete, and then click **Delete**.
- 7 After the Status screen informs you that the policy class has been successfully deleted, click **Close**.

Applying a Policy Class

After you have created at least one policy class, you can select that policy class and apply it immediately.

- 1 From the SNSM home page, choose **Policy Classes** from the **Admin** menu. The **Manage Policy Classes** screen ([Figure 193](#)) appears.
- 2 Select from the **Select Policy Class** list the policy class you want to apply.

- 3 Click **Run** to apply the policy class immediately.
- 4 After the Status screen informs you that the operation was performed successfully, click **OK**.



Chapter 12

StorNext Reports

StorNext enables you to generate and view a number of reports that provide information about your StorNext system. These reports are selectable from the StorNext home page, the SNFS home page, and the SNSM home page.

The following reports are available:

- [The Backup Information Report](#): Provides information about primary and secondary backups for the StorNext system.
- [The Drive State Information Report](#): Provides information about the drives configured in StorNext, including drive IDs and types, archives, usage, errors, and mount data.
- [The File Information Report](#): Provides file information, including the current location of a file, owner, size, and number of copies for the file.
- [The Library Information Report](#): Provides information about libraries, including the type and state, associated media associated, and imported media classes.
- [The Library Space Used Report](#): Provides total storage capacity being used in all configured libraries.
- [The Media Information Report](#): Provides media configuration and statistics.

- [The Policy Class Information Report](#): Provides information about policy classes in the StorNext system, including media, noontimes, hard limits, and cleanup.
- [The Relation Information Report](#): Provides information about the directory-to-policy class relationships in the file system.
- [The Request Information Report](#): Provides information about requests, including the type of request, priority, and current status.
- [The Scheduler Information Report](#): Provides information about StorNext schedules.
- [The Storage Disk Information Report](#): Provides information about the configured storage disks on your system.

SNFS Reports

SNFS reports are also accessible from the StorNext home page, the SNFS home page, and the SNSM home page.

The following reports are available:

- [The Directory Affinity Report](#): This report shows the existing affinities for a selected directory in the file system.
- [The File System Statistics Report](#): This report provides file system statistics including active clients, space, size, disks, and stripe groups.
- [The Stripe Group Statistics Report](#): This report provides statistics for the stripe group, such as space, affinities, and current statuses.
- [The File System Client Report](#): This report provides statistics for StorNext clients, including the number of StorNext SAN clients and distributed LAN clients, and client performance.
- [The File System LAN Client Report](#): This report provides information about distributed LAN clients, including read and write speed.

The Backup Information Report

The Backup Information Report provides the following information:


- **Date**: The date and time the last backup was run

- **Type:** The type of backup that was run: Full or Partial
- **Status:** The current status of backup: PASS, STORED, or NOT. (NOT means all copies of the backups were not stored to media.)
- **Media:** The media ID on which the backup was stored

Use the following procedure to run the Backup Information Report.

- 1 Choose **Backups** from the **Reports** menu. The **Backup Information Report** appears, showing information for the backups.

Figure 1 Backup Information Report



- 2 Click **Close** when you are finished viewing the report.

The Drive State Information Report

The Drive State Information Report provides the following information:


- **Drive ID:** The drive ID (1, 2, 3, etc.)
- **Drive Type:** The type of drive being used (such as LTO)

- **Device Pathname:** The drive's pathname
- **Tape Compression:** Indicates whether tape compression is currently turned on or off
- **Associated Library:** The name of the library associated with the drive
- **State:** Indicates whether the drive is currently online or offline
- **Status:** Indicates the drive's current status, which is one of these conditions:
 - **FREE:** No medium is mounted
 - **IN USE:** Medium is in use
 - **FAILED:** Drive failed
 - **DELAYED:** Drive is in a dismount-delayed state
 - **CLEANING:** Cleaning medium is mounted
 - **USER MOUNT:** Medium was mounted using the GUI's **Media > Mount** command or the CLI `fsmount` command
 - **OTHER:** Medium was not mounted directly by the StorNext Tertiary Manager, but by other means such as through the Media Manager
- **Dismount Delay:** The time delay (in seconds) before media is dismounted
- **Assignment:** Indicates whether the drive is currently free or mounted
- **Mount Count:** The number of times the tape has been mounted
- **Error Count:** The number of errors (mostly write errors) a tape has incurred
- **Mount State:** Indicates whether the drive is currently mounted or unmounted
- **Mounted Media ID:** The ID of the mounted drive

Use the following procedure to run the Drive State Information Report.


- 1 Choose **Drives** from the **Reports** menu. The **Drive States Report** screen appears.

Figure 2 Drive States Report
Screen



- 2 Select from the drives list one or more drives to include in the report, and then click **Apply**. The **Drive State Information Report** appears with state information for the selected drives.

Figure 3 Drive State Information Report



The screenshot shows a web browser window titled "Drive State Information Report - Microsoft Internet Explorer". The page header includes the Quantum logo and the title "Drive State Information Report". The report content is organized into two sections, one for each drive.

Drive Name: scsi_archive1_dr1 Drive Pools: (fs_F0drivepool)

Drive ID	Drive Type	Device Pathname	Tape Compression	Associated Library	State	Status
1	LTO	/dev/sg13	On	scsi_archive1	Online	FREE
Dismount Delay	Assignment	Mount Count	Error Count	Mount State	Mounted Media ID	
300	Free	0	0	Unmounted	None	

Drive Name: scsi_archive1_dr2 Drive Pools: (fs_F0drivepool)

Drive ID	Drive Type	Device Pathname	Tape Compression	Associated Library	State	Status
2	LTO	/dev/sg14	On	scsi_archive1	Online	FREE
Dismount Delay	Assignment	Mount Count	Error Count	Mount State	Mounted Media ID	
300	Free	0	0	Unmounted	None	

At the bottom of the report, there is a "Close" button with a red 'X' icon.

3 Click **Close** when you are finished viewing the report.

The File Information Report

The File Information Report provides the following information:


- **Last Modification Timestamp:** The timestamp (date) when the file was last modified
- **Owner (Access):** The access permissions assigned to the file owner
- **Group (Access):** The access permissions that group has to the file: Read, Write, Execute (rwx), or all three permissions
- **Public Access:** Indicates public access permissions
- **Policy Class:** An associated policy class that manages the file lifecycle
- **Size (bytes):** The size of the file in bytes
- **Checksum:** Indicates whether a checksum exists for the file
- **Truncation Immediately After Store:** Indicates whether files are truncated immediately after a store
- **File Location:** The media on which the file is stored

- **Copies:** The number of copies of the file
- **Affinity:** The affinity with which the file is associated
- **Set Stub Size (KB):** If the Stub File feature is enabled, this is the target size (in kilobytes) for the stub file
- **Actual Stub Length (KB):** If the Stub File feature is enabled, this is the actual size (in kilobytes) of the stub file

Use the following procedure to run the File Information Report.

- 1 Choose **Files** from the **Reports** menu. The **Files Report** screen appears.


Figure 4 Files Report Screen



- 2 Select the files to include in the report by doing one of the following:
 - In the **Enter Individual File** field, type a file name. Proceed to [Step 5](#) on page 280.


- Select multiple files from the **Select Files** list. (You can click the **Select All** button to select all files in the list, or click **Deselect All** to deselect all selected files.) Proceed to [Step 5](#) on page 280.
 - Click **Browse** to select files from a managed directory.
- 3 When you click **Browse**, the **StorNext File Browser** screen appears. In the **Select Directory** list, select a directory. The selected directory appears in the **Current Directory** field. A list of directories appears in the **Select Directory** list, and a list of files in the **Current Directory** field appears in the **Select Files** list.

Figure 5 StorNext File Browser Screen



- 4 Select one or more files from the **Select Files** list, and then click **OK**. (To select multiple files, hold down the **Control** key while clicking the file you want to select.) The Files report screen appears with the selected files in the **Select Files** list.
- 5 Click **Apply** to run the report. The **File Information Report** appears with information for the selected files.

Figure 6 File Information Report



6 Click Close when you are finished viewing the report.

The Library Information Report


The Library Information Report provides the following information:

- **Library Type:** The library type, such as SCSI, Network, or Vault
- **Current State:** Indicates whether the library is online or offline
- **Library Mode:** Displays two modes:
 - **Attended:** The LOI will be associated with required actions
 - **Unattended:** Fails any actions that require user interaction
- **Media Type:** The media type the library uses
- **Slot Count:** The number of media slots available for the library type
- **Current Fill Count:** The current number of media slots associated with the library
- **Import Media Class:** The media type to import

Use the following procedure to run the Library Information Report.


- 1 Choose **Libraries** from the **Reports** menu. The **Libraries Report** screen appears.

Figure 7 Libraries Report Screen



- 2 Select from the **Library List** one or more libraries on which to run the report, and then click **Apply**. The **Library Information Report** appears with information about the selected libraries.

Figure 8 Library Information Report



Library Type		Current State	Library Mode	
SCSI		Online	Attended	
Media Type	Slot Count	Current Fill Count	Import Media Class	
LTO	18	16	F0_LTO_ADDBLANK	
Media Type	Slot Count	Current Fill Count	Import Media Class	
LTOW	18	0	None	

X Close

3 Click **Close** when you are finished viewing the report.


The Library Space Used Report

The Library Space Used Report shows the amount (in gigabytes) of storage currently used by all configured libraries.

Use the following procedure to run the Library Space Used Report.

- 1 Choose **Library Space** from the **Reports** menu. The **Library Space Used Report** appears.

Figure 9 Library Space Used Report



2 Click **Close** when you are finished viewing the report.

The Media Information Report

The Media Information Report provides the following information:

- **Copy #:** The number next to the media ID in the upper left corner that refers to the corresponding copy. (In the illustration, notice the (1) to the right of 000091.) If the copy number does not exist, this media has not been allocated to a policy class.
- **Show Details** link: Click this link to view media dead space and files on that media. (See [Figure 13](#) on page 288.) Depending on the number of files on the media, it could take a long time before data appears after you click the **Show Details** link.
- **Media Type:** The type of media
- **Media Class:** The media class designated for the tape
- **Policy Class:** Indicates whether a policy class is associated with the media
- **Last Access Time:** The date and time when the media was last used

- **Media Status:** Indicates media status: Available, Unavailable, or In Use
- **Write Protect:** Indicates whether the media is write-protected
- **Formatted:** Indicates whether the media is formatted
- **Import Date:** The date the media was added to the library
- **Export:** This attribute is currently not utilized in StorNext, and should remain at status UNMARKED
- **# Files:** The number of files associated with the media
- **Space Used:** The amount of space written
- **% Used:** The percentage of used space
- **Space Remaining (bytes):** The available space on the media (in bytes)
- **Mount Count:** The number of times the tape has been mounted
- **Move Count:** The number of times the tape has been moved
- **Suspect Count:** Indicates whether the media has any errors. If so, errors are marked as suspect
- **Current Action:** This is currently not utilized in StorNext and should remain at status NONE
- **Location:** The location of the tape
- **Current Archive:** The current library in which the media is located
- **Pending Archive:** Indicates whether the media is associated with another library


This Detailed Media Information Report (displayed after you click the **Show Details** link) provides the following information:

- **Dead Space:** The amount of unused space on the media
- **File Pathname:** The file's path location
- **File Size:** The file's size
- **Version:** The file's current version
- **Status:** The file's current status: Active or Inactive
- **Modify/Delete Date:** The date the file was last modified or deleted

Use the following procedure to run the Media Information Report.


- 1 Choose **Media** from the **Reports** menu. The **Media Report** screen appears.

Figure 10 Media Report Screen



- 2 Select the media on which to run the report by doing one of the following:
 - In the **Enter Media ID** field, type the ID of the media on which to run the report. Proceed to [Step 5](#) on page 287.
 - Select media IDs from the Selected Media ID(s) list. (You can click **Select All** to select all media IDs, or **Deselect All** to deselected selected IDs.) Proceed to [Step 5](#) on page 287.
 - Click **Browse** to select media from specific libraries
- 3 When you click **Browse**, the **Media Browser** screen appears. Select from the **View by:** list a library or media class.

Figure 11 Media Browser screen



The selected library or media class appears in the **Current Selection** field, and the media associated with the selected library are shown in the **Select Media** list.

- 4 In the **Select Media** list, select one or more media, and then click **OK**. The **Media Report** screen shows the selected media in the **Entered Media IDs** list.

Note: If you want to limit the media list to a specific selection, type the attributes of the media name in the **Media Filter** field to only browse these items. For example, typing *8 lists all media ending with an 8.

- 5 Click **Apply** to continue. The **Media Information Report** appears.

Figure 12 Media Information Report

The screenshot shows a web browser window titled "Media Information Report - Microsoft Internet Explorer". The page header includes the Quantum logo and the title "Media Information Report".

Media ID: MED001 [Show Details](#)

Media Type	Media Class	Policy Class	Last Access Time	Media Status	Write Protect	Formatted
LTO	F0_LTO_REMOVE	None	17-Jan-2007 14:07:18	None	N	N
Import Date	Export	# Files	Space Used	% Used	Space Remaining (bytes)	Mount Count
17-Jan-2007 14:07:18	UNMARKED	0	0	0.0	0	0
Move Count	Suspect Count	Current Action	Location	Current Archive	Pending Archive	
0	0	Export	Archive	vault1	None	

Media ID: MED002 [Show Details](#)

Media Type	Media Class	Policy Class	Last Access Time	Media Status	Write Protect	Formatted
LTO	F0_LTO_DATA	system blank	17-Jan-2007 14:07:18	AVAIL	N	N
Import Date	Export	# Files	Space Used	% Used	Space Remaining (bytes)	Mount Count
17-Jan-2007 14:07:18	UNMARKED	0	0	0.00	0	0
Move Count	Suspect Count	Current Action	Location	Current Archive	Pending Archive	
0	0	Move	Archive	vault1	scsi_archive1	

Close

- If desired, click the Show Details link to view detailed media information. The **Detailed Media Information Report** appears.

Figure 13 Detailed Media Information Report

The screenshot shows a web browser window titled "Media Information Report - Microsoft Internet Explorer". The page header includes the Quantum logo and the title "Media Information Report".

Media ID: MED003

Dead Space	0.0
------------	-----

No Files On Media MED003

Close

- 7 Click **Close** when you are finished viewing the report.

The Policy Class Information Report

The Policy Class Information Report provides the following information:

- **File Copy (1-4):** The number of copies stored for each file. A media type for File Copy "N" indicates the copy is to be stored; a value of N/A for the media type indicates this copy is not to be stored. File Copy 1 (the primary copy) will always have an associated media, and will be stored. The maximum number of file copies is four.
- **# Media Associated:** The number of media associated with the class
- **Drive Pool:** The name assigned to the pool of associated tape drives
- **Minimum Store Time (minutes):** The number of minutes after the last modification when the file becomes available for storage to tape
- **Minimum Trunc Time (days):** The number of days after the last modification when the files on tape become available for truncation
- **Max Backup Sets:** The maximum number of backup sets to keep for a file
- **Truncate Immediately After Store:** Indicates whether files truncate immediately after a store
- **Checksum Validation:** If enabled, checksums are compared to retained values for the files retrieved by the corresponding policy class
- **Checksum Generation:** If enabled, checksums are generated and retained in the database for files stored by the corresponding policy class
- **Minimum Set Store Size (1 to 999 in MB or GB):** The minimum size that all valid store candidates in the policy class combined must reach before they are stored

- **Maximum File Store Age (1 to 720 in hours):** If any valid store candidate in the policy class reaches this value, all valid candidates are stored
- **Disk-to-Disk:** Indicates whether the disk-to-disk functionality is enabled or disabled for the selected policy class
 - **Affinity From:** If disk-to-disk is enabled, the name of the primary affinity where a file resides
 - **Affinity To:** If disk-to-disk is enabled, the name of the secondary affinity to which the file relocates
 - **File Age Before Relocation:** The age a file must reach before it becomes eligible for relocation
- **Media Clean Pool:** The class name with which the media is associated after it is logically blank
- **Stub Files:** Indicates whether the Stub File feature is enabled or disabled for the storage policy
- **Stub File Size (KB):** If the Stub File feature is enabled, this is the target stub file size specified when the storage policy was created or modified
- **AutoStore:** Indicates whether the Autostore feature is enabled or disabled. The Autostore option automatically stores files for the current policy class. If this feature is disabled, Quantum recommends that the files for the policy class be stored by scheduled events (see below).
- **Retrieve to Affinity:** If you enabled the Retrieve to Affinity feature when you created or modified the policy class, the affinity to which files are retrieved is shown.

The **Schedules** section of the report provides the following information about the schedule associated with the Policy Class:


- **Name:** Name of the associated schedule
- **Run Days:** The days on which the schedule is set to run
- **Start Time:** The time when the schedule is set to begin running
- **Start Window:** The timeframe within which the scheduler attempts to begin the process. (For example, 30 minutes.) If the process cannot begin, it tries again during the next cycle.
- **Last Run:** The date and time the schedule was last run

- **Last Run Status:** The last status of the previous schedule
- **Media IDs:** If you select the Show Media report option, this section of the report shows the media associated with the policy class on which the report was run

Use the following procedure to run the Policy Classes report.


- 1 Choose **Policy Classes** from the **Reports** menu. The **Policy Classes Report** screen appears.

Figure 14 Policy Classes Report Screen



- 2 Select from the **Select Policy Class** list one or more policy classes on which to run the report.
- 3 To also view media associated with these policy classes, select the **Show Media** option.
- 4 Click **Apply**. The **Policy Class Information Report** appears.

Figure 15 Policy Class Information Report



File Copy 1	File Copy 2	File Copy 3	File Copy 4	# Media Associated	Drive Pool
LTO (default)	N/A	N/A	N/A	0	fs_F0drivepool
Minimum Store Time (minutes)	Minimum Trunc Time (days)	Max Backup Sets	Truncate File Immediately After Store	Checksum Validation	Checksum Generation
10	5	10	Off	DISABLED	DISABLED
Minimum Set Store Size (MB)	Maximum File Store Age (hours)	Disk-to-Disk	Affinity From	Affinity To	File Age Before Relocation
n/a	n/a	DISABLED	N/A	N/A	N/A
Media Clean Pool	Stub Files	Stub File Size(KBytes)	AutoStore	Retrieve to Affinity	
SYSTEM	Disabled	0	yes	a1	

X Close

5 Click **Close** when you are finished viewing the report.


The Relation Information Report

The Relation Information Report shows the pathname of the selected directory and the corresponding policy class name for the directory.

Use the following procedure to run the Directory/Policy Class Relationship report.


- 1 Choose **Relations** from the **Reports** menu. The **Directory/Policy Class Relationships Report** screen appears.

Figure 16 Directory/Policy Class Relationships Report Screen



- 2 Select from the **Select Managed Directories** list the directory on which the report is run, and then click **Apply**. The **Relation Information Report** appears.

Figure 17 Relation Information Report




3 Click **Close** when you are finished viewing the report.

The Request Information Report

Use the following procedure to run the Requests report.

- 1 Choose **Requests** from the **Reports** menu. The **Requests Report** screen appears.

Figure 18 Request Report
Screen



- 2 Select from the **Request ID** list one or more request IDs on which to run the report, and then click **Apply**. The **Request Information Report** appears.
- 3 Click **Close** when you are finished viewing the report.

The Scheduler Information Report

The Scheduler Information Report provides the following information:


- **Name:** The name of the schedule
- **Run Days:** The days on which the feature runs
- **Start Time:** The time when the feature runs

- **Window:** If StorNext was down at the time of the scheduled event and is restarted during the window time period, the event is launched; otherwise it is skipped until the next scheduled time
- **Last Run:** The time the schedule was last run

Use the following procedure to run the Scheduler report.


- 1 Choose **Scheduler** from the **Reports** menu. The **Scheduler Report** screen appears.

Figure 19 Scheduler Report Screen



- 2 Select one or more of the schedules on which to report, and then click **Apply**. The **Scheduler Information Report** appears.

Figure 20 Scheduler Information Report



Quantum. Scheduler Information Report

Scheduler Feature: Full Backup

Name	Run Days	Start Time	Window	Last Run	Last Run Status
f_backup_default	sun	22:00	3 hours	Jan-14-2007 22:00:20	Successful

Scheduler Feature: Partial Backup

Name	Run Days	Start Time	Window	Last Run	Last Run Status
p_backup_default	mon,tue,wed,thu,fri,sat	23:00	3 hours	Jan-15-2007 23:00:21	Successful

Scheduler Feature: Rebuild Policy

Name	Run Days	Start Time	Window	Last Run	Last Run Status
rebuild_default	sat	01:05	3 hours	Jan-13-2007 01:05:00	Successful

X Close

3 Click **Close** when you are finished viewing the report.

The Storage Disk Information Report

The Storage Disk Information Report provides the following information.

- **Name:** The name of the storage disk for which the report was run
- **Mount Point:** The storage disk's mount point of the storage disk
- **Copy #:** The copy ID for media on the storage disk
- **# of Streams:** The number of streams that can simultaneously access the storage disk
- **Status:** Shows if the storage disk is currently available or offline

- **Write Protected:** Indicates whether the storage disk is write protected
- **Last Access Time:** The date and time the storage disk was last accessed
- **# Files:** The number of files on the storage disk
- **Space Used:** The amount of space consumed on the storage disk
- **Space Remain:** The amount of space remaining on the storage disk
- **State:** The storage disk's current state (e.g., ONLINE, OFFLINE, or ONLINE-PENDING.)
- **Deduplication:** Indicates whether deduplication is enabled for the storage disk
- **% Savings:** If deduplication is enabled, this field reports the percentage of space saved by deduplication
- **Total Data Stored (MB):** If deduplication is enabled, this field shows total amount of data stored, in megabytes. (This field does not appear when deduplication is disabled.)
- **Unique Data Stored (MB):** If deduplication is enabled, this field shows amount of unique data stored, in megabytes. (This field does not appear when deduplication is disabled.)

Use the following procedure to run the Storage Disks report.

- 1 Choose **Storage Disks** from the **Reports** menu. The **Storage Disk Report** screen appears.


Figure 21 Storage Disk Report Screen



- 2 Select the disks on which to run the report and click **Apply**. The **Storage Disk Information Report** appears.

Note: If you enable the **Show Files on Storage Disk** option, the report could be very large and take a long time to run, depending on the number of files on the storage disk.

Figure 22 Storage Disk Information Report



Storage Disk ID: Storage_Disk_1

Name	Mount Point	Copy #	# of streams	Status
Storage_Disk_1	/stornext/snfs1:DDISK/adic.spock	1	2	AVAIL
Write Protected	Last Access Time	# Files	Space Used (GB)	Space Remain (GB)
N	Mar-8-2007 08:27:40	0	4.31	185.35
State	Deduplication	% Savings	Total Data Stored (MB)	Unique Data Stored (MB)
FREE	Enabled	0.00	0.00	0.00

Close


3 Click **Close** when you are finished viewing the report.

The Directory Affinity Report

Use the following procedure to run the Affinities report.

- 1 Choose **SNFS > Affinities** from the **Reports** menu. The **Affinities Report** screen appears.


Figure 23 Affinities Report
Screen



2 Do one of the following:


- In the **Select Individual Directory** text box, type the full path of the directory on which the to run the report. Proceed to Step 4.
- Click **Browse** to select a directory. The **Directory Browser** screen appears.

Figure 24 Directory Browser Screen



- 3 On the **Directory Browser** screen, select the directories on which to run the report.
- 4 Click **Apply**. The **Directory Affinity Information** report appears.

Figure 25 Directory Affinity Report



- 5 Click **Close** when you are finished viewing the report.

The File System Statistics Report


The File System Statistics Report provides the following information:

- **Creation Date:** The date and time when the file system was created
- **# Active SAN Clients:** The number of SAN client machines associated with the file system
- **File System Block Size:** The file system's block size
- **Message Buffer Size:** The size of the file system message buffer
- **# Disk Devices:** The number of disk devices on the file system
- **# Stripe Groups:** The number of stripe groups on the file system
- **Total Space:** The file system's total size
- **Available Space:** The amount of space still available on the file system
- **Managed:** Indicates whether the file system is managed (Yes) or unmanaged (No)

Use the following procedure to run the File System report.


- 1 Choose **SNFS > Systems** from the **Reports** menu. The **File System Report** screen appears.

Figure 26 File System Report Screen



- 2 In the **Active File Systems** list, select one or more active file systems on which to run the report, and then click **Apply**. The **File System Statistics Report** appears, showing statistical data for the selected file systems.

Figure 27 File System Statistics Report



- 3 Click **Close** when you are finished viewing the report.

The Stripe Group Statistics Report


The Stripe Group report provides the following information:

- **Total Space:** The total space available in the stripe group
- **Available Space:** The available space that has not been used in the stripe group
- **Breadth (blocks):** The number of contiguous blocks in the stripe group
- **Affinity:** Indicates whether an affinity is associated with the stripe group
- **Status:** Indicates whether the stripe group is currently up or down, and whether the data is Metadata, Journal, or Exclusive
- **Read:** Indicates whether the stripe group is read-enabled
- **Write:** Indicates whether the stripe group is write-enabled
- **Read Method:** Indicates whether the read method is assigned to the stripe group
- **Disks in Group:** The list of disks assigned to the stripe group

Use the following procedure to run the Stripe Groups report.

- 1 Choose **SNFS > Stripe Groups** from the **Reports** menu. The **Stripe Groups Report** screen appears.

Figure 28 Stripe Groups
Report Screen



- 2 Select from the **Active File Systems** menu the file system containing the stripe group for which to create the report.
- 3 In the **Stripe Groups** list, select one or more stripe groups on which to run the report, and then click **Apply**. The **Stripe Group Statistics Report** appears.

Figure 29 Stripe Group Statistics Report

Quantum Stripe Group Statistics Report									
File System: snfs2 Stripe Group: StripeGroup1									
Total Space	Available Space	Breadth (blocks)	Affinity	Status	Read	Write	Read Method	Disks in Group	
4.05 GB	3.99 GB (98%)	4	None	Up MetaData Journal Exclusive	Enabled	Enabled	None	spodsmall	
File System: snfs2 Stripe Group: StripeGroup2									
Total Space	Available Space	Breadth (blocks)	Affinity	Status	Read	Write	Read Method	Disks in Group	
189.66 GB	183.17 GB (96%)	4	a1	Up Exclusive	Enabled	Disabled	None	spod01	
File System: snfs2 Stripe Group: StripeGroup3									
Total Space	Available Space	Breadth (blocks)	Affinity	Status	Read	Write	Read Method	Disks in Group	
189.66 GB	185.53 GB (97%)	4	None	Up	Enabled	Enabled	None	spod03	

4 Click Close when you are finished viewing the report.

The File System Client Report

The File System Statistics Report provides the following information:


- **File System:** The name of the file system supporting the clients.
- **Mounted on:** The name of the file system mount point.
- **# SAN Clients:** The total number of physically connected StorNext SAN clients.
- **# LAN Servers:** The total number of distributed LAN servers for the indicated file system.
- **# LAN Clients:** The total number of StorNext distributed LAN clients.
- **SAN Clients:** The names of physically connected SAN clients.
- **Distributed LAN Servers:** The names of the distributed LAN servers.
- **Distributed LAN Clients:** The names of distributed LAN clients.
- **LAN Servers:** The name of the distributed LAN server for which the subsequent details apply.

- **Listening IP:Port:** The IP address and port number through which the distributed LAN server communicates with StorNext.
- **TCP Window Size:** The TCP window size (in KB) used by the distributed LAN server. (Default: 64)
- **Transfer Buffer Size:** The transfer buffer size (in KB) used by the distributed LAN server. A larger buffer may increase performance for larger files. (Default: 256)
- **Transfer Buffer Count:** The number of transfer buffers used by the distributed LAN server. This parameter is used only by Windows servers and clients. Linux servers pass the value of this parameter to Windows clients. (Default: 16)
- **Daemon Threads:** The maximum number of daemon threads used by the distributed LAN server. (Default: 8)
- **Luns:** The disk name; physical device name; number of sectors; and sector size.

Use the following procedure to run the File System Client Statistics Report.

- 1 Choose **SNFS > Clients** from the **Reports** menu. The **File System Client Report** screen appears.

Figure 30 File System Client Report Screen



- 2 Select from the **Active File Systems** list one or more file systems to include in the report. Click **Apply** to continue. The **File System Client Report** appears.

Figure 31 File System Client Report

# SAN Clients		# LAN Servers		# LAN Clients	
5 of 15		4		2 of 15	

SAN Clients	Distributed LAN Servers	Distributed LAN Clients
ha-stormext2.adic.com	stormext4.adic.com	paradise.adic.com
stormext4.adic.com	sooop.adic.com	Dubnium.adic.com
sooop.adic.com	stormext8.adic.com	
stormext8.adic.com	stormext3.adic.com	
stormext3.adic.com		

LAN Servers	Listening IP:Port	TCP Window Size	Transfer Buffer Size	Transfer Buffer Count	Daemon Threads
stormext3.adic.com	10.16.50.121:32779	64K	256K	16	8

Luns

ha-stormext1 on device:sdbf(A),sdo(S) sectors: 622823584 sector size: 512

LAN Servers	Listening IP:Port	TCP Window Size	Transfer Buffer Size	Transfer Buffer Count	Daemon Threads
sooop.adic.com	172.16.99.240:32781	64K	256K	16	8

Luns

ha-stormext1 on device:sdbg(A),sdp(S) sectors: 622823584 sector size: 512

LAN Servers	Listening IP:Port	TCP Window Size	Transfer Buffer Size	Transfer Buffer Count	Daemon Threads
stormext8.adic.com	10.16.50.14:47181	64K	256K	16	8

Luns

ha-stormext1 on device:sdbe(A),sdr(S) sectors: 622823584 sector size: 512

X Close

3 Click Close when you are finished viewing the report.

The File System LAN Client Report

The File System LAN Client Report provides the following information:


- **File System Name:** The name of the file system that supports the indicated distributed LAN server.
- **Distributed LAN Server Name:** The name of the distributed LAN server on the indicated file system.

- **Distributed LAN Client Name:** The name of the distributed LAN client for the indicated file system and distributed LAN server.
- **Read:** The speed at which the distributed LAN client is currently reading data.
- **Write:** The speed at which the distributed LAN client is currently writing data.

Use the following procedure to run the Distributed LAN Client Performance Report.

- 1 Choose **SNFS > LAN Client Performance** from the **Reports** menu. The **File System Distributed LAN Client Statistics Report** screen appears.

Figure 32 File System Distributed LAN Client Statistics Report Screen



- 2 Select from the **Active File Systems** list one or more file systems to include in the report. Click **Apply** to continue. The **File System LAN Client Report** appears.

Figure 33 File System LAN Client Report

Quantum File System LAN Client Report

LAN Client Monitor Auto-Refresh Rate

File System: snfs1 LAN Server: scoop.adic.com

LAN Client	Read	Write
Dubnium.adic.com:172.16.99.80:32793		
paradise.adic.com:172.16.99.90:43312		
stornext4.adic.com:172.16.99.122:33883		
Dubnium.adic.com:172.16.99.80:32784		
paradise.adic.com:172.16.99.90:43308		
stornext4.adic.com:172.16.99.122:33881		


File System: snfs1 LAN Server: stornext3.adic.com

LAN Client	Read	Write
Dubnium.adic.com:172.16.99.80:33049		
Dubnium.adic.com:172.16.99.80:43528		
paradise.adic.com:172.16.99.90:56074		
stornext4.adic.com:172.16.99.90:2290		
Dubnium.adic.com:172.16.99.122:34015		
paradise.adic.com:10.16.50.122:34019		
stornext4.adic.com:172.16.99.80:33052		
:172.16.99.90:56072		
:172.16.99.122:34013		
:10.16.50.122:34017		

3 If desired, click **Refresh** to manually update (refresh) the report data. You can also use the **Auto Refresh Rate** field to specify one of these automatic refresh intervals:

- No Refresh
- 5 seconds
- 10 seconds
- 30 seconds
- 1 minute
- 5 minutes

4 Click **Close** when you are finished viewing the report.



Chapter 13

Service Management

This chapter describes how to use the StorNext Service Management tools to run a health check on your system, create a log that captures the current state of your system, or check current system status. This chapter contains these topics:

- [Using Health Check](#)
- [Using State Capture](#)
- [Using the System Status Tool](#)

Using Health Check

You can run these health checks on your StorNext system:

- **Archive:** Verify that all configured archives are online
- **Config:** Verify that affinities are configured correctly in SNSM for managed file systems, and that SNSM-managed file systems are identified and configured correctly
- **Dedup SDISK:** Verifies blockpool data integrity for all blocketized storage disks (i.e., deduplication-enabled storage disks).
- **Disk Space:** Verify that enough disk space exists for the SNSM database tables, logging, and other functions

- **Drive:** Verify that all configured drives are online
- **Media:** Verify that there are enough media available for all policies to store all file copies, and that SNSM media are configured correctly
- **Policies:** Verify that SNSM is keeping up with file system events and store candidate processing

Running a Health Check

Use the following procedure to run a health check. All of the health options are linked to specific information about that particular health check.

- 1 From the StorNext home page, choose **Health Check** from the **Service** menu. The **Health Check Tests** screen appears.

Figure 1 Health Check Tests Screen




- 2 Select one or more health checks to run, and then click **Run**.
- 3 When the Status screen informs you that the check have run successfully, click **Close**.

Viewing the Health Check History

After running a health check, you can view a five-run history of each health check that has been run.


- 1 Select a health check from the **Health Check Tests** screen.
- 2 Click **History**. The **Health Check History** screen appears.

Figure 2 Health Check History Screen



- 3 Do one of the following:
 - Select a specific report to view and click **Details**.
 - Click the **Pass** or **Fail** link in the **Results** column. Proceed to [Viewing the Health Check Results](#) on page 14.
- 4 When you click **Details** on the **Health Check History** screen, the **Health Check Results** screen appears.

Figure 3 Health Check Results
Screen



- 5 Click **Back** to view more archives, or close the window when finished.

Viewing the Health Check Results

When viewing the Health Check Tests or the History screen, you can check the details of the Last Status of a particular health check. The following example shows how to run the last status check from the **Health Check Tests** screen.

- 1 From the StorNext home page, choose **Health Check** from the **Service** menu. The **Health Check Tests** screen ([Figure 233](#)) appears.
- 2 Click the **Last Status** link (Pass or Fail). The **Health Check Results** screen appears, showing the results of the last health check run. This screen is the same as the Results screen that appears when you click **Details** on the **Health Check History** screen. (See [Figure 235](#))
- 3 After you are finished viewing the results, click **Close**.

Using State Capture

The StorNext State Capture tool enables you to create a log that captures the current state of your system. This log assists Quantum support personnel analyze and debug some problems in the storage system.

When you use the State Capture tool, StorNext creates a log file named using the format *snapshot-machinehostname-YYYYMMDDHHMMSS.tar.gz*. This file contains a summary report that is produced by executing the `pse_snapshot` command on all component config/filelist files.


If desired, you can download or delete a previously captured file. When you are finished using the State Capture tool, you can return to the StorNext home page or select another tool or option.

Capturing the Current System State

Use the following procedure to access the StorNext State Capture tool.

- 1 From the StorNext home page, choose **Capture State** from the **Service** menu. The **Capture System State** screen appears. Any previously captured snapshots are shown.

Figure 4 Capture System State
Screen




- 2 Click **Capture**. The **Capture State Status** window is shown.
- 3 When the Status window informs you that the capture was successful, click **Close**.

Downloading a Previous System Capture

After you have created at least one system capture, you can select and download one of those .tar.gz files to view.

- 1 From the StorNext home page, choose **State Capture** from the **Service** menu. A list of capture files stored in the directory `/usr/adc/www/logs/capture_state` is shown. (This directory is where the files are stored on the StorNext server.)
- 2 Locate the capture file you want to download, and then click the corresponding radio button under the **Number** column beside the filename.
- 3 Click **Download**. The **Download Capture File** screen appears.

Figure 5 Download Capture File Screen



- 4 If the download does not start automatically, click the supplied link.
- 5 Specify whether you want to open or save the capture file. (The file is in compressed tar.gz format, so in most cases you will want to save the file and then open it with a file decompression utility such as WinZip.)

Deleting a Previous System Capture

When you are finished viewing and analyzing a capture file, you can delete an unwanted file.

- 1 From the StorNext home page, choose **Capture State** from the **Service** menu. The **Capture System State** screen shows a list of capture files stored in the directory `/usr/adic/www/logs/capture_state`.
- 2 Locate the capture file you want to delete, and then click the radio button beside the filename.
- 3 Click **Delete**.
- 4 When a confirmation screen prompts you to confirm that you want to delete the file, click **OK** to continue.
- 5 After the status screen informs you that the file was successfully deleted, click **Close**.

Using the System Status Tool

The System Status tool creates a list of RAS tickets that relate to system faults or errors. Ticket details provide a summary of the system fault, an area for Analysis notes, and contains a Recommended Actions link to help you correct the fault.

Use the following procedure to use the System Status tool.

- 1 From the StorNext home page, choose **System Status** from the **Service** menu. The **Service - System Status** screen appears.

Figure 6 Service - System Status Screen

Quantum StorNext Home Help

Config Admin Reports Service Help

Service - System Status

Home Below is a list of tickets indicating faults reported by the system.
To close all open tickets, use "Close All" at the bottom of the page.

SNFS

SNSM

Total Number of Tickets: 20

Ticket	State	Priority	Last Update	Summary
1	closed	high	20-dec-2006 23:00:01	ADIC software
2	closed	high	21-dec-2006 23:00:00	ADIC software
3	closed	high	22-dec-2006 23:00:00	ADIC software
4	closed	high	23-dec-2006 23:00:01	ADIC software
5	closed	high	24-dec-2006 22:00:00	ADIC software
6	closed	high	25-dec-2006 23:00:00	ADIC software
7	closed	high	26-dec-2006 23:00:00	ADIC software
8	closed	high	27-dec-2006 23:00:00	ADIC software
9	closed	high	28-dec-2006 23:00:00	ADIC software
10	closed	high	29-dec-2006 23:00:00	ADIC software
11	closed	high	30-dec-2006 23:00:00	ADIC software

Show Tickets: Open Closed All

Previous Page 1 of 2 Next

Details Close All Refresh

StorNext kazar Active


Basic Administration Page - Allows the user to operate the product from one location Local intranet

The **Service - System Status** screen contains the following information:

- **Total Number of Tickets:** The number of RAS tickets the system has generated
- **Ticket:** The RAS ticket number, displayed in the order in which it was created
- **State:** The ticket's current status: OPEN or CLOSED

- **Priority:** The ticket's priority based on system impact: HIGH, MEDIUM, or LOW
 - **Last Update:** The date of the last system status update
 - **Summary:** A short summary of the fault that triggered creating the RAS ticket
 - **Show Tickets:** Controls the type of tickets shown in the display window: OPEN, CLOSED, or ALL tickets
 - **Previous** and **Next:** Click these buttons to toggle between ticket pages (if there is more than one page of tickets)
 - **Details:** Click this button to view a selected ticket's details
 - **Close All:** Click this button to close all tickets shown in the display window
 - **Refresh:** Click this button to refresh the display window
- 2 Highlight the ticket you wish to view, and then click **Details**. The **RAS Ticket Details** screen appears.

Figure 7 RAS Ticket Details
Screen




The **RAS Ticket Details** screen provides the following information:


- **Ticket Number:** The number of the ticket in the displayed ticket list
- **Opened:** The date and time the ticket was created
- **Status:** The current status of the ticket: OPEN or CLOSED
- **Priority:** The ticket's priority based on system impact: HIGH, MEDIUM, and LOW

- **Summaries and Descriptions:** Detailed information about the ticket, including a Recommended Actions link to help you correct the fault or condition
 - **Analysis:** If desired, enter information about the fault or condition, such as a recommended action
 - **Apply Analysis:** Click this button to save information entered in the **Analysis** field
 - **Close Ticket:** Click this button after you have corrected the condition or fault
 - **Cancel:** Click this button to close the **RAS Ticket Details** screen
- 3 When you click the **View Recommended Actions** link on the **RAS Ticket Details** screen, the **Recommended Actions** screen appears. This screen provides information and steps to correct the condition or fault that generated the RAS ticket. Follow the instructions on the **Recommended Actions** screen to correct the condition or fault.

Figure 8 Recommended
Actions Screen



- 4 Click the **Close Window** link at the bottom of the **Recommended Actions** screen.
- 5 Click the **Close** or **Cancel** button on the **RAS Ticket Details** screen.



Chapter 14

Customer Assistance


More information about this product is available on the Quantum Service and Support website at www.quantum.com/ServiceandSupport. The Quantum Service and Support website contains a collection of information, including answers to frequently asked questions (FAQs). You can also access software, firmware, and drivers through this site.

Quantum Technical Assistance Center

For further assistance, or if training is desired, contact the Quantum Technical Assistance Center:

North America	1+800-284-5101 Option 5
EMEA	00800 999 3822
Online Service and Support	www.quantum.com/OSR
Worldwide Web	www.quantum.com/ServiceandSupport

(Local numbers for specific countries are listed on the Quantum Service and Support Website.)



Appendix A

Operating Guidelines

This appendix contains information pertinent to operating StorNext, as well as some operating guidelines and limitations you should consider.

The Reserved Space Parameter

As of StorNext 3.0, the method of accounting for reserved space has changed. The `MaxMBPerClientReserve` parameter from the StorNext file system configuration file (`/usr/cvfs/config/*.cfg`) has been deprecated. All values except 0 are ignored for this parameter. In addition, there is a new parameter, `ReservedSpace`.

The `ReservedSpace` parameter lets the administrator control the use of delayed allocations on clients. `ReservedSpace` is a performance feature that lets clients perform buffered writes on a file without first obtaining real allocations from the metadata controller (MDC).

The `ReservedSpace` parameter can be set to Yes or No:

- Yes - (Default) The MDC reserves enough disk space so that delayed allocations can be converted to real allocations (even when the MDC is restarted and the client is not). The MDC reserves a minimum of about 4GB for each stripe group and up to 280MBs per actively writing client **for each stripe group**.

Note: The amount of reserved space is usually less than 280MB per client. Reserved space is calculated as 110% of the buffer cache size of each particular client. For example, a client with a 64MB buffer cache is allocated 70MBs of reserved space by the MDC. If the client closes all files that are open for write, the 70MBs of space is no longer accounted for. **It is important to remember that reserved space is per stripe group.**

- No - More disk space is available for use, but buffer cache performance is affected, and fragmentation may occur.

If the MaxMBPerClientReserve parameter exists in the configuration file and has a value of 0, ReservedSpace is set to No. Otherwise, ReservedSpace defaults to Yes.

Note: In prior releases of StorNext, MaxMBPerClientReserve defaulted to 100MBs, and reserved space was the product of MaxMBPerClientReserve multiplied by MaxConnections – 1. In StorNext 3.0, the MDC tracks the actual amount of reserved space that clients use but caps the value to about 280MBs per client.

In addition, the MDC maintains a “minimum level” of reserved space. As a result, in some cases, more reserved space may be visible. Reserved space appears as allocated disk space per data stripe group.

The minimum reserved space is enough to accommodate 15 clients or MaxConnections – 1, whichever is lower. For example, if a cluster has a MaxConnections of 3, the reserved space total can be under 1GB.

Distributed LAN Server/Client Network and Memory Tuning

Using the Distributed LAN Server and Client feature places significant additional demands on network capacity and system memory. Before

creating and using a Distributed LAN Server and Client, review the following information:

- [Distributed LAN Server and Client Network Tuning](#)
- [Distributed LAN Server Memory Tuning](#)

Note: For additional information about Distributed LAN Client and server performance tuning, see the *StorNext File System Tuning Guide*.

Distributed LAN Server and Client Network Tuning

Due to significant demands placed on the network, the following network issues can occur when using Distributed LAN Servers and clients:

- **Configuring Dual NICs.** If the Distributed LAN Server has two network interface cards (NICs), each card must have a different address and reside on a different subnet. In addition, to take advantage of a second NIC in a Distributed LAN Server, the Distributed LAN Clients must also have a second connected network interface.
- **Dropped Packets.** Some Ethernet switches may be unable to accommodate the increased throughput demands required by the Distributed LAN Server and client feature, and will drop packets. This causes TCP retransmissions, resulting in a significant performance loss. On Linux, this can be observed as an increase in the Segments Retransmitted count in `netstat -s` output during Distributed LAN Client write operations and Distributed LAN Server read operations.

To address this issue, edit the `/usr/cvfs/config/dpserver` configuration file and reduce the Distributed LAN Server TCP window size from the default value. (Remount the file system after making changes.) This may reduce the amount of packet loss. However, some Ethernet switches are unable to accommodate true GigE bandwidth, especially when multiple ports are transmitting data at the same time.

- **Linux Network Drivers.** For best performance and compatibility, update Intel e1000 drivers to the latest version.

In some cases, enabling TCP offload can cause issues. (Identify these issues by examining `netstat -s` output for bad segments.) If necessary, use `ethtool -K` to disable the offload of checksum calculations.

On some Linux 2.6 versions running on x86 64-bit systems, a console message regarding noirq handler may appear followed by a hard system hang. This is due to a bug in the kernel. To avoid this error, disable the irqbalance service.

- **Mismatched Server Configuration.** Introducing a slower server onto the network reduces overall throughput. This is because the slower server receives some traffic from all clients. For example, adding a server with one NIC in a network where other servers have two NICs, or adding a server with less disk bandwidth or a bad network connection, reduces throughput for the entire network.

Note: On Linux, use ping and the cvadmin latency test tools to identify network connectivity or reliability problems. In addition, use the netperf tool to identify bandwidth limitations or problems.

On Windows, use the **Networking** tab of **Windows Task Manager** to view network utilization.

Distributed LAN Server Memory Tuning

The minimum amount of memory required for a Distributed LAN Server depends on the configuration.

- **Windows.** For a Windows Distributed LAN Server, use the following formula:

Required memory = 1GB +
(# of file systems served
* # of NICs per Distributed LAN Client
* # of Distributed LAN Clients
* transfer buffer count
* transfer buffer size)

For example, suppose a Windows Distributed LAN Server is serving four file systems to 64 clients each using two NICs for data traffic. Also assume the server uses the defaults of sixteen transfer buffers and 256K per buffer. (On Windows, you can view and adjust the transfer buffer settings using the Client Configuration tool's Distributed LAN tab.) Given this configuration, here is the result:

Required memory = 1GB + (4 * 2 * 64 * 16 * 256K) = 3GB

Note: This example assumes that a 64-bit version of Windows is being used on the Server. 32-bit Windows Distributed LAN Servers are restricted to small configurations using 16 or fewer connections.

If not all clients mount all of the file systems, the memory requirement is reduced accordingly. For example, suppose in the previous example that half of the 64 LAN clients mount three of the four file systems, and the other half of the LAN clients mount the remaining file system. Given this configuration, here is the result:

Required memory = 1GB + (3 * 2 * 32 * 16 * 256K) + (1 * 2 * 32 * 16 * 256K) = 1GB + 768MB + 256MB = 2GB

The calculation also changes when the number of NICs used for data traffic varies across clients. For example, in the previous example if the clients that mount only one file system each use three NICs for data instead of two, here is the result:

Required memory = 1GB + (3 * 2 * 32 * 16 * 256K) + (1 * 3 * 32 * 16 * 256K) = 1GB + 768MB + 384K = 2176MB

- **Linux.** For a Linux Distributed LAN Server, use the following formula:

Required memory = 1GB +

(# of file systems served

* # of NICs on the Distributed LAN Server used for
Distributed LAN traffic

* server buffer count

* transfer buffer size)

For example, consider a Linux Distributed LAN Server that has two NICs used for Distributed LAN traffic, serves four file systems, and uses the default eight server buffers and 256K per buffer. (See the `dpserver` and `sndpscfg` man pages for information about viewing and modifying Distributed LAN buffer settings on Linux.) For this case:

Required memory = 1GB + (4 * 2 * 8 * 256K) = 1040MB

Note: This example results in a memory requirement of less than 2GB. However, Quantum recommends that all Distributed LAN Servers contain a minimum of 2GB of RAM.

Configuring LDAP

This sections describes how to configure the StorNext LDAP functionality and describes related features in the Windows configuration utilities.

Using LDAP

StorNext 2.7 introduced support for Light Directory Access Protocol, or LDAP (RFC 2307). This feature allows customers to use Active Directory/LDAP for mapping Windows User IDs (SIDs) to UNIX User ID/Group IDs.

Changes to “Nobody” mapping

If a Windows user cannot be mapped to a UNIX ID, the user is mapped to Nobody. StorNext allows administrators to change the value of Nobody by using the file system configuration parameters:

UnixNobodyUidOnWindows 60001

UnixNobodyGidOnWindows 60001

These parameters are located in the file system configuration file on the server and can be manually modified by the Windows or StorNext Web GUI.

Note: Compatible Active Directory servers include Windows 2003 Server SP1 (with the Windows Services for Unix 3.5 extended LDAP schema applied,) and Windows 2003 Server R2.

UNIX File and Directory Modes

When a file or directory is created on Windows, the UNIX modes are controlled by the following file system configuration parameters:

UnixDirectoryCreationModeOnWindowsDefault 0755

UnixFileCreationModeOnWindowsDefault 0644

StorNext allows one set of values for all users of each file system.

Note: Administrators can manually change these values in the file system configuration file on the server or use the Windows or Web GUI.

LDAP Refresh Timeout

Due to the implementation of the Windows Active Directory user mappings, services for UNIX can take up to 10 minutes to be propagated to StorNext clients.

User ID Mapping Precedence

If multiple mappings are found for a given Windows user, the following precedence takes place:

- **NIS/PCNFSD** - If mapping exists.
- **Fabricated IDs** - If configured "on".
- **LDAP/RFC 2307** - If defined in Active Directory.
- **Nobody** - If no other mapping found.

Navigating to the **Authentication** tab of the **Client Configuration** Windows utility gives the administrator the ability to disable the NIS/PCNFSD mapping on a client-by-client basis by selecting **Use Active Directory**.

The domain server must be at least Windows 2003 or above to support RFC 2307 mapping.

Setting Up Restrictive ACLs

When setting up restrictive ACLs on a SNFS file system, it is important to understand how SNFS system services are run, especially the account under which the services are run. The Windows default account is the local administrator account, but this can be changed on the Properties tab of each system service.

When sharing restricted file systems, the account under which SNFS system services are run must be included in the ACL for the root of the

file system and all other shares associated with the SNFS file system. Doing this allows the shares to be re-shared upon reboot.

Default Single-Path I/O Retry Behavior

The I/O retry behavior has changed as of StorNext 3.1.2. In prior releases, when only a single path to the storage existed and an I/O error was returned by the disk device driver, StorNext failed the I/O operation. Beginning with version 3.1.2, by default StorNext continuously retries I/O operations until they succeed, regardless of the number of I/O paths. If desired, you can override this new behavior by using the new I/O Retry Time feature. For additional information about I/O Retry Time, consult the `mount_cvfs` man page or the Windows help file.

Event Handles for `fsm.exe` on a Windows Metadata Server

The metadata server (FSM) has many data structures that are used internally. Each of the data structures has some locks (`pthread_mutex_lock`). Each lock is initialized as “uninitialized.”

The first time the lock is used, a small amount of memory and an event (i.e., handle) are allocated. The memory and event/handle are retained by the system until the data structure is destroyed. Some locks that are part of structures are seldom used, and exist for rare conditions. If the lock is not used, the memory/event for that structure will never be allocated.

Some data structures are not destroyed during the life of the FSM. These include in-memory inodes and buffers and others.

When the system starts, handle use is minimal. After the FSM has been up for a while, the handle count increases as the inode and buffer cache are used. After a while, the system stabilizes at some number of handles. This occurs after all inodes and buffers have been used.

The maximum number of used handles can be reduced by shrinking the inode and/or buffer cache. However, changing these variables could significantly reduce system performance.

FsBlockSize, Metadata Disk Size, and JournalSize Settings

The FsBlockSize (FSB), metadata disk size, and JournalSize settings all work together. For example, the FsBlockSize must be set correctly in order for the metadata sizing to be correct. JournalSize is also dependent on the FsBlockSize.

For FsBlockSize the optimal settings for both performance and space utilization are in the range of 16K or 64K.

Settings greater than 64K are not recommended because performance will be adversely impacted due to inefficient metadata I/O operations. Values less than 16K are not recommended in most scenarios because startup and failover time may be adversely impacted. Setting FsBlockSize (FSB) to higher values is important for multi terabyte file systems for optimal startup and failover time.

Note: This is particularly true for slow CPU clock speed metadata servers such as Sparc. However, values greater than 16K can severely consume metadata space in cases where the file-to-directory ratio is low (e.g., less than 100 to 1).

For metadata disk size, you must have a *minimum* of 25 GB, with more space allocated depending on the number of files per directory and the size of your file system.

The following table shows suggested FsBlockSize (FSB) settings and metadata disk space based on the average number of files per directory and file system size. The amount of disk space listed for metadata is *in addition* to the 25 GB minimum amount. Use this table to determine the setting for your configuration.

Average No. of Files Per Directory	File System Size: Less Than 10TB	File System Size: 10TB or Larger
Less than 10	FSB: 16KB Metadata: 32 GB per 1M files	FSB: 64KB Metadata: 128 GB per 1M files
10-100	FSB: 16KB Metadata: 8 GB per 1M files	FSB: 64KB Metadata: 32 GB per 1M files
100-1000	FSB: 64KB Metadata: 8 GB per 1M files	FSB: 64KB Metadata: 8 GB per 1M files
1000 +	FSB: 64KB Metadata: 4 GB per 1M files	FSB: 64KB Metadata: 4 GB per 1M files

The best rule of thumb is to use a 16K FsBlockSize unless other requirements such as directory ratio dictate otherwise.

This setting is not adjustable after initial file system creation, so it is very important to give it careful consideration during initial configuration.

Example: FsBlockSize 16K

JournalSize Setting

The optimal settings for JournalSize are in the range between 16M and 64M, depending on the FsBlockSize. Avoid values greater than 64M due to potentially severe impacts on startup and failover times. Values at the higher end of the 16M-64M range may improve performance of metadata operations in some cases, although at the cost of slower startup and failover time.

The following table shows recommended settings. Choose the setting that corresponds to your configuration.

FsBlockSize	JournalSize
16KB	16MB

FsBlockSize	JournalSize
64KB	64MB

This setting is adjustable using the **cvupdatefs** utility. For more information, see the **cvupdatefs** man page.

Example: JournalSize 16M

Allowance for Special Characters

In prior StorNext releases, the StorNext GUI would sometimes operate incorrectly when encountering files whose names contained certain non-standard characters: single quotation marks and characters with ASCII values less than 32, such as tab stops and carriage returns.

Beginning with StorNext 3.1.2, the GUI has been modified to handle all known issues with file names containing these non-standard characters. Still, Quantum *strongly* recommends against deliberately using these characters in filenames.

If the GUI encounters a filename containing a “less-than-32” ASCII character, the character is rendered in the display as in the following examples:

- A horizontal tab is shown as “_0x09_”.
- A line feed is shown as “_0x0A_”.
- A carriage return is shown as “_0x0D_”.

The following additional characters should be avoided in “mixed” environments containing both Windows and non-Windows clients, because Windows will not be able process any file or directory with these characters in the name:

? \ < > " : |

Disk Naming Requirements

When naming disks, names should be unique across all SANs. If a client connects to more than one SAN, a conflict will arise if the client sees two disks with the same name.

General Operating Guidelines and Limitations

[Table 1](#) lists updated information and guidelines for running StorNext, as well as known limitations.

Table 1 Operating Guidelines
and Limitations

Operating System / Affected Component	Description
HA configurations	In HA (high availability) configurations, do not run <code>adic_control start</code> on the standby server. Doing so may start StorNext Storage Manager components on the standby server which are already running on the primary server, and may result in database corruption.

Operating System / Affected Component	Description
IRIX	<p>As of release 3.0, StorNext does not support metadata controllers running the SGI IRIX operating system. To upgrade to StorNext 3.0 or higher, MDCs running IRIX must be converted to a supported operating system. This conversion must be performed by Quantum Professional Services. For more information, or to schedule an IRIX conversion, contact Quantum Global Services (see Quantum Technical Assistance Center on page 11).</p> <p>NOTE: The StorNext 3.1.2 client software continues to support IRIX.</p>
	<p>SNFS uses 64-bit inode numbers. 32-bit applications (programs compiled without 64-bit) support may experience problems working with files with inode numbers greater than 2147483648. 32-bit applications may also encounter problems when dealing with files greater than 2TB in size, This issue is not unique to StorNext, but may be encountered by some legacy applications when working with SNFS on IRIX.</p>
	<p>StorNext supports only SGI 64-bit versions of IRIX on operating systems that use the SGI version of the QLogic QLA2200 or QLA2310 FC-HBAs.</p>
	<p>On many versions of IRIX, the root crontab contains the following entry which is used to remove old application crash dumps and temporary mail files:</p> <pre>find / -local -type f '(' -name core -o -name dead.letter ')' -atime +7 -mtime +7 -exec rm -f '{}';'</pre> <p>If StorNext file systems are mounted, they will be traversed by this find command which can have a dramatic impact on the performance of other applications currently using these file systems. To prevent the traversal of StorNext file systems, modify the find command so it reads:</p> <pre>find / -local '(' -type d -fstype cvfs -prune ')' -o -type f '(' -name core -o -name dead.letter ')' -atime +7 -mtime +7 -exec rm -f '{}';'</pre>

Operating System / Affected Component	Description
Linux	<p>StorNext does not support the use of loopback IP addresses other than 127.0.0.1. SUSE Enterprise Linux 10 systems sometimes have an entry for 127.0.0.2 in the <code>/etc/hosts</code> file. This causes various FSM operations to fail with the error "Not Local."</p>
	<p>Linux Distributed Gateway (Proxy Server) Restrictions: The <code>dpserver</code> file, used to configure network interfaces and addresses for Distributed Gateways, has a counter-intuitive restriction. If an interface has more than one address, an address must be specified in the <code>dpserver</code> file. However, if the interface has only one address, an address <i>cannot</i> be specified or the mount command will fail.</p>
	<p>Many versions of Linux run a cron script nightly to build a database used by the <code>slocate</code> command. If StorNext file systems are mounted, they are traversed by this cron job which can have a dramatic impact on the performance of other applications currently using these file systems. Perform the following steps (based on Linux version) to prevent the cron script from traversing StorNext file systems.</p> <p>RedHat Enterprise Linux 4 and 5 Add "cvfs" to the list of file system types to be skipped. This is usually done by modifying the "PRUNEFS" line in the <code>/etc/updatedb.conf</code> file to read: PRUNEFS="cvfs sysfs selinuxfs usbdevfs devpts NFS nfs nfs4 afs sfs proc smbfs cifs autofs auto iso9660 udf"</p> <p>SUSE Linux Enterprise Server 10 The optional "findutils-locate" package is used to build the <code>slocate</code> database. The default behavior is to disable building the database. If enabled, to prevent cvfs file systems from being scanned, add "cvfs" to the list of file system types to be skipped. This is usually done by modifying the "UPDATEDB_PRUNEFS" line in the <code>/etc/sysconfig/locate</code> file to read: UPDATEDB_PRUNEFS="cvfs"</p>
All UNIX and Linux	<p>The <code>swapon</code> command does not work on StorNext file systems. The Linux/Unix <code>swapon</code> command is used to specify devices on which paging and swapping take place. If <code>swapon</code> is run on a StorNext file system, the command fails with an invalid argument error.</p>

Operating System / Affected Component	Description
All UNIX and Linux	<p>In a file system configuration file, if the Metadata parameter for a stripe group was set to Yes when the file system was created, do not later change the parameter to No. Doing so will cause the metadata to become inaccessible and result in errors. Also, if the parameter is changed to No, the <code>cvupdatefs</code> command might allow bandwidth expansion of the stripe group, resulting in metadata corruption.</p>
Solaris	<p>On Solaris, by default, the <code>/usr/lib/fs/nfs/nfsfind</code> script is run nightly by the cron daemon. This script contains a <code>find</code> command that traverses any local file system that is exported (shared) via NFS. If the running of this script is interfering with the performance of processes accessing StorNext file systems, it can be modified to skip them. To do so, add the following line to the <code>find</code> command in the script:</p> <pre>'(' -type d -fstype cvfs -prune ')' -o</pre> <p>The final <code>find</code> command should look like this:</p> <pre>find \$dir '(' -type d -fstype cvfs -prune ')' -o -type f -name .nfs* -mtime +7 -mount -exec rm -f {} \;</pre> <p>The Solaris Security Toolkit, formally known as JASS, causes the following two issues:</p> <ul style="list-style-type: none"> • It disables RPC by renaming the RPC startup script, disrupting the StorNext interprocess communication. To fix the communication problem, rename the RPC startup script in <code>/etc/init.d</code> from <code>rpc.<illegal extension></code> to <code>rpc</code>. • It turns on IPSec, causing numerous warning messages in the system log file. Either disable IPSec by removing the IPSec startup file in <code>/etc/init.d</code> or contact Sun Technical Support to find out how to reconfigure IPSec to ignore local loopback connections.
Windows	<p>Windows Services for UNIX (SFU) supports only NTFS for NFS exports. Because of this limitation, a Windows system cannot act as an NFS server for StorNext File System.</p> <p>When a StorNext file system is mounted to a drive letter or a directory, configure the Windows backup utility to NOT include the StorNext file system.</p>

Operating System / Affected Component	Description
Windows	<p>On Windows systems, StorNext may log error messages after reboot if it tries to start before some other services. For example, if the network is not yet available, a StorNext client cannot contact a StorNext server and posts an error in the logs. (The client will continue to retry the connection until it is successful.)</p> <p>To prevent this from occurring, make the startup of StorNext services dependant on another Windows service (for example, the Workstation service). For more information, see the Microsoft Knowledge Base article Q193888, "How to delay loading of specific services."</p>
	<p>Windows-based SNFS clients do not support symbolic links that point to a file system outside of SNFS. Symbolic links are created either by a UNIX/Linux client or by a Windows client (and are referred to as junctions).</p>
	<p>Virus-checking software can severely degrade the performance of any file system, including SNFS. If you have anti-virus software running on a Windows Server 2003 or Windows machine, Quantum recommends configuring the software so that it does NOT check SNFS.</p>
	<p>For best performance and reliability, Quantum recommends running Distributed LAN Servers on 64-bit versions of Windows (or on Linux). Memory management limitations in 32-bit versions of Windows affect Distributed LAN Server performance. Due to these limitations, Distributed LAN Servers running on 32-bit versions of Windows are limited to 16 connections.</p> <p>(Each file system and network interface represents a connection. For example, if a Distributed LAN Server with 2 NICs serves 4 file systems to 4 Distributed LAN Clients, it has a total of 32 connections.)</p>
	<p>Do not use Windows Backup (the application included with Windows operating systems) to back up the directory on which a StorNext file system is mounted.</p>
	<p>As of StorNext release 3.5 the Authentication tab has been removed from the Windows Configuration utility. (For several previous StorNext releases a messaged warned that this tab would be removed in an upcoming release: "WARNING: Active Directory will be the only mapping method supported in a future release. This dialog will be deprecated.")</p>

Operating System / Affected Component	Description
Windows	<p>When you are upgrading to StorNext 3.5 from a release prior to version 3.0, you must uninstall StorNext before installing. After uninstalling you must reboot, install StorNext 3.5, and then reboot again. In future releases, upgrading from StorNext 3.5 should not require a reboot after installing.</p> <hr/> <p>If you are using the StorNext client software with Windows Server 2003, Windows Server 2008, Windows XP, or Windows Vista, turn off the Recycle Bin in the StorNext file systems mapped on the Windows machine.</p> <p>You must disable the Recycle Bin for the drive on which a StorNext file system is mounted. Also, each occurrence of file system remapping (unmounting/mounting) will require disabling the Recycle Bin. For example, if you mount a file system on E: (and disable the Recycle Bin for that drive) and then remap the file system to F:, you must then disable the Recycle Bin on the F: drive.</p> <p>As of release 3.5, StorNext supports mounting file systems to a directory. For Windows Server 2003 and Windows XP you must disable the Recycle Bin for the root drive letter of the directory-mounted file system. (For example: For C:\MOUNT\File_System you would disable the Recycle Bin for the C: drive.) For Windows Server 2008 and Windows Vista you must disable each directory-mounted file system.</p> <p>For Windows Server 2003 or Windows XP:</p> <ol style="list-style-type: none"> 1 On the Windows client machine, right-click the Recycle Bin icon on the desktop and then click Properties. 2 Click Global. 3 Click Configure drives independently. 4 Click the Local Disk tab that corresponds to the mapped or directory-mounted file system. 5 Click the checkbox Do not move files to the Recycle Bin. Remove files immediately when deleted. 6 Click Apply, and then click OK.

Operating System / Affected Component	Description
Windows	<p>For Windows Server 2008 and Windows Vista:</p> <ol style="list-style-type: none"> 1 On the Windows client machine, right-click the Recycle Bin icon on the desktop and then click Properties. 2 Click the General tab. 3 Select the mapped drive that corresponds to the StorNext mapped file system. For directory-mounted file systems, select the file system from the list. 4 Choose the option Do not move files to the Recycle Bin. Remove files immediately when deleted. 5 Click Apply. 6 Repeat steps 3-5 for each remaining directory-mounted file system. 7 When finished, click OK. <p>In StorNext releases prior to 3.5, the StorNext Windows client attempted to keep the UNIX uid, gid and mode bits synchronized with similar fields in the Windows security descriptor. However, these Windows and UNIX fields were often not synchronized correctly due to mapping and other problems. One consequence of this problem was that changing the owner in Windows incorrectly changed the UNIX uid and file permissions and propagated these errors into sub-directories.</p> <p>In release 3.5, the StorNext Windows client sets the UNIX uid, gid and mode bits only when Windows creates a file. The StorNext Windows client will no longer change the Unix uid, gid or mode bits when a Windows user changes the Windows security descriptor or Read-Only file attribute.</p> <p>If you change the UNIX mode bits and the file is accessible from Windows, you must change the Windows security descriptor (if Windows Security is configured On) or Read-Only file attribute to ensure the change is reflected on both Windows and UNIX.</p>
All	<p>To avoid parser errors, do not use "up" or "down" when naming items in the configuration file. This applies especially to naming affinities or any other string-type keyword.</p> <hr/> <p>StorNext file systems cannot be smaller than 5GB.</p>

Operating System / Affected Component	Description
All	<p>Be aware of the following limitations regarding file systems and stripe groups:</p> <ul style="list-style-type: none"> • The maximum number of disks per file system is 512 • The maximum number of disks per data stripe group is 128 • The maximum number of stripe groups per file system is 256 • The maximum number of tape drives is 256
	<p>To prevent data loss and aid disaster recovery, Quantum recommends adjusting policy settings to retain two copies of critical data. In addition, data stored to a storage disk should be stored on RAID (or, if only one copy is retained, the RAID should be mirrored). This way, in the event of a disk failure, no data is lost.</p>
	<p>The trashcan feature has been removed in StorNext 3.0. Before upgrading to StorNext 3.0, you must empty the trashcan on all file systems where this feature is enabled. If the trashcan directories are not removed, the upgrade will fail.</p> <p>To empty the trashcan on a file system where it is enabled, change to the file system directory (for example <code>/stornext/snfs1</code>). Then, at the command prompt, type:</p> <pre>rm -rf TrashCan</pre> <p>Repeat this procedure for all file systems.</p>
	<p>StorNext File System is incompatible with third-party portmappers.</p>
	<p>Hot re-zoning of SAN fabrics is not supported.</p>
	<p>Affinity names cannot be longer than eight characters. StorNext truncates affinity names after the eighth character, so if you have two or more affinities whose first eight characters are identical, StorNext considers them the same affinity. For example, if you have affinities called "affdata01" and "affdata02," StorNext sees them both as "affdata0" and treats them as the same affinity.</p>
	<p>For managed file systems only, the maximum directory capacity is 50,000 files per single directory. (This limitation does not apply to unmanaged file systems.)</p>

Operating System / Affected Component	Description
All	<p>The Move Stripe Group Data feature (part of Resource Allocation) does not support moving sparse files. Sparse files are files that lack on-disk allocations for some of the data within the data range indicated by the size of the file.</p> <p>A stripe group is defragmented as part of the data moving process. Because defragmenting a sparse file would make it unspare (and increase disk usage), snfsdefrag skips sparse files when defragmenting. Therefore, all existing sparse files remain on the original stripe group after moving of other files is complete.</p> <p>As an example, the gnu "cp" command uses a heuristic to attempt maintaining the "sparseness" of a file.</p> <p>Because the file system in question has already had the source stripe group marked read-only, the "cp" command has no choice but to move the file off of the original stripe group.</p> <hr/> <p>Clients outside a NAT firewall can no longer access a metadata controller inside the firewall.</p> <p>Support for client access through non-NAT firewalls continues to be supported via the fsports configuration file.</p> <hr/> <p>If you have configured custom mount options in the /etc/fstab file other than rw and diskproxy, if you subsequently add or remove the disk proxy settings using the StorNext GUI, any custom mount options will be lost. (Settings are added or removed in the StorNext GUI by navigating to the SNFS home page and then choosing Filesystems > Modify from the Config menu.)</p> <hr/> <p>In StorNext 3.0, the buffercachemin and buffercachemax parameters are deprecated and are replaced by buffercachecap.</p> <p>Also, in the Client Configuration Windows utility, the Buffer Cache Min and Buffer Cache Max settings are replaced by Buffer Cache Cap. The Non-Paged Pool Percentage setting has been removed.</p> <hr/> <p>If you use an SDisk or DDisk, you should save at least one other copy of the data to tape or to another SDisk or Ddisk. Without the second copy, your stored data will be vulnerable to data loss if you ever have hardware failure.</p>

Operating System / Affected Component	Description
All	<p>In StorNext 3.0, the default buffer cache settings have been modified. Previously, all reads/writes that were 64K or smaller went through the buffer cache while larger I/O requests went direct. In StorNext 3.0, read/writes that are 1MB or smaller go through the buffer cache, while larger I/O requests go direct.</p> <p>The new buffer cache settings may change the I/O behavior of some applications. For example, on managed servers, I/O to and from tape now goes through the buffer cache. To revert to the settings used in previous releases, change the following mount options on StorNext clients:</p> <pre>auto_dma_read_length=65537 auto_dma_write_length=65537</pre> <hr/> <p>In StorNext 3.0, obsolete methods of specifying FSS addresses are no longer supported.</p> <p>If the fsroutes configuration file is present, it is ignored and a warning is generated. In addition, the FSM address field for each entry in the fsmlist configuration file (the second column) is no longer supported and is ignored if present. If the priority field (the third column) is specified, then a period (.) must be used to take the place of the second column.</p> <p>The correct way to specify the metadata network for a cluster is through the addresses in the fsnameservers configuration file. If the address of a name server in fsnameservers is on network <x>, then network <x> will be used as a metadata network.</p> <hr/> <p>Don't run multiple copies of fsqueue.</p> <p>The fsqueue command checks the request queue and displays status on requests awaiting resources.</p> <p>Care should be taken to execute only one fsqueue process at a time. Invoking multiple simultaneous instances of fsqueue can consume system resources.</p> <p>The fsqueue process may take some time to complete if the file system contains large directories. Invoking fsqueue directly from a cron job frequently can result in multiple copies of fsqueue to be running at the same time. To run fsqueue from a cron job, wrap it in a shell script that checks to see if it is already running.</p>

Operating System / Affected Component	Description
All	<p>As of SNFS 2.7, a change was made to the way that the Reserved Extents performance feature affects free space reporting. In the previous release, SNFS would reserve a certain amount of disk space which would cause applications to receive an out of space error before the disk capacity reached 100%.</p> <p>In the current release, this reserved space is treated as allocated space. This allows applications to perform allocations until the file system is nearly full.</p> <p>NOTE: Due to allocation rounding, applications may still receive a premature out of space error, but only when there are just a few megabytes of space remaining. In the worst case, the error will be returned when the reported remaining space is: $(\text{InodeExpandMax} * \text{\#-of-data-stripe-groups})$</p> <p>One side effect of this change is that after creating a new file system, <code>df</code> will show that space has been used, even though no user data has been allocated.</p> <p>The amount of reserved space varies according to client use but does not go below a “floor” of a few gigabytes per data stripe group. The amount of reserved space at any time can be seen using the <code>cvadmin</code> command, selecting the file system, and using <code>show long</code>.</p> <p>While not recommended, the Reserved Extents feature can be disabled by applying the following setting to the Globals section of the FSM configuration file: <code>ReservedSpace No</code></p> <p>This will cause the file system to not reserve space for buffered I/O, thereby reducing buffer cache performance and possibly causing severe fragmentation.</p> <p>For more information, see The Reserved Space Parameter on page 21 and the <code>cvfs_config(4)</code> man page.</p>

Operating System / Affected Component	Description
All	<p>As a result of log rolling changes in StorNext 3.0, logs are now rolled every 6 hours. For each log, 28 instances (7 days of logs) are retained. Log instances are retained in the same directory as the original log.</p> <p>All log files which are rolled are affected by this change, including TSM logs (tac_00, bp_*.log, hist_01, etc.), MSM logs (tac_00, hist_01, etc.), and any other components configured for rolling. The <component>/config/filelist file contains roll_log entries that determine which files are rolled (where <component> is /usr/adic/TSM, /usr/adic/MSM, etc.).</p> <hr/> <p>The StorNext Library Space Used Report (accessible from the StorNext home page by choosing Library Space from the Reports menu,) shows the amount of nearline space used.</p> <p>The nearline space amount does not include dead space, but <i>does</i> include the following:</p> <ul style="list-style-type: none"> • All used space on all media in all libraries except vaults • All space used by files that were put on a storage disk or de-duplicated storage disk <hr/> <p>The StorNext GUI does not support an isolated metadata network topology with a system configured for HA failover. If the browser does not have connectivity to the isolated metadata network, then it will fail to connect after a failover event. If you have an isolated metadata network topology, make sure the browser has connectivity to the isolated metadata network or correct the network topology so the browser has connectivity to the isolated metadata network.</p> <p>For more information, contact Quantum Global Services.</p>

Appendix A: Operating Guidelines
General Operating Guidelines and Limitations



Appendix B HA Failover

StorNext is designed to be a resilient data management solution. StorNext supports operation in degraded mode and provides functionality to guarantee data protection in the event of a storage device failure or total site outage. For certain environments though, additional protection is required to deliver a higher level of availability. To meet these demands, StorNext includes MetaData Controller (MDC) failover.


MDC failover allows a secondary MDC to take over StorNext operations in the event a primary MDC fails. Failover is supported for all StorNext management operations including client IO requests (File System) and data mover operations (Storage Manager). MDCs in a failover pair typically run in an active / passive configuration, but both MDCs can be configured to run active File System processes. In the event one MDC fails, the other continues to perform its current operations, as well as those of the failed MDC.

Note: Active / Active Storage Manager processes are not currently supported in MDC failover.

Like all failover solutions, StorNext must provide functionality to prevent a damaged or inaccessible MDC from incorrectly processing IO requests that should be handled by the active MDC (often referred to as a “split brain” scenario). To handle this, StorNext utilizes a special failover methodology call STONITH - shoot the other node in the head. STONITH shuts down a degraded MDC and then reboots it so that, on recovery, it becomes the standby MDC in the event future MDC failover is required.

While other methods of failover are available, such as quorums, STONITH is the only widely acknowledged method of reliable failover control - especially for software that runs on multiple operating systems. Without STONITH, there is the possibility that a damaged MDC could continue acting as the primary MDC and attempt to handle client IO requests or move data between storage tiers. For shared file systems STONITH is especially critical because a MDC controls access for multiple hosts reading and writing to a single volume - and potentially the same file.

Currently MDC failover is only supported when configured by Quantum Solutions Engineering staff.



Appendix C

Using The Command Line Interface

Quantum recommends using the GUI to complete most StorNext tasks, but there may be situations where you need to use the command line interface instead. This appendix describes how to use the command line interface for the following tasks:

- [Labeling Disk Devices](#)
- [Modifying Global Settings](#)
- [Making a File System](#)
- [Starting and Stopping SNFS](#)
- [Unmounting or Mounting a File System](#)
- [Creating a File System Server](#)
- [Adding a File System Client](#)
- [Configuring a Stripe Group](#)
- [Adding an Affinity](#)
- [Creating a Disk-to-Disk Policy Class](#)
- [Enabling Stub File Support](#)
- [Managing Storage Disks with Deduplication Enabled](#)
- [Using the Dynamic Resource Allocation Feature](#)
- [Specifying an Alternate Retrieval Location](#)

Labeling Disk Devices

Each drive used by SNFS must be labeled. A new drive only needs to be labeled one time. A drive can be labeled from any StorNext server or client that has a Fibre Channel (FC) connection to the drive. Use this procedure to label a disk device using CLI.

Caution: The process of disk labeling re-partitions the drives. If you select an incorrect drive, you may lose data.

Also, it's a good practice to save a copy of your labels file. If you lose labels (and you using a Linux or Windows system,) the Quantum Technical Assistance Center will need this copy to help you relabel your devices.

All output examples shown in this appendix differ from the actual output, but the structure and information provided is similar.

- 1 On a SNFS client, at the system prompt, display a list of connected drives. Type: `/usr/cvfs/bin/cvlabel -l`

The command output is similar to this:

```

>
      Device          Volume Type/          Sectors  Sector
                   Label                               Size
ib [SEAGATE ST19171FS 0018] unknown Sectors: 17691712. Sector Size: !
lc [SEAGATE ST19171FS 0018] unknown Sectors: 17691712. Sector Size: !
ld [SEAGATE ST19171FS 0018] unknown Sectors: 17691712. Sector Size: !
```

- 2 Looking at the output information, identify any drives that are unused or do not have a recognized Volume Type. For these drives, write down their associated device names.
- 3 Create `/usr/cvfs/config/cvlabels` by typing the following:

```
/usr/cvfs/bin/cvlabel -c > /usr/cvfs/config/cvlabels
```

The created file displays an entry for disks located by the `/usr/cvfs/bin/cvlabel` command.

```
CvfsDisk_UNKNOWN /dev/sdb # host 4 lun 1 sectors  
639570752 ...  
CvfsDisk_UNKNOWN /dev/sdc # host 4 lun 2 sectors  
639570752 ...  
CvfsDisk_UNKNOWN /dev/sdd # host 4 lun 3 sectors  
639570752 ...
```

Caution: Identify any drives that already contain a recognized Volume Type. Do not write a label to these drives or you may lose data.

- 4 In the `/usr/cvfs/config/cvlabels` file, delete any lines that refer to disks that will not be labeled or have already been labeled.
- 5 Edit `/usr/cvfs/config/cvlabels` file to provide a unique name for each drive to be used by SNFS.

In this example, the UNKNOWN variable in the drive name associated with disk device `/dev/sdb` has been renamed to a numeral (in sequence) 0,1,and 2. The disk devices have also been alphabetically ordered.

```
CvfsDisk0 /dev/sdb # host 4 lun 1 sectors 639570752  
...  
CvfsDisk1 /dev/sdc # host 4 lun 2 sectors 639570752  
...  
CvfsDisk2 /dev/sdd # host 4 lun 3 sectors 639570752  
...
```

- 6 Save the `/usr/cvfs/config/cvlabels` file.
- 7 Use the `cvlabel` command to label the disks. Type:
`/usr/cvfs/bin/cvlabel /usr/cvfs/config/cvlabels`
- 8 For each disk, you are prompted to verify that you want to label each disk. Type Y for yes.

Modifying Global Settings

The global section of the file system configuration file contains general parameters that control enabling and disabling features, system performance, and components related to the file system's resource consumption.

For most of these parameters, the only thing necessary for the modified parameter to take effect is to restart the File System Manager (FSM). However, the following parameters require that the file system be fully re-initialized (which will result in data loss,) before they take effect:

- FSBlockSize
- WindowsSecurity

If a parameter change requires file system re-initialization, the system notifies the administrator in the system log. In order to reduce the number of file system remakes, be sure to plan the initial configuration of the FSBlockSize and WindowsSecurity parameters carefully.

The global section also contains several parameters that can dramatically improve or degrade system performance. Exercise caution when modifying performance parameters. One key performance parameter is InodeCacheSize.

Before making any changes to the file system's configuration, carefully review the `cvfs_config(4)` man pages or the "CVFS Configuration File" help file.

Use this procedure to modify system global settings using CLI.

- 1 Unmount the file system by typing the following:

```
umount <file_system_name>
```

Where the file system name is the name of the file system where the settings are being modified.

- 2 Stop the file system by typing the following:

```
/usr/cvfs/bin/cvadmin  
snadmin> stop <file_system_name>  
snadmin> quit
```

where snadmin is the prompt shown after invoking the cvadmin command.

Note: When the file system is down, file system operations will pause and some applications could fail. Plan accordingly to minimize disruptions.

3 Edit the configuration file by typing the following:

```
edit /usr/cvfs/config/<file_system_name>.cfg
```

4 Make the appropriate edits to the configuration file.

5 Restart the file system. Type:

```
/usr/cvfs/bin/cvadmin  
snadmin> start <file_system_name>  
snadmin> activate <file_system_name>  
snadmin> select  
snadmin> quit
```

6 Mount the file system by typing the following:

```
mount -t cvfs <file_system_name> <mount_point>
```

For example: `mount -t cvfs snfs1 /stornext/snfs1`

Making a File System

Following are some reasons to make or re-make a file system:

- Creating a new file system
- Removing a stripe group from the file system
- Removing a disk from a stripe group
- Changing a stripe group's stripe breadth
- Changing the sector count of a disk

Caution: Making or re-making a file system will result in a complete loss of user data.

After creating a relation point on a managed file system, you must delete and recreate the file system, not just remake. For information about deleting a file system, see [Deleting a File System](#) on page 25. For information about creating a file system, see [Adding a File System](#) on page 12.

Use this procedure to make a file system using the CLI.

- 1 Log on to the machine as root.
- 2 Make sure you are the root user by typing the following:

```
cd /
```
- 3 Unmount the system by typing the following:

```
umount <mount_point>
```

where the *<mount_point>* is where you have mounted the SNFS.
For example:

```
umount /stornext/snfs1
```
- 4 Stop the system by typing the following:

```
/usr/cvfs/bin/cvadmin  
snadmin> stop <file_system_name>
```

where the *<file_system_name>* is the name of the file system to be stopped.
- 5 Make the file system by typing the following:

```
/usr/cvfs/bin/cvmkfs <file_system_name>
```
- 6 Start the system by typing the following:

```
/usr/cvfs/bin/cvadmin  
snadmin> start <file_system_name>
```
- 7 Mount the system by typing the following:

```
mount -t cvfs snfs1 /stornext/snfs1
```

Starting and Stopping SNFS

To start SNFS using the CLI, type the following:

```
/usr/cvfs/bin/cvadmin  
snadmin> select  
snadmin> start <file_system_name>  
snadmin> activate <file_system_name>
```

where snadmin is the prompt shown after invoking the cvadmin command.

To stop SNFS using the CLI, type the following:

```
/usr/cvfs/bin/cvadmin  
snadmin> select  
snadmin> stop <file_system_name>
```

Unmounting or Mounting a File System

To unmount a file system using the CLI, follow these steps:

- 1 Change the login user to root. Type:

```
cd /
```

- 2 Unmount the system. Type:

```
umount <mount_point>
```

where the *<mount_point>* is where you are mounting the SNFS.

For example:

```
umount /stornext/snfs1
```

To mount a file system using the CLI, follow these steps:

- 1 Change the login user to root. Type:

```
cd /
```

- 2 Mount the system, type:

```
mount -t cvfs <file_system_name> <mount_point>
where the <mount_point> is where you are mounting the SNFS.
For example:
mount -t cvfs snfs1 /stornext/snfs1
```

Creating a File System Server

The follow procedure describes how to create a file system server using the CLI.

Before initially executing any SNFS command line programs, you are required to source either the `.profile` or the `.cshrc` file. This updates the user environment with the SNFS environment variables.

- If you are running `sh`, `ksh`, or `bash`, type:

```
. /usr/adic/.profile
```

- For all other shells, type:

```
source /usr/adic/.cshrc
```

Caution: Do not attempt to perform the procedures in this section unless you have completed Quantum's StorNext training and are confident you understand all procedures. (You may perform these procedures if Professional Services is assisting you.)

- 1 Install StorNext as described in the *StorNext Installation Guide*. Follow the instructions that pertain to your operating system.
- 2 Create a list of system and FC disks by writing to a file in a format recognized by the `cvlabel` command. Type the following:

```
/usr/cvfs/bin/cvlabel -c > /usr/cvfs/config/cvlabels
```

The created file displays an entry for disk located by the `/usr/cvfs/bin/cvlabel` command.

```
CvfsDisk_UNKNOWN /dev/sdb # host 4 lun 1 sectors
639570752 ...
```



```
CvfsDisk_UNKNOWN /dev/sdc # host 4 lun 2 sectors  
639570752 ...  
CvfsDisk_UNKNOWN /dev/sdd # host 4 lun 3 sectors  
639570752 ...
```

- 3 Edit the cvlabels file that has a list of all system and FC disks visible on the machine. Edit the file to remove all the system disks and any FC disks you do not want labeled, as well as FC disks that are already labeled.

- 4 Label the FC drives by typing the following:

```
/usr/cvfs/bin/cvlabel /user/cvfs/config/cvlabels
```

- 5 Copy the example file system configuration file to the config directory by typing the following:

```
cp /usr/cvfs/examples/example.cfg /usr/cvfs/config/  
<file_system_name>.cfg
```

- 6 Edit the StorNext configuration file you just created to include the desired settings, disks, and stripe groups.

- 7 Copy the fsnameservers file to the config directory by typing the following:

```
cp /usr/cvfs/examples/fsnameservers.example /usr/  
cvfs/config/fsnameservers
```

- 8 Edit the fsnameservers file to include the host's IP address.

- 9 Copy the example fsmlist file to the config directory by typing the following:

```
cp /usr/cvfs/examples/fsmlist.example /usr/cvfs/  
config/fsmlist
```

- 10 Edit the fsmlist file created in step 9 to include the name of file systems you want to start at boot time.

- 11 Obtain your license.dat from the Quantum Technical Assistance Center. For contact information, refer to [Quantum Technical Assistance Center](#) on page 11.

- 12 Place the license.dat file in the /usr/cvfs/config directory.

- 13 Make the file system by typing the following:

```
/usr/cvfs/bin/cvmkfs <file_system_name>
```

Caution: When you run the `cvmkfs` command, you will lose any data currently on the file system.

- 14 Reboot the machine.
- 15 Verify that the labeled drives are available to the file system by typing the following:

```
/usr/cvfs/bin/cvlabel -l
```

- 16 Restart the file system. Type:

```
/usr/cvfs/bin/cvadmin  
snadmin> start <file_system_name>  
snadmin> activate <file_system_name>  
snadmin> select <file_system_name>  
snadmin> show  
snadmin> who  
snadmin> quit
```

- 17 Mount the system by typing the following:

```
mount -t cvfs <file_system_name> <mount_point>
```

For example: `mount -t cvfs snfs1 /stornext/snfs1`

Adding a File System Client

Use this procedure to add a file system client using the CLI.

- 1 Install StorNext as described in the *StorNext Installation Guide*. Follow the instructions that pertain to your operating system.
- 2 Copy the example `fsnameservers` file to the config directory by typing:

```
cp /usr/cvfs/examples/fsnameservers.example /usr/  
cvfs/config/fsnameservers
```

Caution: The `fsnameservers` file must be the same on all machines.

- 3 Edit the `/etc/fstab` file to mount on boot (for Solaris, edit the `/etc/vfstab` file).
- 4 Reboot the machine.
- 5 Verify that the system has mounted by typing the following:
`df -k`
If your machine did not mount on boot, refer to “Resolving Installation Problems” in the *StorNext Installation Guide*.

Configuring a Stripe Group

Following is an example of a stripe group configuration in a file system configuration file. You can use a text editor to modify any of these variables. The values shown might not be representative of typical definitions.

Caution: Modifying the following variables might result in a complete loss of user data and a re-make of the file system:

- Metadata
- StripeBreadth
- Removing a Node

```
[StripeGroup StripeGroup1]
Status UP
Exclusive Yes
MetaData Yes
Journal Yes
Read Enabled
Write Enabled
StripeBreadth 16
MultiPathMethod Rotate
Node fortune1 0
```

```
[StripeGroup StripeGroup2]
Status UP
Affinity media1
Read Enabled
Write Enabled
StripeBreadth 16
Rtios 100
Rtmb 100
RtiosReserve 100
RtmbReserve 100
MultiPathMethod Rotate
Node fortune2 0
```

Adding an Affinity

This procedure lets you set stripe group affinities for assigning file locations to a specific file system stripe group. All subsequent allocations to a file that have been assigned a valid stripe group affinity will occur on the specified stripe group.

Affinities may also be assigned to files and directories using the `cvaffinity` command. For more information about the `cvaffinity` command, refer to the CLI Reference Guide.

- 1 Unmount the file system by typing the following:

```
umount <mount_point>
```

where the `<mount_point>` is where you have mounted the SNFS.

For example: `umount /stornext/snfs1`

- 2 Use `cvadmin` to stop the file system by typing the following:

```
/usr/cvfs/bin/cvadmin
snadmin> stop <file_system_name>
```

(where the `<file_system_name>` is the file system to which the affinity will be added.)

```
snadmin> quit
```

- 3 Using a text editor, open the file system configuration file (*<file_system_name>.cfg*), where *<file_system_name>* is the name of the file system.
- 4 Locate the stripe group section and select the stripe group to which you want to add the affinity.

Following is part of a stripe group configuration with an affinity definition:

```
[StripeGroup StripeGroup2]
Status UP
Affinity aff1
Read Enabled
Write Enabled
StripeBreadth 16
Node disk2 0

StripeGroup StripeGroup3]
Status UP
Affinity aff2
Read Enabled
Write Enabled
StripeBreadth 16
Node disk3 0
Node disk4 1
Node disk5 2
```

If you create a directory association with Affinity aff1, all data written to that directory is written to StripeGroup StripeGroup2, and therefore only to Disk disk2. If you make an association with Affinity aff2 and a separate directory in the file system, all data is directed to StripeGroup StripeGroup3, which contains three disks: disk3, disk4, and disk5. All data is written to these disks when directed to the associated directory with Affinity aff2.

- 5 Add the affinity to the selected stripe group with a line entry in this format:

```
Affinity <affinity_name>
```

where *<affinity_name>* is the name of the affinity. (Affinity names cannot be longer than eight characters.)

Following is an example of a stripe group configuration after an affinity was added. The affinity line is the flag for the Data1 stripe group.

```
[StripeGroup Data1]
Status UP
Read Enabled
Write Enabled
Affinity data1aff
StripeBreadth 512
Node CvfsDisk2 0
Node CvfsDisk3 1
```

- 6 Save and close the configuration file.
- 7 Use cvadmin to start the file system by typing the following:

```
/usr/cvfs/bin/cvadmin
snadmin> start <file_system_name>
snadmin> activate <file_system_name>
snadmin> select <file_system_name>
snadmin> show
```

- 8 Mount the file system by typing the following:

```
mount -t cvfs <file_system_name> <mount_point>

For example: mount -t cvfs snfs1 /stornext/snfs1
```

- 9 Create an affinity/directory association by typing the following:

```
/usr/cvfs/bin/cvmkdir -k <affinity_name>
<directory_name>
```

```
For example: /usr/cvfs/bin/cvmkdir -k data1_aff /
stornext/snfs1/video_data
```

This associates all data written to /stornext/snfs1/video_data with stripe group Data1 and, therefore, disks CvfsDisk2 and CvfsDisk3.

Creating a Disk-to-Disk Policy Class

Use this procedure to create a disk-to-disk-specific policy class. You must have at least two affinities configured to create a disk-to-disk policy class.

Before initially executing any StorNext command line programs, you are required to source either the `.profile` or the `.cshrc` file. This will update the user environment with the StorNext environment variables.

- If you are running `sh`, `ksh`, or `bash`, type:

```
. /usr/adic/.profile
```
- For all other shells, type:

```
source /usr/adic/.cshrc
```

Use the `fsaddclass` command to create a new policy class by typing the following:

```
/usr/adic/TSM/exec/fsaddclass <policy_class_name> -a  
<default_affinity> <destination_affinity> -i  
<relocation_time_in_days>
```

For example: `/usr/adic/TSM/exec/fsaddclass dtdclass1 -a Aff1 Aff2 -i 1`

The list of affinities includes the same affinities that were defined through the GUI.

If you do not use the `-i` option (`MinRelocTime`), the default relocation time of seven days is used.

If only one affinity is listed with the `-a` option, no relocation occurs because a destination affinity is not defined. The first affinity listed after the `-a` option is the default affinity.

Modifying a Disk-to-Disk Policy Class

Use the `fsmodclass` command to modify an existing policy class by typing the following:

```
/usr/adic/TSM/exec/fsmodclass <policy_class_name> -a  
<default_affinity> <destination_affinity> -i  
<relocation_time_in_days>
```

For example: `/usr/adic/TSM/exec/fsmodclass dtdclass1 -a Aff1 Aff2 -i 1`

Manual Disk-to-Disk Relocation

Use the `fsrelocate` command to perform manual disk-to-disk migration. You can use this command to relocate a file from the current affinity to another affinity, provided it meets these criteria:

- The file must be a non-zero sized file
- The file cannot be truncated
- The file cannot be specifically excluded from relocation via the `fschfiat` command

For example:

```
/usr/adic/TSM/exec/fsrelocate /stornext/snfs1/data1/  
file1 -a Aff2
```

In the previous example, the file `/stornext/snfs1/data1/file1` will be relocated to affinity `Aff2`.

You can also use metacharacters when defining the file name parameter.

For example:

```
/usr/adic/TSM/exec/fsrelocate -a Aff2 /stornext/snfs1/  
relocate/*
```

Enabling Stub File Support

The following file system commands have been enhanced to support StorNext's Stub File feature:

- `fsaddclass`
- `fsmodclass`
- `fschfiat`
- `fschdiat`

Specifically, a **-S stubsize** option has been added to each command. This option allows you to specify the desired file stub size in kilobytes.

The following example illustrates usage for each command:

- `fsaddclass class -S 1000`
- `fsmodclass class -S 1000`
- `fschfiat -S 1000`

Additionally, the following commands remain unchanged, but their output now includes the target stub size and actual stub size:

- `fsclassinfo`: Shows Stub File Size
- `fsfileinfo`: Shows Stub Size (target stub file size) and Stub Length (actual stub file size)

Managing Storage Disks with Deduplication Enabled

Quantum recommends using the StorNext GUI to manage the storage disk deduplication feature, but you can use the CLI to add, modify, or delete storage disks with deduplication enabled (Dedup Sdisk).

Adding a Dedup Sdisk

The `fsdiskcfg` command now includes a `-b` option to indicate that the storage disk uses blockletized storage to store data. (Blockletized storage is how this feature accomplishes space savings.)

For example, you would enter the following command to create a storage disk with deduplication enabled:

```
Input: fsdiskcfg -a -p /apps -b sdisk1
```

Before you add a dedup storage disk, you can also use the `fsdiskcfg` command's `-b -l` options to view a list of available locations where a dedup storage disk can be created. The correct syntax is `fsdiskcfg -b -l`.

(You can also use the `fsdiskcfg` command's `-l` option (without the `-b` option) to view all available location where a standard storage disk can be created. The correct syntax is `fsdiskcfg -l`.)

Modifying a Dedup Sdisk

You can modify the location for the dedup sdisk, but the dedup sdisk must be blank and the new location must be under the same file system.

For example, this is how you would invoke the `fsdiskcfg` command for `sdisk1`, which was previously created under `/apps`:

```
Input: fsdiskcfg -m sdisk1 -p /apps/subdir
```

Deleting a Dedup Sdisk

You can delete a dedup sdisk by using the `fsdiskcfg` command's `-d` option.

For example, this is how you would invoke the `fsdiskcfg` for `sdisk1` that was previously created:

```
Input: fsdiskcfg -d sdisk1
```

Obtaining Dedup Sdisk Information

You can obtain information for a dedup sdisk by running the `fsmedinfo` command on the dedup sdisk.

For example, if you invoke the `fsmedinfo` command for a dedup sdisk named `sdisk1`, the output looks similar to this:

```
### fsmedinfo sdisk1
```

```
-----  
-----  
Media Information ReportTue Feb 6 13:17:32 2007
```

```
Media ID:   ddisk(0)
```

```
Media Type: DDISK  
-----  
-----
```

```
Storage Area: VolSub
```

```
Class ID: <system blank>Bytes Used: 4,780,195,840
```

```
Last Accessed: 06-feb-2007 12:04:52Space Remaining:  
68,623,007,744
```

```
Media Status: AVAILPercent Used: 6.51
```

```
Write Protect: NSuspect Count: 0
```

```
Mark Status: UNMARKEDMount Count: 0
```

```
Medium Location: SLOT/BIN
```

```
Formatted: Y
```

```
Number of Segments: 0
```

External Location: N/A
Total Blob Bytes: 0
Unique Blob Bytes: 0
Percent Eliminated: 0.00

FS0000 06 1703716962 fsmedinfo completed: Command Successful.

Note: The Space Remaining amount shown does not take into account the percentage of redundancy elimination; it shows only the physical space remaining on the disk.

Obtaining Distributed LAN Client Information

If your StorNext configuration includes distributed LAN clients, you can obtain information through three commands supported by the cvadmin command:

- proxy
- proxy long
- proxy who

The proxy Command

Use the proxy command to display information about the distributed LAN servers for the file system. In particular, the IP address and port number on which the distributed LAN server is listening is shown.

This command requires that a file system be selected.

Command usage and output looks similar to this:

```
snadmin (yy) > proxy
```

```
Disk Proxy Server 172.16.82.130 (pmport 49152, pmflags  
0x0)
```

```
Listening on 172.16.82.130 port 1036
```

The proxy long Command

Use the proxy long command to display the same information about distributed LAN servers included in the proxy command output, plus the tuning parameters for the distributed LAN server systems and a list of the disks available for distributed LAN usage.

This command requires that a file system be selected.

Command usage and output looks similar to this:

```
snadmin (yy) > proxy long
```

```
Disk Proxy Server 172.16.82.130 (pmport 49152, pmflags  
0x0)
```

```
Listening on 172.16.82.130 port 1036
```

```
windowsize=1024K sbsize=256K sbcount=32
```

```
CvfsDisk0 on device:PhysicalDrive1 sectors:156214012  
sector size: 512
```

The proxy who Command

Use the proxy who command to display the active disk distributed LAN connections for the specified host, which can be either a disk distributed LAN server or client. Output also includes average I/O statistics for each connection. The syntax for this command is proxy who *hostname*.

Command usage and output looks similar to this (both distributed LAN client and distributed LAN server output is shown):

```
snadmin (yy) > proxy who y
```

```
FS 'yy'
```

```
Disk Proxy Client connection from 172.16.82.62
```

```
Remote address 172.16.82.62 port 1052 flags 0x2
```

```
Read 1.2 Mbytes/s, write 0.0 bytes/s
```

```
snadmin (yy) > proxy who fie
FS 'yy'
  Disk Proxy Server connection to 172.16.82.130
  Remote address 172.16.82.130 port 1036 flags 0x1
  Read 0.0 bytes/s, write 1.2 bytes/s
```

Using the Resource Allocation Feature

Quantum recommends that you perform resource allocation using the StorNext GUI. However, if your operating system does not support using the GUI for this feature (or if you are operating in a failover environment,) you can accomplish the following tasks from the command line:

- [Adding a Stripe Group Without Moving](#)
- [Adding and Moving a Data Stripe Group](#)
- [Moving a Metadata/Journal Stripe Group](#)

Caution: When you add a new disk or stripe group to your SAN, often an OS-dependent operation must be run to make the added device recognizable by a host. Some of these utilities can disrupt access to existing disks, causing access hangs or failures. To avoid this, stop all file system operations on the affected host *before* rescanning for the new device.

Checking the File System

Before you use the Resource Allocation feature, Quantum strongly recommends running the `cvfsck` command on the file system you will be using. This step could take a considerable amount of time to complete, but your file system should be in good condition before you attempt to expand it or move stripe groups.

Caution: If you do not run the `cvfsck` command to check your file system before attempting file system expansion, **irreparable file system damage could occur.**

Adding a Stripe Group Without Moving

Use the following procedure to expand the file system by adding a stripe group, and not migrating.

- 1 Label disks for the new stripe groups you want to add to the file system.
- 2 If your StorNext configuration includes a failover environment, you must first shut down any standby FSMs that would start when you shut down the primary FSM. The movement procedure will not complete successfully unless all FSMs are shut down.

Caution: If you do not shut down standby FSMs, **file system corruption or data loss could occur.**

- 3 (Optional) Run the `cvfsck` command on the file system. See [Checking the File System](#).
- 4 Add the new stripe groups to the file system.
- 5 Stop the File System Manager (FSM).
- 6 Run the `cvupdatefs` command.
- 7 Restart the FSM.

Adding and Moving a Data Stripe Group

New functionality has been added to the `snsdefrag` utility to support operations on multiple stripe groups.

Note: During Stripe Group Movement, affinities are preserved when files are moved from one stripe group to another. When you create a new stripe group to use with the Stripe Group Movement feature, the new stripe group must include sufficient space for its affinities. (You must add any affinities from the source stripe group to the new stripe group.)

Use the following procedure to add new stripe groups, and then move data off of the old stripe group.

- 1 Label disks for the new stripe groups you want to add to the file system.
- 2 If your StorNext configuration includes a failover environment, you must first shut down any standby FSMs that would start when you shut down the primary FSM. The move procedure will not complete successfully unless all FSMs are shut down.

Caution: If you do not shut down standby FSMs, file system corruption or data loss could occur.

- 3 (Optional) Run the `cvfsck` command on the file system. See [Checking the File System](#).
- 4 Add the new stripe groups to the file system configuration and mark the old stripe groups as read-only. (Make sure the old stripe group is write disabled.)
- 5 Stop the File System Manager (FSM) for the desired file system.
- 6 Run `cvupdatefs`.
- 7 Restart the FSM.
- 8 Run `snfsdefrag -G <n> -m 0 -r /filesystemroot`
where `<n>` is the zero-based number of the source stripe group from which the move starts, and `filesystemroot` is the file name of the file system tree's root. You can specify multiple `-G` options to use multiple source stripe groups.
- 9 Verify that no data remains on the original stripe groups.
- 10 Edit the file system configuration to mark the old stripe groups as "Down."
- 11 Stop the FSM.
- 12 Restart the FSM.

Note: The old stripe groups marked "Down/Readonly" must be left in the file system configuration file.

Moving a Metadata/ Journal Stripe Group

Metadata movement is performed on a LUN level, meaning you must specify the source LUN and the destination LUN. The new `sndiskmove` command that accomplishes metadata movement has two arguments: a source and destination LUN.

After movement is complete, the physical source disk can be removed.

Note: Although a stripe group can consist of multiple disks or LUNs, the `sndiskmove` command moves only a single disk or LUN. Consequently, references to “stripe group” in this section refer to a single disk or LUN when migrating metadata with `sndiskmove`.

Caution: The metadata/journal stripe group you want to move cannot contain data.

`Sendiskmove` treats metadata and journal stripe groups the same way, so it doesn't matter whether the stripe group you want to move is a metadata stripe group, a journal stripe group, or a combined metadata and journal stripe group. The only caveat is that stripe groups used for movement cannot contain data.

If you attempt to move a metadata/journal stripe group that contains data, **data loss could occur**.

Use the following procedure to move a metadata/journal stripe group from a source LUN to a destination LUN.

- 1 Stop the File System Manager (FSM) for the file system.
- 2 If your StorNext configuration includes a failover environment, you must shut down any standby FSMs that would start when you shut down the primary FSM. The movement procedure will not complete successfully unless all FSMs are shut down.

Caution: If you do not shut down standby FSMs, **file system corruption or data loss could occur**.

- 3 (Optional) Run the `cvfsck` command on the file system. See [Checking the File System](#).

- 4 Run `sndiskmove <source-LUN-label-name> <destination-LUN-label-name>`

where `<source-LUN-label-name>` is the source stripe group from which the move starts, and `<destination-LUN-label-name>` is the destination stripe group to which you want to move data.

During the move process StorNext appends “.old” to the source stripe group name. This is to avoid confusion because the destination stripe group is given the same name as the original stripe group. Both stripe group names remain in the configuration file.

For example:

`source-LUN-label-name` (the original stripe group name) becomes `source-LUN-label-name.old`

`destination-LUN-label-name` (the new stripe group name) becomes `source-LUN-label-name` (the same name as the original stripe group)

Note: When you run `sndiskmove`, it could take a considerable amount of time to copy the data between disks, depending on disk size and performance.

- 5 Only if your system includes a standby FSM: After you run `sndiskmove`, rescan the disks on the standby FSM's host by running `cvadmin -e 'disks refresh'`. You must run `cvadmin -e 'disks refresh'` on all systems on which you have a configured FSM for the file system involved in the move.
- 6 Restart the FSM.
- 7 Only if your system includes a standby FSM: Restart the standby FSM.

Specifying an Alternate Retrieval Location

The Alternate Retrieval Location feature applies to situations where file retrieval fails because the normal file copies cannot be retrieved from the machine on which StorNext Storage Manger resides. This feature

enables you retrieve a copy of the truncated file from a different machine.

Quantum strongly recommends using the StorNext GUI to enable this feature. If you choose to use the command line interface instead, you can use the `CP1altnode` command to add, modify, or delete alternate retrieval location information. See the `fsaltnode` man page for more information.

Also, a new `-x` argument has been added for `fsretrieve`, which enables you to make changes when this feature is enabled. For more information, see the man page for `fsretrieve`.



Appendix D RAS Messages

RAS messages appear when StorNext encounters an error condition. The RAS window shows symptoms of the condition, plus workarounds you can try to resolve the condition before calling the Quantum Technical Assistance Center.

This appendix shows the different RAS messages you might see. Messages are separated into the following categories:

- [Media and Drive RAS Messages](#)
- [SNFS RAS Messages](#)
- [Other RAS Messages](#)

Media and Drive RAS Messages

This section describes RAS messages that might appear as a result of a media-related error condition, such as no media detected or media format failure.

Figure 1 No Media Found RAS

Recommended Actions	
No Media Found	
IF	THEN
If the service ticket indicates that no media is found to satisfy the request:	<ul style="list-style-type: none"> • Add more media to StorNext. <p>or</p> <ul style="list-style-type: none"> • Check for suspect media.
The problem IS resolved:	Close the service ticket. Refer to Closing Service Tickets .
The problem has NOT been resolved:	<ol style="list-style-type: none"> 1. Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. 2. Contact the Quantum Technical Assistance Center. In the USA: 1-800-284-5101 UK, France and Germany: 00800 4 QUANTUM EMEA: +44 1256 848 766 On the Web: http://www.quantum.com/support

Figure 2 Possible Drive/Media Mount Discrepancy RAS

Recommended Actions	
Possible Drive/Media Mount Discrepancy	
IF	THEN
A service ticket indicates the drive is mounted and the media mounted in the drive cannot be verified:	<p>This is a caution regarding drive and media mounts, and might require user intervention. In this situation, StorNext assigns to the unverified media the barcode (media ID) of the last tape mounted in the drive, and continues to operate.</p> <p>If StorNext cannot dismount this drive at a later time, dismount it manually.</p> <ol style="list-style-type: none"> 1. Check the drive to see if the media has been ejected. 2. If the media has not been ejected, press Eject on the drive. 3. Try to dismount the drive again using the GUI. (From the SNSM home page, choose Library > Dismount from the Media menu.) 4. If the dismount fails using the GUI, dismount the drive using the operator panel on the physical library. 5. Using the GUI, perform an audit to make sure that the Remap Audit checkbox is selected. (From the SNSM home page, choose Library > Audit Library from the Admin menu.)
A service ticket indicates the drive is NOT mounted and the media mounted in the drive cannot be verified:	<ol style="list-style-type: none"> 1. Dismount the drive via the operator panel on the physical library. 2. Using the GUI, perform an audit to making sure that the Remap Audit checkbox is selected. (From the SNSM home page, choose Library > Audit Library from the Admin menu.)
The problem IS resolved:	Close the service ticket. Refer to Closing Service Tickets .
The problem has NOT been resolved:	<ol style="list-style-type: none"> 1. Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. 2. Contact the Quantum Technical Assistance Center. In the USA: 1+800-284-5101 UK, France and Germany: 00800 4 QUANTUM EMEA: +44 1256 848 766 On the Web: http://www.quantum.com/support

Figure 3 Tape Drive Alerts RAS
part 1

Recommended Actions

Tape Drive Alerts

Follow the recommendations below after the tape drive has issued a tape alert. Some alerts are fatal and indicate that the drive is no longer useful. Other alerts indicate that user intervention (such as cleaning) will correct the problem. Note the flag number from the ticket for use in troubleshooting.

The host application should have received the same tape alert message. Not all host applications respond with the same behavior.

The recommendations below are based on best practices for a typical host application.

IF	THEN
Flag 1 (01h) - Read warning Flag 2 (02h) - Write warning	Contact the Quantum Technical Assistance Center .
Flag 3 (03h) - Hard Error The problem could be the tape or the drive. The drive cannot isolate the source.	<ol style="list-style-type: none"> 1. Close the ticket and retry the read/write operation on the original drive and media. 2. Monitor operation for a reoccurrence of the ticket. 3. Insert the suspect media into an alternate drive and retry the read/write operation. 4. If the error follows the media, retire the media. 5. If the error stays with the drive, contact the Quantum Technical Assistance Center to replace the drive.
Flag 4 (04h) - Media	<ol style="list-style-type: none"> 1. Copy the data to another piece of media. 2. Remove the original media from the library and discard.
Flag 5 (05h) - Read Failure Flag 6 (06h) - Write Failure The problem could be the tape or the drive. The drive cannot isolate the source.	<ol style="list-style-type: none"> 1. Close the ticket and retry the read/write operation on the original drive and media. 2. Monitor operation for a reoccurrence of the ticket. 3. Insert the suspect media into an alternate drive and retry the read/write operation. 4. If the error follows the media, retire the media. 5. If the error stays with the drive, contact the Quantum Technical Assistance Center to replace the drive.
Flag 7 (07h) - Media life	<ol style="list-style-type: none"> 1. Copy the data to another piece of media. 2. Remove the original media from the library and discard.
Flag 9 (09h) - Write protect	If the media is used for writing data, adjust the write-protect tab on the media and clear the write-protect setting.
Flag 10 (0Ah) - No removal	The prevent media removal flag is on, so it must be turned off.
Flag 11 (0Bh) - Cleaning media	If cleaning media does not reside in the cleaning media pool, move it there.

Figure 4 Tape Drive Alerts RAS
part 2

Flag 12 (0Ch) - Unsupported format	The media is not supported by the drive. If the media is blank, remove it. Otherwise, contact the Quantum Technical Assistance Center .
Flag 13 (0Dh) - Recoverable snapped tape Flag 14 (0Eh) - Unrecoverable snapped tape	Contact the Quantum Technical Assistance Center .
Flag 15 (0Fh) - Memory chip in cartridge failure	<ol style="list-style-type: none"> 1. Write protect the media. 2. Copy the data to a new piece of media. 3. Discard the old media.
Flag 16 (10h) - Forced eject	Investigate whether the StorNext administrator's actions might have initiated an eject operation.
Flag 17 (11h) - Read-only format	The loaded cartridge is a read-only type in this drive. If the media contains data, write protect it.
Flag 18 (12h) - Tape directory corrupted on load	The tape directory must be rebuilt. Contact the Quantum Technical Assistance Center .
Flag 19 (13h) - Nearing media life	<ol style="list-style-type: none"> 1. Copy the data to another piece of media. 2. Remove the original media from the library and discard.
Flag 20 (14h) - Clean now Flag 21 (15h) - Clean periodic	Clean the drive.
Flag 22 (16h) - Expired cleaning media Flag 23 (17h) - Invalid cleaning media	<ol style="list-style-type: none"> 1. Remove the media. 2. Add new media.
Flag 24 (18h) - Retension requested Flag 25 (19h) - Dual-port interface error Flag 26 (1Ah) - Cooling fan failure Flag 27 (1Bh) - Power supply failure Flag 28 (1Ch) - Power consumption Flag 29 (1Dh) - Drive maintenance	Contact the Quantum Technical Assistance Center .
Flag 30 (1Eh) - Hardware A Flag 31 (1Fh) - Hardware B A hardware error has occurred that should be captured and returned for failure analysis.	Contact the Quantum Technical Assistance Center .

Figure 5 Tape Drive Alerts RAS
part 3

Flag 32 (20h) - Interface	<ol style="list-style-type: none"> 1. Check the cabling between the library and the attached tape library. 2. If the problem is unresolved, contact the Quantum Technical Assistance Center.
<p>Flag 33 (21h) - Eject media</p> <p>The drive has experienced an issue that can be resolved by unloading and reloading media.</p>	<ol style="list-style-type: none"> 1. Retry the operation. 2. If the problem persists, contact the Quantum Technical Assistance Center.
Flag 34 (22h) - Firmware download via SCSI or FC has failed	StorNext does not support user firmware updates. Contact the Quantum Technical Assistance Center for upgrade information.
<p>Flag 35 (23h) - Drive Humidity</p> <p>Flag 36 (24h) - Drive Temperature</p> <p>Flag 37 (25h) - Drive Voltage</p> <p>Flag 38 (26h) - Predictive Failure</p> <p>Flag 39 (27h) - Diagnostics Required</p> <p>Flag 40 (28h) - Loader hardware A</p> <p>Flag 41 (29h) - Loader stray tape</p> <p>Flag 42 (2Ah) - Loader Hardware B</p> <p>Flag 43 (2Bh) - Loader door</p> <p>Flag 44 (2Ch) - Loader hardware C</p> <p>Flag 45 (2Dh) - Loader magazine</p> <p>Flag 46 (2Eh) - Loader predictive failure</p>	Contact the Quantum Technical Assistance Center .
Flag 51 (33h) - Tape directory invalid at unload	The tape directory must be rebuilt. Contact the Quantum Technical Assistance Center .
<p>Flag 52 (34h) - Tape system area</p> <p>Flag 53 (35h) - Tape system area read failure</p> <p>Flag 54 (36h) - No start of data</p> <p>Flag 55 (37h) - Loading failure</p>	Contact the Quantum Technical Assistance Center .
The problem <u>IS</u> resolved:	Close the service ticket. Refer to Closing Service Tickets .

Figure 6 Drive Reported Drive Error RAS

Recommended Actions									
Tape Drive - Drive Reported Drive Error									
IF	THEN								
The service ticket indicates the tape drive reported a drive error:	<ol style="list-style-type: none"> 1. Check the tape library's control panel to determine if any other errors exist. <ul style="list-style-type: none"> ○ If other errors exist, correct them before proceeding. Refer to the documentation for this type of tape library. ○ If no other errors exist and the media is mounted, dismount the media. 2. If the media is not dismounted, check the drive to see if it has been ejected. 3. If the media has not been ejected: <ul style="list-style-type: none"> ○ Press the Eject button on the drive to eject the media. ○ Try to dismount the media again. 								
The problem <u>IS</u> resolved:	Close the service ticket. Refer to Closing Service Tickets .								
The problem has <u>NOT</u> been resolved:	<ol style="list-style-type: none"> 1. Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. 2. Contact the Quantum Technical Assistance Center. <table style="margin-left: 20px; border: none;"> <tr> <td>In the USA:</td> <td>1+800-284-5101</td> </tr> <tr> <td>UK, France and Germany:</td> <td>00800 4 QUANTUM</td> </tr> <tr> <td>EMEA:</td> <td>+44 1256 848 766</td> </tr> <tr> <td>On the Web:</td> <td>http://www.quantum.com/support</td> </tr> </table> 	In the USA:	1+800-284-5101	UK, France and Germany:	00800 4 QUANTUM	EMEA:	+44 1256 848 766	On the Web:	http://www.quantum.com/support
In the USA:	1+800-284-5101								
UK, France and Germany:	00800 4 QUANTUM								
EMEA:	+44 1256 848 766								
On the Web:	http://www.quantum.com/support								

Figure 7 Cleaning of Drive
Failed RAS

Recommended Actions

Tape Drive - Cleaning of Drive Failed

IF	THEN
Drive cleaning failed:	<p>The cleaning media might be defective or expired, or there is a problem with the drive.</p> <ol style="list-style-type: none"> 1. Replace existing cleaning media. 2. Attempt to clean the drive using the StorNext GUI. (From the SNSM home page, choose Drive > Clean Drive from the Admin menu.) 3. If the drive still indicates that cleaning is required, contact the Quantum Technical Assistance Center using the contact information below.
The problem <u>IS</u> resolved:	<p>Close the service ticket. Refer to Closing Service Tickets.</p>
The problem has <u>NOT</u> been resolved:	<ol style="list-style-type: none"> 1. Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. 2. Contact the Quantum Technical Assistance Center. In the USA: 1+800-284-5101 UK, France and Germany: 00800 4 QUANTUM EMEA: +44 1256 848 766 On the Web: http://www.quantum.com/support

Figure 8 Wrong Firmware Level/Invalid Drive Type RAS

Recommended Actions	
Tape Drive - Wrong Firmware Level/Invalid Drive Type	
IF	THEN
The service ticket indicates the tape drive's firmware level is wrong:	Contact the Quantum Technical Assistance Center using the contact information below.
The service ticket indicates the drive type is invalid:	Disconnect the drive, and then contact the Quantum Technical Assistance Center using the contact information below.
The problem <u>IS</u> resolved:	Close the service ticket. Refer to Closing Service Tickets .
The problem has <u>NOT</u> been resolved:	<ol style="list-style-type: none"> 1. Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. 2. Contact the Quantum Technical Assistance Center. In the USA: 1+800-284-5101 UK, France and Germany: 00800 4 QUANTUM EMEA: +44 1256 848 766 On the Web: http://www.quantum.com/support <p>OR</p> <ol style="list-style-type: none"> 3. If you are a properly-trained service professional, perform the procedures required for this type of tape library.

Figure 9 Drive Removed RAS

Recommended Actions	
Tape Drive - Drive Removed	
IF	THEN
The service ticket indicates a tape drive was removed:	<ol style="list-style-type: none"> 1. Make sure the connections to the physical drive (for example, fibre, SCSI) are secure. 2. Make sure the physical (tape) library can see the drive.
The problem <u>IS</u> resolved:	Close the service ticket. Refer to Closing Service Tickets .
The problem has <u>NOT</u> been resolved:	<ol style="list-style-type: none"> 1. Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. 2. Contact the Quantum Technical Assistance Center. In the USA: 1+800-284-5101 UK, France and Germany: 00800 4 QUANTUM EMEA: +44 1256 848 766 On the Web: http://www.quantum.com/support

Figure 10 Tape Drive -
Configuration Failed RAS

Recommended Actions

Tape Drive - Configuration Failed

IF	THEN
Drive configuration failed:	<ol style="list-style-type: none"> 1. Capture the StorNext system state. Refer to Capturing a System State. 2. Contact Quantum Technical Assistance Center using the information below.
The problem <u>IS</u> resolved:	Close the service ticket. Refer to Closing Service Tickets .
The problem has <u>NOT</u> been resolved:	<ol style="list-style-type: none"> 1. Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. 2. Contact Quantum Technical Assistance Center. In the USA: 1+800-284-5101 UK, France and Germany: 00800 4 QUANTUM EMEA: +44 1256 848 766 On the Web: http://www.quantum.com/support

Figure 11 Tape Drive -
Reported Media Error RAS

Recommended Actions	
Tape - Drive Reported Media Error	
IF	THEN
The drive reported a media error (sense data, tape alert):	<ol style="list-style-type: none"> 1. Check the tape library's control panel to determine if any other errors exist. <ul style="list-style-type: none"> o If other errors exist, correct them before proceeding. Refer to the documentation for this type of tape library. o If no other errors exist and the media is mounted, dismount the media. 2. If the media is not dismounted, check the drive to see if it has been ejected. 3. If the media has not been ejected: <ul style="list-style-type: none"> o Press the Eject button on the drive to eject the media. o Try to dismount the media again.
The problem <u>IS</u> resolved:	Close the service ticket. Refer to Closing Service Tickets .
The problem has <u>NOT</u> been resolved:	<ol style="list-style-type: none"> 1. Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. 2. Contact the Quantum Technical Assistance Center. <ul style="list-style-type: none"> In the USA: 1+800-284-5101 UK, France and Germany: 00800 4 QUANTUM EMEA: +44 1256 848 766 On the Web: http://www.quantum.com/support

Figure 12 Cleaning Media Expired RAS

Recommended Actions	
Cleaning Media - Expired	
IF	THEN
The service ticket indicates the cleaning media for the tape library has expired:	<ol style="list-style-type: none"> 1. If the tape library has exported the cleaning media to the entry port, remove the cleaning media. 2. If the tape library has NOT exported the cleaning media to the entry port, export it. 3. If no other cleaning media is available in the tape library, add a new one.
The problem IS resolved:	Close the service ticket. Refer to Closing Service Tickets .
The problem has NOT been resolved:	<ol style="list-style-type: none"> 1. Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. 2. Contact the Quantum Technical Assistance Center. In the USA: 1+800-284-5101 UK, France and Germany: 00800 4 QUANTUM EMEA: +44 1256 848 766 On the Web: http://www.quantum.com/support

Figure 13 No Cleaning Media Available RAS

Recommended Actions	
Cleaning Media - No Cleaning Media Available	
IF	THEN
The service ticket indicates the tape library does not have any available cleaning media:	Add new cleaning media to the tape library.
The problem IS resolved:	Close the service ticket. Refer to Closing Service Tickets .
The problem has NOT been resolved:	<ol style="list-style-type: none"> 1. Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. 2. Contact the Quantum Technical Assistance Center. In the USA: 1+800-284-5101 UK, France and Germany: 00800 4 QUANTUM EMEA: +44 1256 848 766 On the Web: http://www.quantum.com/support

Figure 14 Media Suspect
Threshold Count Exceeded RAS

Recommended Actions

Media Suspect Threshold Count Exceeded

StorNext allows media to be marked suspect a certain number of times before the threshold is met. (The default threshold is 3 times.) If the suspect threshold count is exceeded, media is treated as logically write protected.

IF	THEN
If a particular piece of media was marked suspect and exceeded the suspect count threshold:	<ol style="list-style-type: none"> 1. Use the Capturing a System State feature to create a system snapshot. 2. Contact the Quantum Technical Assistance Center.
The problem <u>IS</u> resolved:	Close the service ticket. Refer to Closing Service Tickets .
The problem has <u>NOT</u> been resolved:	<ol style="list-style-type: none"> 1. Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. 2. Contact the Quantum Technical Assistance Center. <ul style="list-style-type: none"> In the USA: 1+800-284-5101 UK, France and Germany: 00800 4 QUANTUM EMEA: +44 1256 848 766 On the Web: http://www.quantum.com/support

Figure 15 Media Format Failure RAS

Recommended Actions

Media Format Failure

IF	THEN
The problem indicates that a particular piece of media failed to format because the volume ID (valid) was already in use:	<ol style="list-style-type: none"> 1. Use the Capturing a System State feature to create a system snapshot. 2. Contact the Quantum Technical Assistance Center immediately.
The problem indicates that a particular piece of media failed to format:	<p>Validate that the problem was the media or the drive:</p> <ol style="list-style-type: none"> 1. Use <code>fschmedstate</code> to clean up the media's <i>suspect and marked</i> state. 2. Use <code>fschstate</code> to take offline the drive in which the media failed to format. 3. Attempt to reformat the media using <code>fsformat</code>. <p>If the format fails again, the media is unusable and should be discarded. Otherwise the previous drive or connectivity might be problematic. Contact the Quantum Technical Assistance Center.</p>
The problem persists:	<ol style="list-style-type: none"> 1. Use the Capturing a System State feature to create a system snapshot. 2. Contact the Quantum Technical Assistance Center.
The problem <u>IS</u> resolved:	Close the service ticket. Refer to Closing Service Tickets .
The problem has <u>NOT</u> been resolved:	<ol style="list-style-type: none"> 1. Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. 2. Contact the Quantum Technical Assistance Center. In the USA: 1+800-284-5101 UK, France and Germany: 00800 4 QUANTUM EMEA: +44 1256 848 766 On the Web: http://www.quantum.com/support

Figure 16 Invalid Media Label
Detected RAS

Recommended Actions

Invalid Media Label Detected

IF	THEN
A specific piece of media fails label validation:	<p>Run <code>fsCheckSlotMapping</code> to ensure that all paths to tape drives are configured properly.</p> <p>If <code>fsCheckSlotMapping</code> fails (indicating that tape drive paths are not configured properly):</p> <ol style="list-style-type: none"> 1. Reboot the server 2. Change the media state to available using <code>fschmedstate</code> 3. Try the failed operation again. <p>If <code>fsCheckSlotMapping</code> runs successfully, determine whether the drive is not the cause of the error:</p> <ol style="list-style-type: none"> 1. Use <code>fschstate</code> to place offline the tape drive for which the tape label verification failed. 2. Use <code>fschmedstate</code> to change the media state to Available. 3. Try the failed operation again. <ul style="list-style-type: none"> • If the operation succeeds: <ul style="list-style-type: none"> - Use <code>fschstate</code> to place the original drive back online - Contact the Quantum Technical Assistance Center about the suspected drive. • If the operation fails again with a label validation failure: <ul style="list-style-type: none"> - Use <code>fschstate</code> to place the original drive back online - Contact the Quantum Technical Assistance Center about the failed media. <p>If the problem persists:</p> <ol style="list-style-type: none"> 1. Use the Capturing a System State feature to create a system snapshot. 2. Contact the Quantum Technical Assistance Center.
The problem <u>IS</u> resolved:	Close the service ticket. Refer to Closing Service Tickets .
The problem has <u>NOT</u> been resolved:	<ol style="list-style-type: none"> 1. Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. 2. Contact the Quantum Technical Assistance Center. <p>In the USA: 1+800-284-5101 UK, France and Germany: 00800 4 QUANTUM EMEA: +44 1256 848 766 On the Web: http://www.quantum.com/support</p>

Figure 17 Media Not Found
RAS

Recommended Actions									
Media Not Found									
IF	THEN								
The system indicates that a particular piece of media is not found (unavailable):	<p>Check the library console for the indicated media.</p> <ul style="list-style-type: none"> If the media is found in the library, issue <code>vsaudit <library></code> If the media is NOT found in the library, locate the indicated media and insert it into the library's mailbox. Import media using the GUI. (From the StorNext home page, choose Remove/Move Media from the Admin menu.) <p>If the problem persists:</p> <ol style="list-style-type: none"> Use the Capturing a System State feature to create a system snapshot. Restart StorNext. Contact the Quantum Technical Assistance Center. 								
The problem <u>IS</u> resolved:	Close the service ticket. Refer to Closing Service Tickets .								
The problem has <u>NOT</u> been resolved:	<ol style="list-style-type: none"> Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. Contact the Quantum Technical Assistance Center. <table border="0"> <tr> <td>In the USA:</td> <td>1+800-284-5101</td> </tr> <tr> <td>UK, France and Germany:</td> <td>00800 4 QUANTUM</td> </tr> <tr> <td>EMEA:</td> <td>+44 1256 848 766</td> </tr> <tr> <td>On the Web:</td> <td>http://www.quantum.com/support</td> </tr> </table> 	In the USA:	1+800-284-5101	UK, France and Germany:	00800 4 QUANTUM	EMEA:	+44 1256 848 766	On the Web:	http://www.quantum.com/support
In the USA:	1+800-284-5101								
UK, France and Germany:	00800 4 QUANTUM								
EMEA:	+44 1256 848 766								
On the Web:	http://www.quantum.com/support								

Figure 18 Duplicate Physical
Media Found RAS

Recommended Actions									
Duplicate Physical Media Found									
IF	THEN								
If the service ticket indicates that duplicate physical media has been found:	Remove the duplicate media using the library's operator panel. Refer to your library's reference manual for operator panel instructions.								
The problem <u>IS</u> resolved:	Close the service ticket. Refer to Closing Service Tickets .								
The problem has <u>NOT</u> been resolved:	<ol style="list-style-type: none"> Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. Contact the Quantum Technical Assistance Center. <table border="0"> <tr> <td>In the USA:</td> <td>1+800-284-5101</td> </tr> <tr> <td>UK, France and Germany:</td> <td>00800 4 QUANTUM</td> </tr> <tr> <td>EMEA:</td> <td>+44 1256 848 766</td> </tr> <tr> <td>On the Web:</td> <td>http://www.quantum.com/support</td> </tr> </table> 	In the USA:	1+800-284-5101	UK, France and Germany:	00800 4 QUANTUM	EMEA:	+44 1256 848 766	On the Web:	http://www.quantum.com/support
In the USA:	1+800-284-5101								
UK, France and Germany:	00800 4 QUANTUM								
EMEA:	+44 1256 848 766								
On the Web:	http://www.quantum.com/support								

Figure 19 Storage Disk Taken Offline RAS

Recommended Actions	
Storage Disk Taken Offline	
IF	THEN
A storage disk exceeds its failure threshold and is taken offline.	<ol style="list-style-type: none"> 1. Verify that the file system can be reached (NFS), and is still mounted and accessible. 2. If the storage disk is located on a CFVS file system, check the File System Manager (FSM).
A deduplication-enabled storage disk is taken offline:	<ol style="list-style-type: none"> 1. Try bringing the dedup sdisk back online. 2. If this doesn't work, run the health checks on the media to validate the blockpool used for deduplication. After verifying, bring the dedup sdisk back online. 3. If the storage disk is located on a CFVS file system, check the File System Manager (FSM). <p>(The above recommended actions apply only to deduplication-enabled storage disks.)</p>
The problem IS resolved:	Close the service ticket. Refer to Closing Service Tickets .
The problem has NOT been resolved:	<ol style="list-style-type: none"> 1. Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. 2. Contact the Quantum Technical Assistance Center. In the USA: 1+800-284-5101 UK, France and Germany: 00900 4 QUANTUM EMEA: +44 1256 848 766 On the Web: http://www.quantum.com/support

SNFS RAS Messages

This section describes RAS messages that might appear as a result of a file system-related error condition, such as an I/O error or a missing LUN.

Figure 20 Configuration Not Supported RAS

Recommended Actions	
<h3>Configuration Not Supported</h3>	
IF	THEN
The file system configuration file is corrupt, missing, or causes a syntax error to be reported:	Verify that a valid file system configuration file exists for the specified file system. Also, check the system logs for additional configuration file error details.
The problem <u>IS</u> resolved:	Close the service ticket. Refer to Closing Service Tickets .
The problem has <u>NOT</u> been resolved:	<ol style="list-style-type: none"> 1. Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. 2. Contact the QuantumTechnical Assistance Center. In the USA: 1+800-284-5101 UK, France and Germany: 00800 4 QUANTUM EMEA: +44 1256 848 766 On the Web: http://www.quantum.com/support

Figure 21 Label Validation Failure RAS

Recommended Actions	
<h3>Label Validation Failure</h3>	
IF	THEN
Disk label verification has failed:	Use the cvlabel command to check for corrupt, incorrect, or missing disk labels. Also inspect system logs for I/O errors, and check SAN integrity.
The problem <u>IS</u> resolved:	Close the service ticket. Refer to Closing Service Tickets .
The problem has <u>NOT</u> been resolved:	<ol style="list-style-type: none"> 1. Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. 2. Contact the QuantumTechnical Assistance Center. In the USA: 1+800-284-5101 UK, France and Germany: 00800 4 QUANTUM EMEA: +44 1256 848 766 On the Web: http://www.quantum.com/support

Figure 22 Connection Rejected
RAS

Recommended Actions									
<h2>Connection Rejected</h2>									
IF	THEN								
A client connection has been rejected unexpectedly:	Check the system logs to determine the root cause. If the problem is caused by exceeding the maximum number of connections, increase MaxConnections in the file system configuration file.								
The problem <u>IS</u> resolved:	Close the service ticket. Refer to Closing Service Tickets .								
The problem has <u>NOT</u> been resolved:	<ol style="list-style-type: none"> 1. Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. 2. Contact the QuantumTechnical Assistance Center. <table border="0"> <tr> <td>In the USA:</td> <td>1+800-284-5101</td> </tr> <tr> <td>UK, France and Germany:</td> <td>00800 4 QUANTUM</td> </tr> <tr> <td>EMEA:</td> <td>+44 1256 848 766</td> </tr> <tr> <td>On the Web:</td> <td>http://www.quantum.com/support</td> </tr> </table> 	In the USA:	1+800-284-5101	UK, France and Germany:	00800 4 QUANTUM	EMEA:	+44 1256 848 766	On the Web:	http://www.quantum.com/support
In the USA:	1+800-284-5101								
UK, France and Germany:	00800 4 QUANTUM								
EMEA:	+44 1256 848 766								
On the Web:	http://www.quantum.com/support								

Figure 23 File System Failover
RAS

Recommended Actions									
<h2>File System Failover</h2>									
IF	THEN								
A file system failed over unexpectedly:	Inspect the system log and the FSM cvlog to determine the root cause.								
The problem <u>IS</u> resolved:	Close the service ticket. Refer to Closing Service Tickets .								
The problem has <u>NOT</u> been resolved:	<ol style="list-style-type: none"> 1. Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. 2. Contact the QuantumTechnical Assistance Center. <table border="0"> <tr> <td>In the USA:</td> <td>1+800-284-5101</td> </tr> <tr> <td>UK, France and Germany:</td> <td>00800 4 QUANTUM</td> </tr> <tr> <td>EMEA:</td> <td>+44 1256 848 766</td> </tr> <tr> <td>On the Web:</td> <td>http://www.quantum.com/support</td> </tr> </table> 	In the USA:	1+800-284-5101	UK, France and Germany:	00800 4 QUANTUM	EMEA:	+44 1256 848 766	On the Web:	http://www.quantum.com/support
In the USA:	1+800-284-5101								
UK, France and Germany:	00800 4 QUANTUM								
EMEA:	+44 1256 848 766								
On the Web:	http://www.quantum.com/support								

Figure 24 I/O Error RAS

Recommended Actions									
I/O Error									
IF	THEN								
An I/O error has occurred:	Check LUN and disk path health, as well as overall SAN integrity. Also, inspect the system logs for driver-level I/O errors.								
The problem IS resolved:	Close the service ticket. Refer to Closing Service Tickets .								
The problem has NOT been resolved:	<ol style="list-style-type: none"> 1. Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. 2. Contact the QuantumTechnical Assistance Center. <table border="0"> <tr> <td>In the USA:</td> <td>1+800-284-5101</td> </tr> <tr> <td>UK, France and Germany:</td> <td>00800 4 QUANTUM</td> </tr> <tr> <td>EMEA:</td> <td>+44 1256 848 766</td> </tr> <tr> <td>On the Web:</td> <td>http://www.quantum.com/support</td> </tr> </table> 	In the USA:	1+800-284-5101	UK, France and Germany:	00800 4 QUANTUM	EMEA:	+44 1256 848 766	On the Web:	http://www.quantum.com/support
In the USA:	1+800-284-5101								
UK, France and Germany:	00800 4 QUANTUM								
EMEA:	+44 1256 848 766								
On the Web:	http://www.quantum.com/support								

Figure 25 Journaling Error Detected RAS

Recommended Actions									
Journaling Error Detected									
IF	THEN								
Journal recovery has failed:	Contact the Quantum Technical Assistance Center and open a service request.								
The problem IS resolved:	Close the service ticket. Refer to Closing Service Tickets .								
The problem has NOT been resolved:	<ol style="list-style-type: none"> 1. Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. 2. Contact the QuantumTechnical Assistance Center. <table border="0"> <tr> <td>In the USA:</td> <td>1+800-284-5101</td> </tr> <tr> <td>UK, France and Germany:</td> <td>00800 4 QUANTUM</td> </tr> <tr> <td>EMEA:</td> <td>+44 1256 848 766</td> </tr> <tr> <td>On the Web:</td> <td>http://www.quantum.com/support</td> </tr> </table> 	In the USA:	1+800-284-5101	UK, France and Germany:	00800 4 QUANTUM	EMEA:	+44 1256 848 766	On the Web:	http://www.quantum.com/support
In the USA:	1+800-284-5101								
UK, France and Germany:	00800 4 QUANTUM								
EMEA:	+44 1256 848 766								
On the Web:	http://www.quantum.com/support								

Figure 26 SNFS License
Required RAS

Recommended Actions	
<h3>SNFS License Required</h3>	
IF	THEN
You receive a warning that your SNFS license will expire within 48 hours:	Contact the Quantum Technical Assistance Center to obtain a valid license.
The problem <u>IS</u> resolved:	Close the service ticket. Refer to Closing Service Tickets .
The problem has <u>NOT</u> been resolved:	<ol style="list-style-type: none"> 1. Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. 2. Contact the Quantum Technical Assistance Center. In the USA: 1+800-284-5101 UK, France and Germany: 00800 4 QUANTUM EMEA: +44 1256 848 766 On the Web: http://www.quantum.com/support

Figure 27 SNFS License Failure
RAS

Recommended Actions	
<h3>SNFS License Failure</h3>	
IF	THEN
You receive a warning that your SNFS license will expire within 48 hours:	Contact the Quantum Technical Assistance Center to obtain a valid license.
OR Your SNFS license has expired:	
The problem <u>IS</u> resolved:	Close the service ticket. Refer to Closing Service Tickets .
The problem has <u>NOT</u> been resolved:	<ol style="list-style-type: none"> 1. Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. 2. Contact the Quantum Technical Assistance Center. In the USA: 1+800-284-5101 UK, France and Germany: 00800 4 QUANTUM EMEA: +44 1256 848 766 On the Web: http://www.quantum.com/support

Figure 28 LUN Mapping Changed RAS

Recommended Actions									
<h2>LUN Mapping Changed</h2>									
IF	THEN								
A disk scan has detected a change in an existing LUN path:	If the LUN mapping change is unexpected, run the cvadmin "disks" and "paths" commands to confirm that all LUN paths are present. Also, check SAN integrity and inspect the system logs to determine the root cause.								
The problem <u>IS</u> resolved:	Close the service ticket. Refer to Closing Service Tickets .								
The problem has <u>NOT</u> been resolved:	<ol style="list-style-type: none"> 1. Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. 2. Contact the QuantumTechnical Assistance Center. <table border="0"> <tr> <td>In the USA:</td> <td>1+800-284-5101</td> </tr> <tr> <td>UK, France and Germany:</td> <td>00800 4 QUANTUM</td> </tr> <tr> <td>EMEA:</td> <td>+44 1256 848 766</td> </tr> <tr> <td>On the Web:</td> <td>http://www.quantum.com/support</td> </tr> </table> 	In the USA:	1+800-284-5101	UK, France and Germany:	00800 4 QUANTUM	EMEA:	+44 1256 848 766	On the Web:	http://www.quantum.com/support
In the USA:	1+800-284-5101								
UK, France and Germany:	00800 4 QUANTUM								
EMEA:	+44 1256 848 766								
On the Web:	http://www.quantum.com/support								

Figure 29 Communication Failure RAS

Recommended Actions									
<h2>Communication Failure</h2>									
IF	THEN								
A client has disconnected unexpectedly:	Check the health of the network used for metadata traffic. Also, inspect the FSM log and the system logs on the clients and metadata controller to determine the root cause.								
The problem <u>IS</u> resolved:	Close the service ticket. Refer to Closing Service Tickets .								
The problem has <u>NOT</u> been resolved:	<ol style="list-style-type: none"> 1. Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. 2. Contact the QuantumTechnical Assistance Center. <table border="0"> <tr> <td>In the USA:</td> <td>1+800-284-5101</td> </tr> <tr> <td>UK, France and Germany:</td> <td>00800 4 QUANTUM</td> </tr> <tr> <td>EMEA:</td> <td>+44 1256 848 766</td> </tr> <tr> <td>On the Web:</td> <td>http://www.quantum.com/support</td> </tr> </table> 	In the USA:	1+800-284-5101	UK, France and Germany:	00800 4 QUANTUM	EMEA:	+44 1256 848 766	On the Web:	http://www.quantum.com/support
In the USA:	1+800-284-5101								
UK, France and Germany:	00800 4 QUANTUM								
EMEA:	+44 1256 848 766								
On the Web:	http://www.quantum.com/support								

Figure 30 Metadata Inconsistency Detected RAS

Recommended Actions									
Metadata Inconsistency Detected									
IF	THEN								
The FSM has detected a metadata inconsistency:	Check SAN integrity and inspect the system logs for I/O errors. If the SAN is healthy, run <code>cvfsck</code> on the affected file system at the earliest convenient opportunity.								
The problem <u>IS</u> resolved:	Close the service ticket. Refer to Closing Service Tickets .								
The problem has <u>NOT</u> been resolved:	<ol style="list-style-type: none"> 1. Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. 2. Contact the QuantumTechnical Assistance Center. <table border="0"> <tr> <td>In the USA:</td> <td>1+800-284-5101</td> </tr> <tr> <td>UK, France and Germany:</td> <td>00800 4 QUANTUM</td> </tr> <tr> <td>EMEA:</td> <td>+44 1256 848 766</td> </tr> <tr> <td>On the Web:</td> <td>http://www.quantum.com/support</td> </tr> </table> 	In the USA:	1+800-284-5101	UK, France and Germany:	00800 4 QUANTUM	EMEA:	+44 1256 848 766	On the Web:	http://www.quantum.com/support
In the USA:	1+800-284-5101								
UK, France and Germany:	00800 4 QUANTUM								
EMEA:	+44 1256 848 766								
On the Web:	http://www.quantum.com/support								

Figure 31 Bad File System Metadata Dump RAS

Recommended Actions									
Bad File System Metadata Dump									
IF	THEN								
The system has detected that a new metadata dump is required:	Run <code>snmetadump</code> for the affected file system as soon as possible. Note: This condition could occur if <code>cvfsck</code> or <code>cvupdatefs</code> was recently run, or if a Restore Journal error occurred and the Restore Journal was shut down.								
The problem <u>IS</u> resolved:	Close the service ticket. Refer to Closing Service Tickets .								
The problem has <u>NOT</u> been resolved:	<ol style="list-style-type: none"> 1. Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. 2. Contact the QuantumTechnical Assistance Center. <table border="0"> <tr> <td>In the USA:</td> <td>1+800-284-5101</td> </tr> <tr> <td>UK, France and Germany:</td> <td>00800 4 QUANTUM</td> </tr> <tr> <td>EMEA:</td> <td>+44 1256 848 766</td> </tr> <tr> <td>On the Web:</td> <td>http://www.quantum.com/support</td> </tr> </table> 	In the USA:	1+800-284-5101	UK, France and Germany:	00800 4 QUANTUM	EMEA:	+44 1256 848 766	On the Web:	http://www.quantum.com/support
In the USA:	1+800-284-5101								
UK, France and Germany:	00800 4 QUANTUM								
EMEA:	+44 1256 848 766								
On the Web:	http://www.quantum.com/support								

Figure 32 Metadata Dump Failure RAS

Recommended Actions	
Metadata Dump Failure	
IF	THEN
The system has detected either a stale or missing metadata dump for a managed file system.	<p>StorNext backup and restore operations (and also some file system scanning operations such as a rebuild policy.) require the existence of a current, valid metadata dump. Use the following procedure to perform a metadata dump.</p> <p>Unmount the system:</p> <ol style="list-style-type: none"> 1. From the SNFS Home Page, choose Unmount from the Admin menu. 2. Select from the Mounted File Systems list the file system to unmount. 3. Click Unmount. <p>Stop the file system:</p> <ol style="list-style-type: none"> 1. From the SNFS Home Page, choose Start/Stop File System from the Admin menu. 2. Select from the Active File Systems list the file system you want to stop. 3. Click Stop. <p>Perform the metadata dump:</p> <ol style="list-style-type: none"> 1. From the SNFS Home Page, choose Metadata Dump from the Admin menu. 2. Select the file system on which to perform the metadata dump. 3. Click Apply. <p>After the dump is complete, restart and mount the file system.</p>
The problem IS resolved:	Close the service ticket. Refer to Closing Service Tickets .
The problem has NOT been resolved:	<ol style="list-style-type: none"> 1. Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. 2. Contact the Quantum Technical Assistance Center. <ul style="list-style-type: none"> In the USA: 1+800-284-5101 UK, France and Germany: 00800 4 QUANTUM EMEA: +44 1256 848 766 On the Web: http://www.quantum.com/support

Figure 33 File System or Metadata Capacity Warning RAS

Recommended Actions	
File System or Metadata Capacity Warning	
IF	THEN
You receive a warning about your file system exceeding <code>FsCapacityThreshold</code> :	Add additional storage capacity or reduce file system usage.
You receive a warning that the file system is running out of metadata capacity:	Add additional metadata storage capacity.
The problem IS resolved:	Close the service ticket. Refer to Closing Service Tickets .
The problem is NOT resolved:	<ol style="list-style-type: none"> 1. Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. 2. Contact the Quantum Technical Assistance Center. <ul style="list-style-type: none"> In the USA: 1+800-284-5101 UK, France and Germany: 00800 4 QUANTUM EMEA: +44 1256 848 766 On the Web: http://www.quantum.com/support

Figure 34 File Processing Failure RAS

Recommended Actions	
File Processing Failure	
IF	THEN
A failure occurred while trying to process an internal file.	<p>See the error details for more complete information about the failure.</p> <p>Possible reasons for the failure:</p> <ul style="list-style-type: none"> • An attempt to roll the file (close the current file and open a new one for use) failed • A corruption in the file was detected <p>In general, the system can continue without intervention after one of these errors occurs. However, if you experience these failures on a regular basis it could be indicative of a more serious situation, and you should contact the Quantum Technical Assistance Center.</p>
The problem <u>IS</u> resolved:	Close the service ticket. Refer to Closing Service Tickets .
The problem has <u>NOT</u> been resolved:	<ol style="list-style-type: none"> 1. Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. 2. Contact the Quantum Technical Assistance Center. In the USA: 1+800-284-5101 UK, France and Germany: 00800 4 QUANTUM EMEA: +44 1256 848 766 On the Web: http://www.quantum.com/support

Figure 35 Missing LUNs RAS

Recommended Actions	
Missing LUNs	
IF	THEN
A client fails to mount because a LUN is missing:	<p>Check the system logs to determine the root cause.</p> <p>Run the cvadmin "disks" and "paths" commands, and then check for missing LUNs.</p>
The problem <u>IS</u> resolved:	Close the service ticket. Refer to Closing Service Tickets .
The problem has <u>NOT</u> been resolved:	<ol style="list-style-type: none"> 1. Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. 2. Contact the Quantum Technical Assistance Center. In the USA: 1+800-284-5101 UK, France and Germany: 00800 4 QUANTUM EMEA: +44 1256 848 766 On the Web: http://www.quantum.com/support

Figure 36 Disk Space Allocation Failure RAS

Recommended Actions

Disk Space Allocation Failure

IF	THEN
A disk space allocation has failed:	Free up disk space by removing unnecessary disk copies of files, or add disk capacity. If the allocation failure is unexpected, contact the Quantum Technical Assistance Center.
The problem IS resolved:	Close the service ticket. Refer to Closing Service Tickets .
The problem has NOT been resolved:	<ol style="list-style-type: none"> 1. Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. 2. Contact the Quantum Technical Assistance Center. <p>In the USA: 1+800-284-5101 UK, France and Germany: 00800 4 QUANTUM EMEA: +44 1256 848 766 On the Web: http://www.quantum.com/support</p>

Figure 37 System Resource Failure RAS

Recommended Actions

System Resource Failure

IF	THEN
SNFS has failed to allocate memory:	Determine the cause of memory depletion and correct the condition by adding memory or paging space to your system. If SNFS is using excessive amounts of memory, adjusting the configuration parameters might resolve the problem. For information about adjusting parameters, refer to the Release Notes, the <code>cvfs_config(4)</code> and <code>mount_cvfs(1)</code> man pages, and the SNFS Tuning Guide.
The FSM detects exhaustion of a resource controlled by an adjustable parameter:	Modify the file system configuration file as needed, and then restart the file system.
The problem IS resolved:	Close the service ticket. Refer to Closing Service Tickets .
The problem has NOT been resolved:	<ol style="list-style-type: none"> 1. Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. 2. Contact the Quantum Technical Assistance Center. <p>In the USA: 1+800-284-5101 UK, France and Germany: 00800 4 QUANTUM EMEA: +44 1256 848 766 On the Web: http://www.quantum.com/support</p>

Figure 38 Affinity Configuration Violations RAS

Recommended Actions

Affinity Configuration Violations

When a configuration violation occurs in the StorNext application, it must be repaired by stopping the system, editing the configuration, and then restarting the system. Below are specific configuration violations and recommended actions to repair each specific issue.

IF	THEN
There is more than one affinity on one stripe group:	<p>You cannot have more than one affinity on one stripe group.</p> <p>Examine all DSM configuration files (/usr/adc/dsm/config/*.cfg). In any file that has the Data Migration Flag set to YES, and for every stripe group with more than one affinity, remove the extra affinities.</p>
The file system does not contain at least one non-exclusive data stripe group:	<p>The file system has at least one affinity, and therefore must contain at least one non-exclusive data stripe group.</p> <p>Examine all DSM configuration files (/usr/adc/dsm/config/*.cfg). In any file that has the Data Migration Flag set to YES, make sure at least one stripe group has the following configuration:</p> <pre> Metadata No Journal No Exclusive No </pre>
A file system contains both data stripe groups with affinities and data stripe groups without affinities:	<p>A file system can contain data stripe groups with affinities, or data stripe groups without affinities, but it cannot contain both.</p> <p>Examine all DSM configuration files (/usr/adc/dsm/config/*.cfg). In any file that has the Data Migration Flag set to YES, make sure that either every stripe group has an affinity, or that every stripe group does not have an affinity.</p>
There are more than two affinities across all managed file systems:	<p>No more than two affinities across all managed file systems are allowed.</p> <p>Examine all DSM configuration files (/usr/adc/dsm/config/*.cfg). For all of the configuration files that have the Data Migration Flag set to YES, change the stripe group Affinities so there are no more than a total of two.</p>
The number of affinities on managed file systems do not match for TSM and CVFS:	<p>Examine all DSM configuration files (/usr/adc/dsm/config/*.cfg). For all of the configuration files that have the Data Migration Flag set to YES, make sure the complete list of affinities matches the TSM affinity names found in the TIERDEF database table.</p>
TSM does not recognize the CVFS managed file system affinity name:	<p>Examine all DSM configuration files (/usr/adc/dsm/config/*.cfg). For all of the configuration files that have the Data Migration Flag set to YES, make sure the complete list of affinities matches the TIERNAME fields found in the TSM TIERDEF database table.</p>
An affinity in a policy class is not found in the TIERDEF table:	<p>Make sure the non-zero elements in the TIERLIST field of the TSM CLASSDEF tables all match the TIERNUM fields in the TSM TIERDEF table.</p>
The problem IS resolved:	<p>Close the service ticket. Refer to Closing Service Tickets.</p>
The problem has NOT been resolved:	<ol style="list-style-type: none"> Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. Contact the Quantum Technical Assistance Center. <p>In the USA: 1+800-284-5101 UK, France and Germany: 00800 4 QUANTUM EMEA: +44 1256 848 766 On the Web: http://www.quantum.com/support</p>

Figure 39 Quota Limit or Fragmentation Warnings RAS

Recommended Actions	
Quota Limit or Fragmentation Warnings	
IF	THEN
You receive a warning that the quota hard limit is reached for a user:	Either increase the user's quota, or notify the user.
You receive a warning that fragmentation has been detected in an inode:	<ol style="list-style-type: none"> 1. Consult the snfsdefrag man page for instructions on performing fragmentation analysis and defragmenting files. 2. See ExtentCountThreshold in the cvfs_config documentation for information on adjusting this RAS event.
The problem <u>IS</u> resolved:	Close the service ticket. Refer to Closing Service Tickets .
The problem is <u>NOT</u> resolved:	<ol style="list-style-type: none"> 1. Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. 2. Contact the Quantum Technical Assistance Center. <ul style="list-style-type: none"> In the USA: 1+800-284-5101 UK, France and Germany: 00800 4 QUANTUM EMEA: +44 1256 848 766 On the Web: http://www.quantum.com/support

Figure 40 Shutdown Error RAS

Recommended Actions	
Shutdown Error	
IF	THEN
SNFS shutdown errors have occurred:	Inspect the file system and system logs to determine the root cause.
The problem <u>IS</u> resolved:	Close the service ticket. Refer to Closing Service Tickets .
The problem has <u>NOT</u> been resolved:	<ol style="list-style-type: none"> 1. Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. 2. Contact the Quantum Technical Assistance Center. <ul style="list-style-type: none"> In the USA: 1+800-284-5101 UK, France and Germany: 00800 4 QUANTUM EMEA: +44 1256 848 766 On the Web: http://www.quantum.com/support

Figure 41 Initialization Failure
RAS

Recommended Actions	
Initialization Failure	
IF	THEN
An FSM or FSMPM process has failed to start.	Correct the system configuration as suggested by the event detail, or examine system logs to determine the root cause.
OR	If the detail text suggests a problem with starting the fsmmpm process, run "cvlabel -f" to verify that disk scanning is working properly.
An attempt to mount an SNFS file system has failed:	
All CVFS file systems are not mounted:	Mount the file systems that are not mounted.
The problem <u>IS</u> resolved:	Close the service ticket. Refer to Closing Service Tickets .
The problem has <u>NOT</u> been resolved:	<ol style="list-style-type: none"> 1. Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. 2. Contact the Quantum Technical Assistance Center. In the USA: 1+800-284-5101 UK, France and Germany: 00800 4 QUANTUM EMEA: +44 1256 848 766 On the Web: http://www.quantum.com/support

Other RAS Messages

This section describes RAS messages that might appear as a result of an error condition that is not related to media or the file system.

Figure 42 Checksum Error RAS

Recommended Actions

Checksum Error

When a checksum error occurs during a file retrieve operation, the error is generally due to a hardware failure in a tape drive, host bus adapter, or the cabling between them. The error can also be caused by damaged media. If the checksum error is due to a drive or media failure, there might be an associated "Tape Alert" service ticket.

IF	THEN
The drive or media is suspected:	<ol style="list-style-type: none"> 1. Close the ticket and retry the read/write operation on the original drive and media. 2. Monitor operation for a reoccurrence of the ticket. 3. Insert the suspect media into an alternate drive and retry the read/write operation. 4. If the error follows the media, retire the media. 5. If the error stays with the drive, contact the Quantum Technical Assistance Center to replace the drive.
The media is bad:	<ol style="list-style-type: none"> 1. Copy the data to another piece of media. 2. Remove the original media from the library and discard.
The host bus adapter is suspected:	<ol style="list-style-type: none"> 1. Check the host's system log for HBA-related errors. 2. Replace the HBA with a spare. 3. Close the ticket and retry the read/write operation on the original drive and media. 4. Monitor operation for a reoccurrence of the ticket. 5. If the problem is unresolved, contact the Quantum Technical Assistance Center.
Cabling is suspected:	<ol style="list-style-type: none"> 1. Check the cabling between the drive and the host bus adapter. 2. If the problem is unresolved, contact the Quantum Technical Assistance Center.
There is an associated "Tape Alert" service ticket:	<ol style="list-style-type: none"> 1. Follow the instructions in the Tape Alert service ticket. 2. If the problem is unresolved, contact the Quantum Technical Assistance Center.
The problem <u>IS</u> resolved:	Close the service ticket. Refer to Closing Service Tickets .
The problem has <u>NOT</u> been resolved:	<ol style="list-style-type: none"> 1. Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. 2. Contact the Quantum Technical Assistance Center. <ul style="list-style-type: none"> In the USA: 1+800-284-5101 UK, France and Germany: 00800 4 QUANTUM EMEA: +44 1256 848 766 On the Web: http://www.quantum.com/support

Figure 43 Troubleshooting the StorNext Software RAS

Recommended Actions	
Troubleshooting the StorNext Software	
IF	THEN
A service ticket indicates software issues including incorrect firmware levels:	<ol style="list-style-type: none"> Capture the StorNext system state. Refer to Capturing a System State. Download the captured state. Refer to Downloading a System State Capture. Stop and restart the StorNext software.
An I/O error occurs on a path in a multipath environment (LUN communication failure):	Check the system and RAID logs for SAN integrity.
A process or task dies and does not restart:	Check the FSM logs and system logs to determine the root cause. If possible, take corrective action. If you suspect a software bug, contact the Quantum Technical Assistance Center.
You receive a message that the FSM is delayed or the file system is not responding:	Verify that the FSM process for the specified file system is running on the metadata controller. Also check the health of the metadata network.
A health check operation is launched with invalid command arguments:	<p>Contact the Quantum Technical Assistance Center.</p> <p>To temporarily disable the invalid health check operation until a fix is delivered, follow these steps:</p> <ol style="list-style-type: none"> Locate the 'filelist' file under /usr/adic that contains the health check entry for the failed operation. Place a '#' character at the front of that health check entry line.
<p>You receive an "Internal Software Error" message describing one of the following conditions:</p> <ul style="list-style-type: none"> Process initialization failed A database operation failed An unhandled software error has occurred A CLI command failure has caused the process that invoked it to abort A Blockpool Verify command failed to complete successfully 	<ol style="list-style-type: none"> If you suspect there might have been a temporary problem with system resources, rerun the health check operation via the StorNext GUI's Health Check Service. If the problem persists, contact the Quantum Technical Assistance Center.
<p>You receive one of the following "TSM Control Error" messages:</p> <ul style="list-style-type: none"> The TSM software could not be started The TSM software could not be stopped 	<ol style="list-style-type: none"> Try restarting the TSM software. (On the command line, run "tmsstop; tmsstart") If restarting succeeds, rerun the Health Check using the StorNext GUI. If restarting does not succeed, contact the Quantum Technical Assistance Center.
You cannot connect to the SNAPi server process:	<ol style="list-style-type: none"> Restart the system software. Run the Health Check service using the StorNext GUI. If the problem persists, contact the Quantum Technical Assistance Center.
The problem <u>IS</u> resolved:	Close the service ticket. Refer to Closing Service Tickets .
The problem is <u>NOT</u> resolved:	<ol style="list-style-type: none"> Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. Contact the Quantum Technical Assistance Center. <p style="margin-left: 40px;"> In the USA: 1+800-284-5101 UK, France and Germany: 00800 4 QUANTUM EMEA: +44 1256 848 766 On the Web: http://www.quantum.com/support </p>

Figure 44 Software Resource Violations RAS

Recommended Actions	
Software Resource Warning	
IF	THEN
If the service ticket indicates that no media is found to satisfy the request:	<ul style="list-style-type: none"> Add more media to StorNext <p>OR</p> <ul style="list-style-type: none"> Check for suspect media
A StorNext file system has exceeded the maximum allowable threshold for percent used:	Identify any obsolete or unneeded files in that file system. Either delete those files or move them to a different file system.
A tape drive is in an unusable state:	<ol style="list-style-type: none"> If this is a new or temporary problem, run the <code>execute</code> command to change the tape drive state. If the problem persists, restart StorNext. If the problem still persists, contact the Quantum Technical Assistance Center.
A tape drive is off-line:	<ol style="list-style-type: none"> If this is a new or temporary problem, run the <code>execute</code> command to change the tape drive state. If the problem keeps occurring, contact the Quantum Technical Assistance Center.
Not enough media are available for TSM Data Policies:	Add more media to the StorNext system.
A policy class does not store automatically and is not scheduled:	<ul style="list-style-type: none"> Change the policy class so that 'Store Automatically' is set to 'Yes' <p>OR</p> <ul style="list-style-type: none"> Add the policy to the list of scheduled features, specifying the <policyClassName> as the Option
A device path could not be resolved:	<ol style="list-style-type: none"> Verify that the device is powered on. Verify that the device has not registered any errors. Verify that the device is physically connected by checking all cabling and connections. Verify that the device is correctly mapped to the host server. (Caution: if your system configuration allows you to logically remap all devices, ensure that StorNext has been stopped prior to beginning remapping. Restart StorNext after remapping is complete.)
The problem IS resolved:	Close the service ticket. Refer to Closing Service Tickets
The problem has NOT been resolved:	<ol style="list-style-type: none"> Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. Contact the Quantum Technical Assistance Center. In the USA: 1+800-284-5101 UK, France and Germany: 00800 4 QUANTUM EMEA: +44 1256 848 766 On the Web: http://www.quantum.com/support

Figure 45 Closing Service Tickets RAS

Closing Service Tickets

Use this procedure to close a service ticket.

 **NOTE:** You can analyze a service ticket after it has been closed. For more information, refer to [Analyzing Service Tickets](#).

1. Analyze the service ticket.
Refer to [Analyzing Service Tickets](#).
2. Select the **Close Ticket** check box.
3. Click **Apply** to close the service ticket.

The Progress window appears, showing the status of the service ticket being closed.

 **NOTE:** Once the Progress window appears, you cannot cancel or stop this action. However, you can close the window by clicking on the X in the upper-right corner of the window, but confirmation of success or failure is NOT shown.

IF	THEN
The Progress window shows Success :	The service ticket was successfully closed.
The Progress window shows Failure :	The service ticket was NOT closed. To view the troubleshooting procedures, click View Recommended Actions . To view the error details, click Error Logs for information on why the ticket was not closed. To close a service ticket, repeat Step 1 through Step 3.

4. Click **Close** to close the Progress window.

Figure 46 Analyzing Service Tickets RAS

Analyzing Service Tickets

Use this procedure to add information to a service ticket related to system troubleshooting, and to view the current status of a problem reported by StorNext. All modified entries are kept with the ticket number and ticket summary when the service ticket is closed.

1. Open and view a service ticket.

Refer to [Viewing Service Tickets](#).

2. Click **Analysis**.

The **Ticket Analysis** screen appears.

3. Enter all relevant information regarding actions taken to resolve the issue, and then click **Apply**.

The Progress window appears, showing the status of the ticket being modified.



NOTE: Once the Progress window appears, you cannot cancel or stop this action. However, you can close the window by clicking the **X** in the upper-right corner of the window, but confirmation of success or failure is NOT shown.

IF	THEN
The Progress window shows Success :	The service ticket was successfully modified.
The Progress window shows Failure :	The ticket was NOT modified. To view the troubleshooting procedures, click View Recommended Actions . To view the error details, click Error Logs for information on why the ticket was not modified. To modify a service ticket, repeat Step 1 through Step 3.

4. Click **Close** to close the Progress window.

Figure 47 Viewing Service Tickets RAS

Viewing Service Tickets

View StorNext service tickets to view details of the System Status notification and a suggested resolution of the reported problem.

Do one of the following:

- o Click **System Status** at the bottom of the screen.

OR

- o Access the StorNext home page and choose **System Status** from the **Service** menu.

The **Tools - System Status** screen appears.

- **Ticket** – View service ticket numbers, which can be listed in ascending or descending order by clicking the top of the column.
- **State** – View the service ticket's current state. The state can be either **Open** or **Closed**.
- **Last Updated** – This is the date when the service ticket was last accessed, either opened or closed depending on the selected sort order.
- **Summary** – View a summary of the problem reported by the StorNext.

2. Scroll through the list of service tickets, select the ticket to be viewed, and click **Details**.

The **Ticket Details** screen appears. This screen details the service ticket number, date and time when the ticket was last accessed (either opened or closed), ticket status, and problem description.

3. Click **Cancel** to close the screen.



NOTE: For information on analyzing service tickets and obtaining additional information about a reported problem, refer to [Analyzing Service Tickets](#).

Figure 48 Vault Failure RAS

Recommended Actions

Vault Failure

IF	THEN
The problem indicates that vaulting has failed:	<ol style="list-style-type: none"> 1. Capture the StorNext system state. Refer to Capturing a System State. 2. Download the captured state to a local or network drive. Refer to Downloading a System State Capture. 3. Contact the Quantum Technical Assistance Center. In the USA: 1+800-284-5101 UK, France and Germany: 00800 4 QUANTUM EMEA: +44 1256 848 766 On the Web: http://www.quantum.com/support

Figure 49 Robotics - Not Ready
RAS

Recommended Actions	
Robotics - Not Ready	
IF	THEN
The service ticket indicates that the tape library's robotics is not ready:	<ol style="list-style-type: none"> 1. Verify that the tape library is online and ready. 2. Verify that the tape library is online and ready through the StorNext GUI. 3. Verify that the tape library is connected to the server.
No archives or storage disks exist:	Use the StorNext GUI to add at least one physical archive or storage disk to the system.
An archive is off-line:	<ol style="list-style-type: none"> 1. Use <code>vsarchivevray</code> to change the archive state to online. 2. Run <code>fsstate</code> and verify that all drives are listed correctly. In particular, no drive state should be listed as 'Unknown.' 3. If running <code>fsstate</code> shows 'Unknown' for every drive state, do the following: <ul style="list-style-type: none"> o Run <code>tsmstop; tsmstart</code> to reinitialize software communication pathways. o Repeat step 2 above.
The problem IS resolved:	Close the service ticket. Refer to Closing Service Tickets .
The problem has NOT been resolved:	<ol style="list-style-type: none"> 1. Note any codes displayed on the tape library's control panel. 2. Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. 3. Contact the Quantum Technical Assistance Center. In the USA: 1+800-284-5101 UK, France and Germany: 00800 4 QUANTUM EMEA: +44 1256 848 766 On the Web: http://www.quantum.com/support <p>OR</p> <ol style="list-style-type: none"> 4. If you are a properly-trained service professional, perform the procedures required for this type of tape library.

Figure 50 Robotics - Move Failure RAS

Recommended Actions	
Robotics - Move Failure	
IF	THEN
The service ticket indicates that the tape library's robotics has experienced a move failure:	<ol style="list-style-type: none"> 1. Verify that the tape library is online and ready. 2. Verify the state of the tape library component that failed. 3. Verify the media is in the slot. 4. Verify the drive/library is online using the StorNext GUI.
The problem <u>IS</u> resolved:	Close the service ticket. Refer to Closing Service Tickets .
The problem has <u>NOT</u> been resolved:	<ol style="list-style-type: none"> 1. Note any codes displayed on the tape library's control panel. 2. Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. 3. Contact the Quantum Technical Assistance Center. In the USA: 1+800-284-5101 UK, France and Germany: 00800 4 QUANTUM EMEA: +44 1256 848 766 On the Web: http://www.quantum.com/support <p>OR</p> <ol style="list-style-type: none"> 4. If you are a properly-trained service professional, perform the procedures required for this type of tape library.

Figure 51 Robotics - Wrong Firmware Level/Invalid Library Type RAS

Recommended Actions	
Robotics - Wrong Firmware Level/Invalid Library Type	
IF	THEN
The service ticket indicates the tape library's firmware level is wrong:	Use the tape library's control panel to verify the firmware level for this release against the StorNext Release Notes . Contact your library vendor to obtain the proper firmware.
The service ticket indicates the tape library type is invalid:	Disconnect the tape library and contact the Quantum Technical Assistance Center using the contact information below.
The problem <u>IS</u> resolved:	Close the service ticket. Refer to Closing Service Tickets .
The problem has <u>NOT</u> been resolved:	<ol style="list-style-type: none"> 1. Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. 2. Contact the Quantum Technical Assistance Center. In the USA: 1+800-284-5101 UK, France and Germany: 00800 4 QUANTUM EMEA: +44 1256 848 766 On the Web: http://www.quantum.com/support <p style="text-align: center;">OR</p> <p>If you are a properly-trained service professional, perform the procedures required for this type of tape library.</p>

Figure 52 Backup Failed RAS

Recommended Actions	
Backup Failed	
Backup failure errors typically fall into one of three categories:	
<ul style="list-style-type: none"> Media Issues (for example, out of media, archive offline, no drives available, and so on) System Software Issues (for example, metadata could not be applied) File System Issues (for example, file system not mounted, fsm not running, and so on) 	
To determine the exact cause of the backup failure, see the error log included in the RAS notification, or the email notification. The error log contains the actual output of the <code>anbackup</code> command, and will help the Quantum Technical Assistance Center determine the exact cause of the backup failure.	
The following table lists some common backup failure errors and the corresponding recommended actions. For errors not listed, contact the Quantum Technical Assistance Center.	
IF	THEN
You receive one of the following media errors: "All copies of files not stored" "Store files to media failed" "Store failed for backup files"	Check all media, drives, and archives.
You receive one of the following system software errors: (1) "Error connecting to database" (2) "Error opening fs_sysparm file" (3) "FSM software not running" (4) "Cannot run backup on standby server"	Restart the StorNext software via the GUI. Before you can do a backup, the storage manager, the database, and the file system must be running. (1) This error indicates that the Linter database is not running. (2) This error indicates that SNFS was not installed correctly, or that the Configuration Wizard was terminated prematurely. (3) This error indicates that the storage manager is not running. (4) This error indicates an HA problem where a backup was run on the standby system.
You receive one of the following file system errors: "Could not set store exclude on \$stagingArea" "Backup staging directory could not be created" "Backup temporary directory could not be created" "Invalid arguments" (Caused during manual CLI invocation when incorrect syntax is used) "Application of metadata journals failed" "Missing \$ENV ('DSM_DIR')/config/zamlst file" (SNFS installation issue or corrupt filesystem) "Missing \$ENV ('DSM_DIR')/config/fFileSystem.cfg file" (Backup file system is missing or was deleted) "Missing or bad metadata dump file for "	<ol style="list-style-type: none"> Make sure the file system exists and is mounted, and that the file system manager (FSM) is running. (This might require restarting the file system to get FSM running and the file system mounted.) Metadata issues might require re-dumping metadata. This process involves unmounting the file system, stopping the file system, dumping metadata, restarting the file system, and remounting the file system. If there is access loss to the file system (for example, you cannot create a directory,) repeat step 1.
You receive the error "backup execution could not complete".	<ol style="list-style-type: none"> Make sure you have the email notification you received after the backup failed. Contact the Quantum Technical Assistance Center and send them the email notification so they can determine the cause of the backup failure.
The problem IS resolved:	Close the service ticket. Refer to Closing Service Tickets .
The problem has NOT been resolved:	<ol style="list-style-type: none"> Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. Contact the Quantum Technical Assistance Center. In the USA: 1+800-284-5101 UK, France and Germany: 00800 4 QUANTUM EMEA: +44 1256 848 766 On the Web: http://www.quantum.com/support

Figure 53 Backup Errors RAS

Recommended Actions	
Backup Errors	
<p>The backup status can be obtained on any currently running or last completed backup by running the <code>sbackup -s</code> command. The first line shows the overall status; the status line contains the same string viewed in the RAS message. The log file associated with that backup is shown beneath the status line, and shows any errors that have occurred. All errors in the log file are prefaced with ERR.</p> <p>Below is a list of individual errors and recommended actions.</p>	
IF	THEN
Backup execution could not complete:	This is a generic failure message. Run the <code>sbackup -s</code> command and examine the failure.
There was an error connecting to database:	The database has not been started or is in a state that does not allow communication. Restart the database software.
There was an error opening the <code>fs_sysparm</code> file:	The <code>/usr/adc/TSM/config/fs_sysparm</code> file cannot be located. Contact the Quantum Technical Assistance Center for assistance.
TSM software is not running:	The StorNext TSM software is down. Restart the software.
The backup staging directory could not be created:	<ul style="list-style-type: none"> Verify that the file system used by the <code>sbackup</code> command is active and mounted. Make sure <code>root</code> user has permission to create new directories.
The system could not store exclude on <code><file system></code> :	Contact the Quantum Technical Assistance Center .
The backup temporary directory could not be created:	<ul style="list-style-type: none"> Verify that the file system used by the <code>sbackup</code> command is active and mounted. Make sure <code>root</code> user has permission to create new directories.
There were invalid arguments:	Check the usage by issuing the <code>sbackup -h</code> command, or through the man page.
Application of metadata journals failed:	The metadata for a file system might be corrupt. Contact the Quantum Technical Assistance Center .
The file <code>/usr/adc/DSN/config/asm1st</code> is missing:	A configuration file is missing from the file system software directory. Contact the Quantum Technical Assistance Center .
The file <code>/usr/adc/DSN/config/<file system>.cfg</code> is missing:	A configuration file is missing from the file system software directory. Correct or provide a configuration file for this file system.
The metadata dump file for <code><file system></code> is missing:	A new file system metadata dump must be generated for the file system. Use the GUI to create the metadata dump file. (From the SNFS Home Page, choose Metadata Dump from the Admin menu.)
<ul style="list-style-type: none"> All copies of files not stored: Store files to media failed. Store failed for backup files: 	<p>Run either the <code>sbackup -s</code> or the <code>sbackupreport</code> command, or through the GUI run a Backup Report to see which copy of the backups failed. Check all media and archives associated with that copy to determine the failure.</p> <p>To run a Backup Report from the GUI:</p> <ol style="list-style-type: none"> Access the SNFS home page. Choose Backups from the Reports menu.
The problem IS resolved:	Close the service ticket. Refer to Closing Service Tickets .
The problem has NOT been resolved:	<ol style="list-style-type: none"> Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. Contact the Quantum Technical Assistance Center. <p>In the USA: 1-800-284-5101 UK, France and Germany: 00800 4 QUANTUM EMEA: +44 1256 848 766 On the Web: http://www.quantum.com/support</p>

Figure 54 Configuration Violations RAS

Recommended Actions

Configuration Violations

When a configuration violation occurs in the StorNext application, it must be repaired by stopping the system, editing the configuration, and then restarting the system. Below are specific configuration violations and recommended actions to repair each specific issue.

IF	THEN								
There is more than one affinity on one stripe group.	You cannot have more than one affinity on one stripe group. Examine all DSM configuration files (/usr/adc/DSM/config/*.cfg). In any file that has the Data Migration Flag set to YES , and for every stripe group with more than one affinity, remove the extra affinities.								
The file system does not contain at least one non-exclusive data stripe group.	The file system has at least one affinity, and therefore must contain at least one non-exclusive data stripe group. Examine all DSM configuration files (/usr/adc/DSM/config/*.cfg). In any file that has the Data Migration Flag set to YES , make sure that at least one stripe group has the following configuration: Metadata No Journal No Exclusive No								
A file system contains both data stripe groups with affinities and data stripe groups without affinities.	A file system can contain data stripe groups with affinities, or data stripe groups without affinities, but it cannot contain both. Examine all DSM configuration files (/usr/adc/DSM/config/*.cfg). In any file that has the Data Migration Flag set to YES , make sure that either every stripe group has an affinity, or that every stripe group does not have an affinity.								
There are affinities across all managed file systems.	No more than two affinities across all managed file systems are allowed. Examine all DSM configuration files (/usr/adc/DSM/config/*.cfg). For all of the configuration files that have the Data Migration Flag set to YES , change the stripe group Affinities so there are no more than a total of two.								
The number of affinities on managed file systems do not match for TSM and CVFS.	Examine all DSM configuration files (/usr/adc/DSM/config/*.cfg). For all of the configuration files that have the Data Migration Flag set to YES , make sure the complete list of affinities matches the TSM affinity names found in the TIERDEF database table.								
TSM does not recognize the CVFS managed file system affinity name.	Examine all DSM configuration files (/usr/adc/DSM/config/*.cfg). For all of the configuration files that have the Data Migration Flag set to YES , make sure the complete list of affinities matches the TIERNAME fields found in the TSM TIERDEF database table.								
An affinity in a policy class is not found in the TIERDEF table.	Make sure the non-zero elements in the TIERLIST field of the TSM CLASSDEF tables all match the TIERNAME fields in the TSM TIERDEF table.								
A StorNext file system has exceeded the maximum allowable threshold for percent used.:	<ul style="list-style-type: none"> Examine the TSM logs for a message identifying the full file system, and then remove unwanted files or save them off to a different file system. <p>OR</p> <ul style="list-style-type: none"> Log onto the application host and issue a mount command to obtain a list of mount points. For each mount point, type the following: <code>df -k <mount point></code> <p>After the file system in question has been identified, remove any or all unwanted files, or save them off to a different file system.</p>								
The problem IS resolved.	Close the service ticket. Refer to Closing Service Tickets .								
The problem has NOT been resolved.	<ol style="list-style-type: none"> Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. Contact the Quantum Technical Assistance Center. <table border="0"> <tr> <td>In the USA:</td> <td>1-800-284-5101</td> </tr> <tr> <td>UK, France and Germany:</td> <td>00800 4 QUANTUM</td> </tr> <tr> <td>EMEA:</td> <td>+44 1256 848 766</td> </tr> <tr> <td>On the Web:</td> <td>http://www.quantum.com/support</td> </tr> </table> 	In the USA:	1-800-284-5101	UK, France and Germany:	00800 4 QUANTUM	EMEA:	+44 1256 848 766	On the Web:	http://www.quantum.com/support
In the USA:	1-800-284-5101								
UK, France and Germany:	00800 4 QUANTUM								
EMEA:	+44 1256 848 766								
On the Web:	http://www.quantum.com/support								

Figure 55 Invalid Configuration RAS part 1

Recommended Actions	
Invalid Configuration	
IF	THEN
The configuration file containing the error is <code>filesystems.config</code> , and the error indicates that an entry is missing:	<p>Examine the configuration file to identify the missing entry.</p> <ul style="list-style-type: none"> If the entry should be included in the configuration file, add it to the file. If the entry should not be included in the configuration file, contact the Quantum Technical Assistance Center to clean up the database entries.
The list of managed file system names do not match for TSM and CVFS:	<p>If the file system is supposed to be managed by TSM but is listed as not managed, do the following:</p> <ol style="list-style-type: none"> Stop StorNext. Change the <code>/usr/adc/dsm/config/</code> configuration file for that file system. Change the "Data Migration" property from "No" to "Yes." Note: Verify with your system administrator that other settings in the configuration file are valid. Restart StorNext. <p>If the file system is not supposed to be managed by TSM but is listed as managed, contact the Quantum Technical Assistance Center for assistance locating and deleting all StorNext management data for all data on that file system.</p>
Multiple tape drives have the same serial number and device path:	<ol style="list-style-type: none"> Restart the StorNext software system, which will correct the internal records of the device path and serial number. If the problem persists, contact the Quantum Technical Assistance Center.
A configured tape drive has an invalid slot identifier:	<ol style="list-style-type: none"> Use the StorNext GUI to delete and then re-add the drive. <p style="text-align: center;">OR</p> <ol style="list-style-type: none"> Contact the Quantum Technical Assistance Center.
A configured tape drive has no detectable SCSI device path:	<ol style="list-style-type: none"> Verify that the device is powered on. Verify that the device has not registered any errors. Verify that the device is physically connected by checking all cabling and connections. Verify that the device is correctly mapped to the host server. (Caution: If your system configuration allows you to logically remap all devices, verify that StorNext has been stopped prior to beginning the remapping. Restart StorNext after the remapping is complete.
A managed tape drive does not have a corresponding TSM <code>CfgD1r</code> element:	<ol style="list-style-type: none"> Use the StorNext GUI to delete and then re-add the drive. <p style="text-align: center;">OR</p> <ol style="list-style-type: none"> Contact the Quantum Technical Assistance Center.
You receive one of the following "Invalid Media Configuration" messages: <ul style="list-style-type: none"> MSM does not have any Data media classes A TSM Media type could not be converted to an MSM media type MSM does not have any MediaCriteria for managed media 	Contact the Quantum Technical Assistance Center for help further diagnosing the invalid media configuration.
A TSM Data Policy media type does not match any media existing in the system:	<ul style="list-style-type: none"> If the data policy media type is correct, add that type of media to an archive that supports that media type. <p style="text-align: center;">OR</p> <ul style="list-style-type: none"> If the data policy media type is incorrect, change that policy's media type to the correct media type.

Figure 56 Invalid Configuration RAS part 2

<p>A media ID that exists in TSM does not exist in MSM.</p>	<ol style="list-style-type: none"> 1. Visually locate the medium in a specific archive. 2. Run the <code>vsaudit</code> command and verify that the medium now appears as a New Medium. 3. Run the <code>vazeclassify</code> command, converting the medium from the ADDBLANK to the DATA media class. <p>If the medium does not physically exist in any archive or if the above commands fail, contact the Quantum Technical Assistance Center.</p>								
<p>The SNAPI configuration file is not valid.</p>	<p>Correct the errors in the configuration file <code>/usr/adc/SNAPI/config/snapi.cfg</code>.</p>								
<p>The problem IS resolved:</p>	<p>Close the service ticket. Refer to Closing Service Tickets.</p>								
<p>The problem has NOT been resolved:</p>	<ol style="list-style-type: none"> 1. Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets. 2. Contact the Quantum Technical Assistance Center. <table border="0"> <tr> <td>In the USA:</td> <td>1-800-284-5101</td> </tr> <tr> <td>UK, France and Germany:</td> <td>00800 4 QUANTUM</td> </tr> <tr> <td>EMEA:</td> <td>+44 1256 848 766</td> </tr> <tr> <td>On the Web:</td> <td>http://www.quantum.com/support</td> </tr> </table> 	In the USA:	1-800-284-5101	UK, France and Germany:	00800 4 QUANTUM	EMEA:	+44 1256 848 766	On the Web:	http://www.quantum.com/support
In the USA:	1-800-284-5101								
UK, France and Germany:	00800 4 QUANTUM								
EMEA:	+44 1256 848 766								
On the Web:	http://www.quantum.com/support								

Figure 57 Downloading a System State Capture RAS

Downloading a System State Capture

Use the [Capturing a System State](#) feature to obtain and preserve detailed information about the current StorNext state. When the system state is captured, **ALL** viable logs for StorNext software components are saved and available for review and analysis. Use the Downloading a System State Capture feature to save the capture state to a local or network drive for system troubleshooting.

1. Access the StorNext home page.
2. Choose **Capture State** from the **Service** menu.
The **Service - Capture System State** screen appears.
3. Select a captured system state to download.
4. Click **Download**.
The **File Download** screen appears.
5. Click **Save** to save the zipped file to a local or network drive.
The **Save As** screen appears.
6. Click **Save**.

Figure 58 Capturing a System State RAS

Capturing a System State

Use the Capture System State feature to obtain and preserve detailed information about the current StorNext state. The Capture State feature includes **ALL** viable logs for the hardware and software components. Use the [Downloading a System State Capture](#) feature to save the captured system state to a local or network drive for troubleshooting purposes.

1. Access the StorNext home page.
2. Choose **Capture State** from the **Service** menu.
The **Service - Capture System State** screen appears.
3. Click **Capture**.

The **Progress** window appears.

 **NOTE:** Once the **Progress** window appears, you cannot cancel or stop this action. However, you can close the window by clicking the **X** in the upper-right corner of the window, but confirmation of operation success or failure is **NOT** shown.


IF	THEN
The Progress window shows Success :	The system state was successfully captured.
The Progress window shows Failure :	The system state was NOT captured. To view the troubleshooting procedures, click View Recommended Actions . To view the error details, click Error Logs for information on why the system state was not captured. To capture an additional system state, repeat Steps 1 - 3.

3. Click **Close** to close the **Progress** window.

The **Tools - Capture System State** screen appears.

4. Download the captured state.

Refer to [Downloading a System State Capture](#).



Glossary

A

Affinity : An association between a relation point in the file system and a stripe group. It allows the user to direct data to specific primary disks by writing to the affinities associated relation point.

C

Clean Media The operation of logically removing old file versions from a piece of media. This is a database operation that removes knowledge of managed files that have been updated or removed. A piece of media that contains nothing but removed files will not be considered blank until it is cleaned.

Configuration Wizard A tool for setting up a basic environment for the management of data, both on disk and on removable media (tape or disk). It appears the first time the administrator connects to the browser after installing StorNext.

D

Data Storage Manager (DSM) One of several components that make up StorNext. The DSM corresponds to the StorNext File System.

Drive Pool A grouping of drives for use in storing and retrieving data.

F **Fibre Channel (FC)** A high speed data transfer architecture.

File Transfer Protocol (FTP) The protocol used on the Internet for sending files.

G **GUI** Graphical User Interface.

M **Managed Directory** A directory that has a policy class relationship.

Managed File System A file system that enables automatic data movement managed by StorNext Storage Manager between the primary disk and secondary storage (either disk or tape).

MediaClass A grouping of media used for storing or retrieving data.

Media Storage Manager (MSM) The Media Storage Manager is responsible for controlling media and archives.

Metadata Server The system hosting the SNFS and SNSM server installation and processes.

P **Policy Class** A set of rules and criteria set up by SNSM that control the movement of data between primary disk to secondary storage (either disk or tape).

Q **Quantum Technical Assistance Center** The Quantum customer help desk.

Quota This variable enables or disables the enforcement of the file system quotas.

R **Recover** The process of bringing back to disk a managed file that was previously removed from the disk. This can only be done if the file had been successfully stored to media. Also, the file cannot not exist in the Trash can. (See Undelete) File recovery can be

done regardless of whether the Trash can is enabled, up until the time the containing media is cleaned.

Relation Point/Relations A mapping of a policy class to a directory in a managed file system.

Relocation The process of moving a file from one affinity on a file system to another affinity on that file system.

Restore The process of replacing a file system's contents after some sort of disaster. Also known as disaster recovery.

Retrieve The process of retrieving data for a file from secondary storage (either disk or tape).

RHAS Red Hat Advanced Server

RHEL Red Hat Enterprise Linux

S

SCSI Small Computer System Interface. The interface that is used to talk to most hardware devices such as tape and libraries.

StorNext A scalable, high performance, data management solution that ensures the long-term safety and recoverability of data in SAN environments, while optimizing the use of storage resources. It consists of two components, the StorNext Storage Manager (SNSM) and the StorNext File System (SNFS).

Storage Area Network (SAN) A SAN is a dedicated, high-performance network whose primary purpose is the transfer of data along FC or high-speed Ethernet connections between servers, interconnect devices, and storage peripherals.

StorNext File System (SNFS) One of the two components that make up StorNext. SNFS is primarily used to provide Fibre Channel connections (but supports other types of connections) in a serverless environment which enables clients to access data and share files.

StorNext Storage Manager (SNSM) One of several components that make up StorNext. SNSM combines the functionality of two products, TSM and MSM to provide high-performance file migration and management services, and to manage automated and manual media libraries, including library volumes.

Store The process of copying data for a file to secondary storage (either disk or tape).

Stripe Group A set of similar storage devices that can be maintained as a group.

T

Tertiary Storage Manager (TSM) The Tertiary Storage Manager is responsible for policy management and controlling data movement between primary disk and secondary storage (either disk or tape).

Truncation The process of freeing data blocks stored to secondary storage (either disk or tape). The file name remains visible in the file system.

U

Undelete The process of returning a file from the Trash can to its original location on disk. This can be done only if the Trash can is enabled.

Unmanaged File System A file system that does not have archive capability controlled by SNSM.