

Quantum[®]

User's Guide

StorNext 4.2.1



Quantum StorNext 4.2.1 User's Guide, 6-67370-02 Rev A, December 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

Copyright 2011 by 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, StorNext, the DLT logo, DXi, GoVault, SDLT, StorageCare, Super DLTtape, and SuperLoader are registered trademarks of Quantum Corporation in the U.S. and other countries. Protected by Pending and Issued U.S. and Foreign Patents, including U.S. Patent No. 5,990,810.

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.	2
	About Distributed LAN Clients	2
	About Licensing	3
	Purpose of This Guide	3
	How This Guide is Organized.	3
	Notes, Cautions, and Warnings	4
	Document Conventions	5

Chapter 2	StorNext GUI Overview	7
	Accessing the StorNext GUI	7
	The StorNext Home Page.	10
	StorNext Monitors	11
	StorNext Home Page Dropdown Menus	14

Chapter 3	The Configuration Wizard	19
	High Availability Systems	20
	Step 1: Welcome	20
	Step 2: Licenses	21
	Step 3: Name Servers.	22
	Multiple fsnameserver Hosts and Redundant Metadata Networks	24
	Configuring a Foreign Server	26
	Entering Name Servers	27
	Deleting a Name Server.	28
	Step 4: File Systems.	29
	Allocation Session Reservation	33
	Editing a File System and Performing Other Actions.	34
	Step 5: Storage Destinations	34
	Adding a New Library	35
	Viewing an Existing Library.	36
	Editing a Library	38
	Deleting a Library	39
	Performing Other Library Actions	39
	Storage Disk Overview	40
	Adding a New Storage Disk	41
	Viewing an Existing Storage Disks.	42
	Editing a Storage Disk.	43
	Deleting a Storage Disk.	44
	Adding a New Data Replication Target	44
	Editing a Data Replication Host	45
	Deleting a Data Replication Target	46
	Adding a New Mount Point	46
	Enabling Data Deduplication	47
	Step 6: Storage Policies	48
	Adding a Storage Manager Storage Policy	49
	Adding a Replication Storage Policy	55
	Viewing a Storage Policy.	55
	Running a Storage Policy	56
	Editing a Storage Policy	56
	Deleting a Storage Policy	57
	Step 7: Email Server.	57
	Adding an Email Server.	58

Step 8: Email Notification 58

 Adding an Email Recipient 59

 Viewing Email Recipient Information 60

 Editing an Email Recipient 61

 Deleting an Email Recipient 61

Step 9: Done 62

Chapter 4	File System Tasks	63
	Label Disks.	64
	Labeling a Device	64
	Unlabeling a Device	66
	Understanding Resource Allocation.	66
	About File System Expansion	66
	About Stripe Group Movement.	67
	Expansion and Movement Steps.	68
	Using Resource Allocation From the Command Line	68
	Checking the File System	69
	Adding a Stripe Group Without Moving	69
	Adding and Moving a Data Stripe Group	70
	Moving a Metadata/Journal Stripe Group.	71
	Check File System	73
	Viewing and Deleting a Check Report.	75
	File System Check Output Files	76
	Affinities	76
	Allocation Strategy	77
	Example Use Cases	77
	Adding a New Affinity	78
	Deleting an Affinity.	80
	Migrate Data	80
	Truncation Parameters.	82

Chapter 5	Storage Manager Tasks	85
	Storage Components	86
	Setting Devices Online and Offline	87
	Additional Options for Tape Drives	87
	Drive Pools	87
	Viewing Drive Pool Information	88
	Adding a Drive Pool	89
	Editing a Drive Pool	90
	Deleting a Drive Pool	90
	Media Actions	91
	Viewing Media Information	91
	Filtering Media	92
	Performing Media Actions	92
	Storage Exclusions	100
	Accessing Storage Exclusions	101
	Adding an Exclusion Pattern	102
	Truncation Exclusions	104
	Accessing Truncation Exclusions	105
	Adding an Exclusion Pattern	105
	Tape Consolidation	108
	Setting Tape Consolidation Parameters	108
	Scheduling Tape Cleaning and Defragmentation	110
	Library Operator Interface	111
	Software Requests	113
	Scheduler	114
	Viewing a Schedule	115
	Adding a Schedule	116
	Editing an Existing Schedule	118
	Deleting an Existing Schedule	118
	Alternate Retrieval Location	119

Distributed Data Mover (DDM)	121
Distributed Data Mover Overview	121
Installing the DDM Feature on Clients.	125
Accessing Distributed Data Mover	125
Enabling DDM.	126
Managing DDM Hosts.	127
Host Priority	130
Distributed Data Mover Reporting	130

Chapter 6 Replication and Deduplication 131

Replication Overview	132
Replication Configuration Overview	132
Replication Process Overview	133
Files Excluded From Replication	134
Replication Terms and Concepts	135
Namespace Realization	135
Blockpool	136
Blackout Period	136
Replication Source Policy and Replication Source Directory	136
Replication Target Directory	136
Replication Schedule.	137
Replication Copies.	137
Bandwidth Throttling	137
Multilink	137
Virtual IP (vIP)	138
Some Replication Scenarios	138
Scenario 1: Simplest Replication	139
Scenario 2: Replicating Multiple Copies in the Same Target File System	139
Scenario 3: Replicating to Multiple Target Hosts / File Systems.	141
Additional Replication Possibilities	142
Non-Supported Replication Between Source and Target.	145
"Chained" Replication.	145

Configuring Replication	146
Step 1: Create Source and Target File Systems	146
Step 2: Setting up the Blockpool.	149
Step 3: Creating Replication Targets	151
Step 4: Create a Replication Storage Policy	152
Configuration Steps Summary	158
Scheduling Replication Blackouts (Optional).	159
Optional HA and Multilink Configuration	161
Running Replication Manually (Optional)	165
Replication Statuses and Reporting	166
Replication Reports	166
Replication Administration	166
StorNext Jobs	167
Troubleshooting Replication	168
Data Deduplication Overview	168
How Deduplication Works	169
Deduplication and Replication	169
Setting Up Deduplication	170
Step 1: Creating a Deduplication-Enabled File System	170
Step 2: Specifying the Blockpool.	171
Step 3: Creating a Deduplication-Enabled Storage Policy	171
Data Deduplication Functions	173
Deduplication Administration	173
Deduplication Reports	173

Chapter 7

Tools Menu Functions	175
User Accounts	177
Adding a New User	177
Viewing an Existing User Profile	179
Modifying an Existing User	180
Deleting an Existing User	181
Client Download	182
System Control	184
Starting or Stopping StorNext File System	185
Starting or Stopping StorNext Storage Manager	185
Refreshing System Status	185
Specifying Boot Options	185

- File and Directory Actions 185
 - Store Files 186
 - Change File Version 188
 - Recover Files 189
 - Recover Directories 190
 - Retrieve Files 193
 - Retrieve Directory 194
 - Truncate Files 194
 - Move Files 195
 - Modify File Attributes 196
 - View File Information 198
- File Systems 199
- Storage Manager 199
- Replication and Deduplication 200
- HA 201

Chapter 8	Service Menu Functions	203
	The Health Check Function	203
	Running a Health Check	204
	Viewing the Health Check Results	205
	Viewing Health Check Histories	206
	The Capture State Function	206
	Creating a Capture State Log	207
	Deleting a Previous System State Capture	208
	Creating a Capture State for an HA Secondary Node	208
	The System Backup Function	208
	The Admin Alerts Function	209
	The Tickets Function	211
	Viewing Ticket Information	211
	Editing Ticket Information	213
	Closing Tickets	214
	Logging	214
	Enabling Logging	215

Chapter 9	Converting to HA	217
	HA Overview	217
	HA Terms and Concepts	218
	Failover	218
	Primary Node	218
	Secondary Node	218
	Virtual IP (vIP)	219
	Virtual Netmask	219
	HA Reset	219
	Preparing for HA Conversion	220
	Pre-Conversion Steps	220
	HA and Distributed LAN Clients	221
	Converting to HA.	221
	HA Conversion Procedure	222
	Managing HA	224
	HA Statuses and Reporting	226
	Troubleshooting HA	226
Chapter 10	StorNext Reports	227
	Report Navigation Controls.	228
	StorNext Logs	229
	The Jobs Report.	230
	Viewing Detailed Job Information.	231
	Exiting the Jobs Report screen.	233
	The Files Report.	233
	The Drives Reports.	236
	The File Systems Report.	238
	The SAN Devices Report.	239
	The Media Report	241
	The Relations Report	244
	The Tape Consolidation Report	244
	The SAN and LAN Clients Report	246

The LAN Client Performance Report.	247
The Distributed Data Mover Report.	249
Replication Deduplication Reports.	250
Policy Activity Report	250
Policy Summary Report.	252

Chapter 11	Customer Assistance	257
	Quantum Technical Assistance Center.	257

Appendix A	Operating Guidelines	259
	The Reserved Space Parameter	259
	Linux Configuration File Format	260
	Distributed LAN Server/Client Network and Memory Tuning	261
	Distributed LAN Server and Client Network Tuning.	261
	Distributed LAN Server Memory Tuning	263
	Configuring LDAP	264
	Using LDAP.	264
	UNIX File and Directory Modes	265
	LDAP Refresh Timeout.	265
	Setting Up Restrictive ACLs	266
	Default Single-Path I/O Retry Behavior.	266
	Event Handles for fsm.exe on a Windows Metadata Server.	266
	FSBlockSize, Metadata Disk Size, and JournalSize Settings	267
	Disk Naming Requirements	269
	Changing StorNext's Default Session Timeout Interval	270
	General Operating Guidelines and Limitations	271

Appendix B	Additional Replication and Deduplication Information	289
	Replication Configuration File	289
	Replication Terminology and Conventions.	290
	Copies and Versions	290

Replication Target Directories	291
Number of Replication Copies.	292
Isolating a Replication Target Directory.	293
Final Recommendation For Target Directories.	294
StorNext snpolicyd Policies	295
Replication Copies = 2 (Detail)	298
More About Replication Target Directories	301
Deduplication Overview.	303
Enabling Deduplication.	305
Deduplication Modification Time	305
Deduplication and Blockpools.	306
Deduplication and Truncation	306
Enabling Deduplication and Truncation	307
Storage Manager Truncation	307
Replication, Deduplication and Truncation	308
Replication, Deduplication and Storage Manager	309
Replicating into a Storage Manager Relation Point.	309
Truncation and Deduplication / Replication (with and without SM)	311
The snpolicyd Debug Log	318

Appendix C

High Availability Systems	321
High Availability Overview.	322
HA Internals: HAMon Timers and the ARB Protocol	324
Primary and Secondary Server Status	328
File System Types	328
The <code>ha_peer</code> and <code>fsnameservers</code> File.	330
HA Manager	330
Configuration and Conversion to HA.	335
Conversion to HA	336
SyncHA process.	338
Managing High Availability in the StorNext GUI	339
Configuring Multiple NICs	342
LAN Configuration	342

High Availability Operation	343
Windows and Linux SNFS Installations Without the HaShared File System	344
Linux SNMS and SNFS Installations with the HaShared File System	345
HA Resets	350
HA Resets of the First Kind	350
HA Resets of the Second Kind	351
HA Resets of the Third Kind	352
Using HA Manager Modes	352
HA Tracing and Log Files	352
Single (Singleton) Mode	353
FSM failover in HA Environments	354
Failover Timing	355
Replacing an HA System	359

Appendix D

Web Services API	363
Using the APIs	364
Using APIs With the High Availability MDC Feature	365
WS-API APIs	366
The doCancel API	366
The doMediaMove API	366
The doRetrieve API	367
The doStore API	367
The doTruncate API	367
The getDriveReport API	367
The getFileLocation API	367
The getMediaInfo API	367
The getMediaReport API	367
The getSMQueue API	368
Examples	368
Example: WS-API URLs	368
Sample XML Output	369
Sample JSON Output	371
Sample Text Output	373

Appendix E	Storage Manager Truncation	375
	Truncation Overview	375
	Normal Truncation	376
	Immediate Truncation	376
	Daily Truncation	376
	Space Management	377
	LOSPACE Cycle	377
	Emergency Cycle	378
	Disabling Truncation	379
	Truncation Feature Locking	379
	Disable Truncation Commands	380
	Common Problems	380
	Files Are Not Truncated as Expected	380
	"Old" Files Not Truncating According to Policy Class	381
	Small Files Not Truncating	381
	Miscellaneous Usage Notes	381
	Scheduling Truncation Manually	382
	Setting Truncation	382
	Setting Up the cron Job	383
	Considerations When Scheduling Truncation and Relocation Policies	385
Appendix F	Security	387
	StorNext Security	387
	ACLs on Windows	388
	ACLs on Mac OS X	392
	"Central Control"	395
	Limitations	397
	Example	397
	Cross-Platform Permissions	398
	Config (.cfg) File Options	399
Appendix G	Troubleshooting	401
	Troubleshooting StorNext File System	402
	Troubleshooting StorNext Storage Manager	409
	Troubleshooting OS Issues	411

Troubleshooting Replication 414

Troubleshooting HA 415

Troubleshooting StorNext Installation and Upgrade Issues . . 419

Troubleshooting Other Issues 420

Appendix H	StorNext Offline Notification	425
	StorNext Offline Notification Overview	426
	What is a StorNext Offline File?	426
	Why Block Access to Offline Files?	426
	Offline Notification Configuration Options	427
	How the Notification Feature Operates.	427
	Installing the Notification Application	429
	Installing onto a Standalone Machine.	429
	Starting the Notification Application.	435
	Configuring the Notification Application.	436
	Setting Application Options	436
	Viewing Access Denied Files	438
	Viewing the Application Log.	439
	Exiting Application Preferences.	439
	Uninstalling the Notification Application	440

Appendix I	RAS Messages	441
	Media and Drive RAS Messages.	442
	SNFS RAS Messages.	451
	Other RAS Messages	464



Figures

Figure 1	StorNext Login Window	9
Figure 2	System Control	9
Figure 3	StorNext Home Page.	11
Figure 4	Configuration > Configuration Wizard Welcome Screen .	21
Figure 5	Configuration > Licenses Screen	22
Figure 6	StorNext Foreign Servers.	27
Figure 7	Name Servers Screen	28
Figure 8	Configuration > File System Screen	29
Figure 9	Configuration > File System > New Screen	30
Figure 10	Configuration > File System > New Screen 2.	31
Figure 11	Configuration > File System > New Screen 3.	32
Figure 12	Configuration > File System > New Screen 3.	33
Figure 13	Storage Destinations > Library Screen	35
Figure 14	Storage Destinations > Library > New Screen.	36
Figure 15	Library Details Screen	37
Figure 16	Edit Library Screen	38
Figure 17	Configuration > Storage Destinations > Storage Disk Screen.	41

Figure 18	Storage Destinations > Storage Disk > New Screen	42
Figure 19	View Storage Disk Screen	43
Figure 20	Configuration > Storage Destinations > Replication / Deduplication Screen	45
Figure 21	Configuration > Storage Destinations > Replication / Deduplication Screen (Blockpool)	47
Figure 22	Configuration > Storage Policies Screen	49
Figure 23	Storage Policies > New Screen	50
Figure 24	Storage Policies > New > General Tab	51
Figure 25	Storage Policies > New > Relocation Tab	52
Figure 26	Storage Policies > New > Steering Tab	53
Figure 27	Storage Policies > New > Schedule Tab	54
Figure 28	Storage Policies > New > Associated Directories Tab . . .	55
Figure 29	View Storage Policies Screen	56
Figure 30	Configuration > Email Server Screen	58
Figure 31	Configuration > Email Notifications Screen	59
Figure 32	Configuration > Email Notifications New Screen	60
Figure 33	Configuration > Configuration Wizard Done Screen . . .	62
Figure 34	Label Disks Screen	65
Figure 35	Check File System Screen	74
Figure 36	Check File System Report	75
Figure 37	Affinities Screen	79
Figure 38	New Affinity Screen	79
Figure 39	Migrate Screen	81
Figure 40	Truncation Parameters Screen	83
Figure 41	Storage Components Screen	86
Figure 42	Drive Pools Screen	88
Figure 43	New Drive Pool Screen	89
Figure 44	Media Actions Screen	91
Figure 45	Storage Exclusions Screen	101

Figure 46	Exclusion Screen	102
Figure 47	Truncation Exclusions Screen	105
Figure 48	Exclusion Screen	106
Figure 49	Tape Consolidation Screen	109
Figure 50	Library Operator Interface Screen	112
Figure 51	Software Requests Screen.	114
Figure 52	Scheduler Screen	116
Figure 53	Scheduler > New Screen	117
Figure 54	Alternate Retrieval Location Screen	120
Figure 55	Configuration > Distributed Data Mover Screen	126
Figure 56	DDM Screen New Host	128
Figure 57	Replication Process	134
Figure 58	Replication scenario 1	139
Figure 59	Replication Scenario 2	140
Figure 60	Replication Scenario 3	142
Figure 61	Replicating From One Source to Multiple Targets.	143
Figure 62	Replicating From Multiple Sources to One Target.	144
Figure 63	Non-Supported Replication From Source to Target	145
Figure 64	Configuration > File System > New Screen	147
Figure 65	Configuration > File System > New Screen 2.	148
Figure 66	Configuration > File System > New Screen 3.	149
Figure 67	Configuration > Storage Destinations > Deduplication Screen (Blockpool)	150
Figure 68	Storage Destinations > Replication Targets Screen	151
Figure 69	Configuration > Storage Policies > New Screen	153
Figure 70	Configuration > Storage Policies > New / Source Directories Screen.	154
Figure 71	Storage Policies > New > Outbound Replication Tab	155
Figure 72	Outbound Replication Tab > Replication Schedule.	156

Figure 73	Configuration > Storage Policies Screen (Select "target")	157
Figure 74	Storage Policies > Edit > target > Inbound Replication Tab	158
Figure 75	Configuration > Storage Policies (Run Policy).	159
Figure 76	Storage Policies > New > Blackout Tab	160
Figure 77	Tools > HA Convert Screen.	162
Figure 78	Tools > Replication > Bandwidth Screen	164
Figure 79	Tools > Replication/Deduplication > Administration Screen.	167
Figure 80	Deduplication	169
Figure 81	Replication/Deduplication Policy Screen	172
Figure 82	User Accounts Screen	177
Figure 83	New User Screen.	178
Figure 84	View User Screen	180
Figure 85	Edit User Screen	181
Figure 86	Client Download Screen	183
Figure 87	Client Download Link	183
Figure 88	System Control Screen	184
Figure 89	File and Directory Action Screen	187
Figure 90	Change File Version Screen.	188
Figure 91	Recover Files Browse Screen	189
Figure 92	Recover Directories Browse Screen	191
Figure 93	Retrieve Files Screen	193
Figure 94	Retrieve Directory Screen	194
Figure 95	Truncate Files Screen	195
Figure 96	Move Files Screen	196
Figure 97	Modify File Attributes Screen	197
Figure 98	View File Info Screen	198
Figure 99	Health Check Screen	204

Figure 100 Health Check > View Selected Screen	205
Figure 101 Health Check > View History Screen.	206
Figure 102 Capture State Screen	207
Figure 103 Backup Screen	209
Figure 104 Admin Alerts Screen	210
Figure 105 Tickets Screen	211
Figure 106 Tickets > View Ticket Screen	212
Figure 107 Tickets > Edit Ticket Screen	214
Figure 108 Logging Screen.	215
Figure 109 Tools > HA Screen	223
Figure 110 Manage HA Screen	225
Figure 111 Reports > Logs Screen	230
Figure 112 Jobs Report.	231
Figure 113 Files Report.	234
Figure 114 StorNext File Browser	235
Figure 115 File Info Screen	236
Figure 116 Drives Report	237
Figure 117 Drive Information Report	238
Figure 118 File Systems Report.	239
Figure 119 SAN Devices Report	240
Figure 120 Media Report	241
Figure 121 Media Information Report	243
Figure 122 Relations Report	244
Figure 123 Tape Consolidation Report	245
Figure 124 SAN and LAN Clients Report	246
Figure 125 LAN Client Performance Report	248
Figure 126 Distributed Data Mover Report.	249
Figure 127 Replication/Deduplication Policy Activity Report.	251
Figure 128 Replication/Deduplication Policy Summary Report.	253

Figure 129	Replication/Deduplication Policy Details Report	254
Figure 130	Replication/Deduplication Policy Completion Report . . .	255
Figure 131	High Availability Manage Screen	340
Figure 132	FSM Failover in an HA Cluster	355
Figure 133	Run as Administrator	430
Figure 134	Logging in to the Administrator Account	430
Figure 135	Installing the .NET Framework	431
Figure 136	Offline Notification Setup Wizard	432
Figure 137	Quantum End User License Agreement	432
Figure 138	Select Installation Folder	433
Figure 139	Confirm Installation	434
Figure 140	Installation Complete	434
Figure 141	Manual Start	435
Figure 142	Application Options	437
Figure 143	Access Denied List	438
Figure 144	Application log	439
Figure 145	Removing the Application	440
Figure 146	Possible Drive/Media Mount Discrepancy RAS	442
Figure 147	Tape Drive Alerts RAS part 1	443
Figure 148	Tape Drive Alerts RAS part 2	444
Figure 149	Tape Drive Alerts RAS part 3	445
Figure 150	Drive Reported Drive Error RAS	446
Figure 151	Cleaning of Drive Failed RAS	447
Figure 152	Wrong Firmware Level/Invalid Drive Type RAS	448
Figure 153	Tape Drive - Reported Media Error RAS	449
Figure 154	Cleaning Media Expired RAS	450
Figure 155	Duplicate Physical Media Found RAS	450
Figure 156	Storage Disk Taken Offline RAS	451
Figure 157	Configuration Not Supported RAS	452

Figure 158	Label Validation Failure RAS	452
Figure 159	Connection Rejected RAS	453
Figure 160	File System Failover RAS	453
Figure 161	I/O Error RAS.	454
Figure 162	Journaling Error Detected RAS	454
Figure 163	SNFS License Required RAS.	455
Figure 164	SNFS License Failure RAS	455
Figure 165	LUN Mapping Changed RAS	456
Figure 166	Communication Failure RAS	456
Figure 167	Metadata Inconsistency Detected RAS	457
Figure 168	Bad File System Metadata Dump RAS.	457
Figure 169	Metadata Dump Failure RAS.	458
Figure 170	File System or Metadata Capacity Warning RAS	458
Figure 171	File Processing Failure RAS	459
Figure 172	Missing LUNs RAS.	459
Figure 173	Disk Space Allocation Failure RAS.	460
Figure 174	System Resource Failure RAS	460
Figure 175	Affinity Configuration Violations RAS.	461
Figure 176	Quota Limit or Fragmentation Warnings RAS.	462
Figure 177	Shutdown Error RAS.	462
Figure 178	Initialization Failure RAS	463
Figure 179	SNFS I/O Error RAS	463
Figure 180	Port Failure	464
Figure 181	Checksum Error RAS	465
Figure 182	Troubleshooting the StorNext Software RAS	466
Figure 183	Software Resource Violations RAS	467
Figure 184	Vault Failure RAS	468
Figure 185	Robotics - Not Ready RAS	468
Figure 186	Robotics - Move Failure RAS	469

Figure 187 Robotics - Wrong Firmware Level/Invalid Library Type
RAS. 470

Figure 188 Backup Errors RAS 471

Figure 189 Invalid Configuration RAS part 1 472

Figure 190 Invalid Configuration RAS part 2 473



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, Solaris, 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-server failures. In high availability (HA) configurations, a redundant server is available to access files and pick up processing requirements of a failed system, and carry on processing.

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

In addition to supporting StorNext clients attached via fibre channel, StorNext also supports *Distributed LAN Clients*. Unlike a direct-attached StorNext client, a distributed LAN client connects 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.

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 [Step 2: Licenses](#) on page 21

StorNext provides distributed LAN client information via the status monitors on the StorNext home page. More detailed information is available through the Clients Report and Distributed LAN Client Performance Report. For more information about StorNext reports, see [Chapter 10, StorNext Reports](#).

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*.

About Licensing

Separate licenses are required for various StorNext features, as well as to perform an upgrade to a new release. If you add new StorNext features, you must enter license information for those new features as described in the section [Step 2: Licenses](#) on page 21.

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, The Configuration Wizard](#)
- [Chapter 4, File System Tasks](#)
- [Chapter 5, Storage Manager Tasks](#)
- [Chapter 6, Replication and Deduplication](#)
- [Chapter 7, Tools Menu Functions](#)
- [Chapter 8, Service Menu Functions](#)
- [Chapter 9, Converting to HA](#)
- [Chapter 10, StorNext Reports](#)
- [Chapter 11, Customer Assistance](#)
- [Appendix A, Operating Guidelines](#)

- [Appendix B, Additional Replication and Deduplication Information](#)
- [Appendix C, High Availability Systems](#)
- [Appendix D, Web Services API](#)
- [Appendix E, Storage Manager Truncation](#)
- [Appendix F, Security](#)
- [Appendix G, Troubleshooting](#)
- [Appendix H, StorNext Offline Notification](#)
- [Appendix I, 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
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.

This chapter includes the following topics:

- [Accessing the StorNext GUI](#)
- [The StorNext Home Page](#)

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: The following browsers have been tested to work with StorNext. Browsers not listed may work but are not recommended.

- Internet Explorer 7.x, 8.x and 9.x
- FireFox 3.x

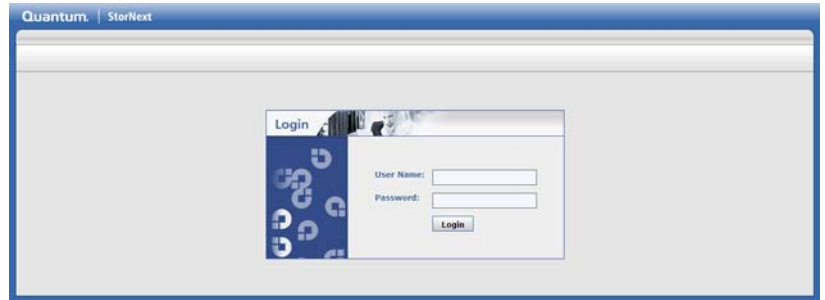
To ensure proper browser operation, all monitors must be set to display at a minimum resolution of 1024 x 768. If you use a popup blocker, be sure to disable pop-ups in order to ensure that StorNext displays properly.

- 2 If you are using Internet Explorer 9, you must verify one setting before accessing StorNext:
 - a Launch Internet Explorer 9.
 - b Choose **Internet Options** from the **Tools** menu.
 - c Click the **Advanced** tab.
 - d Under the **Security** heading, make sure the option **Do not save encrypted pages to disk** is not checked.
 - e Click **OK**.
 - f Close Internet Explorer 9 and then open it again to continue.
- 3 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.)

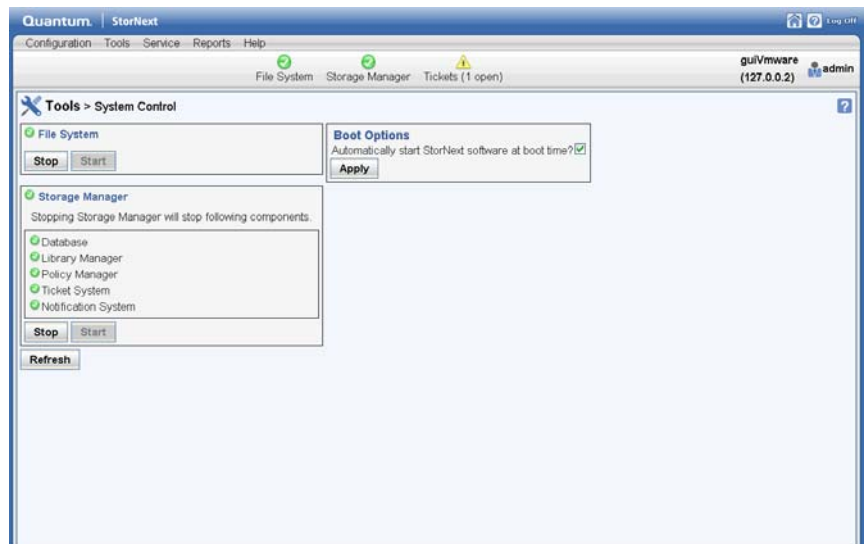
After you enter the machine name and port number, the following window appears:

Figure 1 StorNext Login Window



- 4 In the **User ID** field, type admin.
- 5 In the **Password** field, type password.
- 6 Click **Login**. The initial StorNext **System Control** screen appears.

Figure 2 System Control



- 7 On this screen you can determine if the StorNext File System and Storage Manager components are currently started. If not, click **Start** for each component to start them.
- 8 Click the home (house) icon in the upper right corner to go to the StorNext Home Page.

Note: When you log into StorNext for the first time, you might see a message warning you about a security certificate. Refer to the Quantum Knowledge Base for a permanent workaround to this issue. For a temporary solution, create a certificate exception that will allow you to log into StorNext without seeing the warning message during subsequent logins.

The StorNext Home Page

On the home page you will find the following:

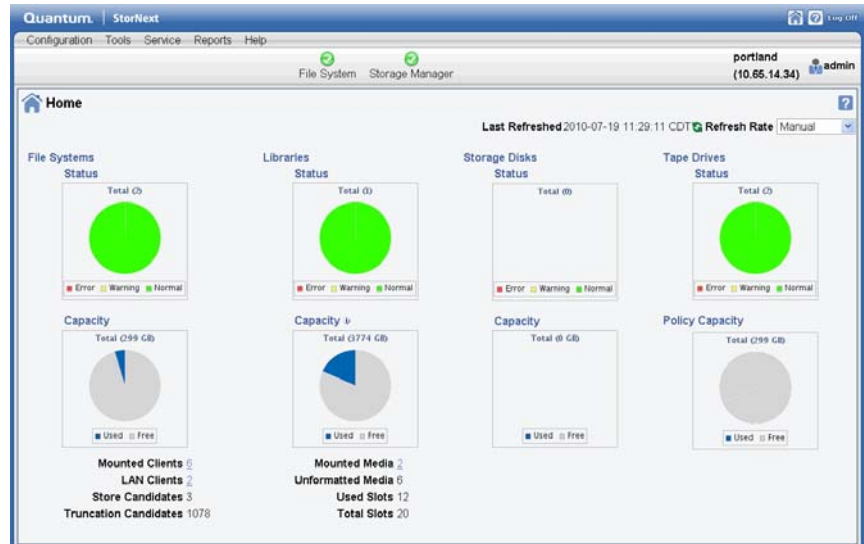
- Status and Capacity Monitors for file systems, libraries, storage disks, and tape drives
- Dropdown Menus: Configuration, Tools, Service, Reports and Help
- Current status indicators for the file system and Storage Manager
- A link to the tickets page (if tickets exist)
- A link to admin alerts (if they exist)
- A link to the Library Operator Actions Required page if actions exist
- A link to blockpool status if the blockpool is in the process of starting up

From any page you can return to the StorNext home page by clicking the Home (house) icon in the upper right corner of the screen.

Beside the Home icon is a question mark icon. Clicking this icon displays a list of StorNext online help topics.

Displayed in the upper right corner beneath the home and help icons is the user name or IP address of the StorNext user currently logged in.

Figure 3 StorNext Home Page



StorNext Monitors

The StorNext Home Page displays the following status and capacity monitors which are used to show the current state of the StorNext system:

- [File Systems Capacity Monitor](#)
- [Libraries Capacity Monitor](#)
- [Storage Disks Capacity Monitor](#)
- [Tape Drive Status](#)
- [Policy Capacity 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. Each of the status monitors provides an at-a-glance view of the total number of components (file systems, libraries, storage disks, or tape drives) and the current state of the file system: green for normal, yellow for warning, and red for error.

The information shown in the monitors is refreshed periodically. You can specify the refresh rate by choosing the desired interval from the Refresh Rate list:

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

File Systems Capacity Monitor

The File Systems Capacity Monitor provides the following information:

- Total space (in GB) for the file system
- A graphical representation of the free and used space amount
- The number of active StorNext SAN clients (connected via fibre channel or iSCSI) for which you are licensed
- 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.
- The number of store candidates, which are files selected for storage to secondary media.
- The number of files that have been stored and meet the criteria to become a truncation candidate.
- Current status (Error, Warning or Normal)

Libraries Capacity Monitor

The Libraries Capacity Monitor provides the following information:

- Total space (in GB) for the library. (This amount is an approximation if the library contains unformatted media.)
- A graphical representation of the library's free and used space amount
- The number of mounted and unmounted media

- The number of used slots
- The total number of slots
- Current status (Error, Warning or Normal)

Storage Disks Capacity Monitor

The Storage Disks Capacity Monitor provides the following information:

- Total number of storage disks
- A graphical representation of the free and used space amount
- Current status (Error, Warning or Normal)

Tape Drive Status

The Tape Drive Status Monitor provides the following information:

- Total number of tape drives
- A graphical representation of the free and used space amount
- Current status (Error, Warning or Normal)

Policy Capacity Monitor

The Policy Capacity Monitor provides the following information:

- Total space (in GB) for policy
- A graphical representation of the free and used space amount

Note: The home page status and capacity monitors are intended to give you an **approximate** at-a-glance view of all the file systems, libraries, storage disks etc. on your system.

For a detailed, more accurate summary of your system's components, click inside one of the Status or Capacity boxes to view all file system, libraries, storage disks, and so on. (For example, click inside either the File Systems Status or Capacity box to view all file systems.)

StorNext Home Page Dropdown Menus

The dropdown menu options located in the bar at the top of every page allow you to access StorNext setup, tools, service, and reporting options.

The StorNext home page contains these drop-down menus and menu options:

The Configuration Menu

Use these menu options to launch the Configuration Wizard or complete individual configuration tasks.

- **Configuration Wizard:** Launches the StorNext Configuration Wizard
- **License:** Enter or view license information for StorNext features
- **Name Servers:** Enter and set order for servers used for StorNext file systems
- **File Systems:** Add a file system to your environment
- **Storage Destinations:** Add a library or storage disk, or set up data replication and deduplication
- **Storage Policies:** Add a storage policy to a file system
- **Email Server:** Configure the email server to use for notifications
- **Email Notifications:** Configure email notifications for Service Tickets, Admin Alerts, StorNext Backups, and Policy Class Alerts

The Tools Menu

Use these options to control day-to-day operations of StorNext.

- **User Accounts:** Control user access to StorNext tasks
- **Client Download:** Download StorNext client software
- **System Control:** Stop or start the file system or StorNext Storage Manager, and specify whether to automatically start StorNext at system startup
- **File and Directory Actions:** Perform file-related and directory-related tasks such as storing and moving files, recovering and retrieving files and directories, and modifying file attributes.

- **File Systems**
 - **Label Disks:** Label disk drives
 - **Check File System:** Run a check on your file system before expanding the file system or migrating a stripe group.
 - **Affinities:** Add affinities to the file system.
 - **Migrate Data:** Migrate the file system's stripe group(s)
 - **Truncation Parameters:** Manage the file system's truncation parameters
- **Storage Manager**
 - **Storage Components:** View current status for libraries, storage disks, and tape drives; place one or more of these components online or offline
 - **Drive Pools:** Add, modify, or delete drive pools
 - **Media Actions:** Remove media from a library or move media from one library to another
 - **Storage Exclusions:** Specify file types to exclude from StorNext processing
 - **Truncation Exclusions:** Specify files to exclude from the truncation process
 - **Tape Consolidation:** Consolidate tape volumes which contain unused space
 - **Library Operator Interface:** Enter or eject media from the Library Operator Interface
 - **Software Requests:** View or cancel pending software requests
 - **Scheduler:** Schedule file system events including Clean Info, Clean Versions, Full Backup, Partial Backup, and Rebuild Policy
 - **Alternate Retrieval Location:** Specify a remote retrieval location to use in situations where files stored on tape or a storage disk cannot be accessed.
 - **Distributed Data Mover:** Distribute data movement operations from the metadata controller to client machines.
- **Replication/Deduplication**
 - **Administration:** View current progress for data replication, data deduplication, and truncation operations

- **Replication Targets:** Add replication hosts and mount points to your replication targets, and edit properties for existing hosts and mount points
- **Replication Bandwidth:** Monitor bandwidth usage for ongoing data replication processes
- **High Availability**
 - **Convert:** Convert to a High Availability (HA) configuration
 - **Manage:** Manage an HA configuration
- **Advanced Reporting:** If you have installed StorNext Advanced Reporting, this menu option allows you to access that feature. This menu option does not appear if you have not installed Advanced Reporting.

The Service Menu

Use these options to monitor and capture system status information.

- **Health Check:** Perform one or more health checks on StorNext and view recent health check results
- **Capture State:** Obtain and preserve detailed information about the current StorNext system state
- **System Backup:** Run a backup of StorNext software
- **Admin Alerts:** View informational messages about system activities
- **Tickets:** View, edit, or close service tickets generated for the system
- **Logging:** Enables robust debugging mode for advanced tracing

The Reports Menu

Use these options to view StorNext reports.

- **Logs:** Access logs of StorNext operations
- **Jobs:** View a list of pending and completed jobs on the system
- **Files:** View information about specific files, such as the owner, group, policy class, permissions, and copy information
- **Drives:** View information about the drives in your libraries, including the serial number and current state and status

- **Media:** View information pertaining to the media in selected libraries or all libraries, including the media ID, type and class, as well as the number of files and the last access date and time
- **Relations:** View the name of the policy class which corresponds to the managed directories in your system
- **File Systems:** View file system statistics including active clients, space, size, disks, and stripe groups
- **SAN Devices:** View information about devices currently attached to your SAN, including disks/LUNs, libraries, and tape drives
- **Tape Consolidation:** View statistics on the tape consolidation (defragmenting) process
- **SAN and LAN Clients:** View statistics for StorNext clients, including the number of connected clients and distributed LAN clients, and client performance
- **LAN Client Performance:** View information about distributed LAN clients and servers, including read and write speed
- **Replication/Deduplication**
 - **Policy Activity:** View replication and deduplication performance statistics
 - **Policy Summary:** View replication and deduplication information for each policy
- **Distributed Data Mover:** View activity related to the Distributed Data Mover feature

The Help Menu

Use these options to access StorNext documentation, find Quantum contact information, or detailed information about this version of StorNext.

- **Documentation:** Access the StorNext documentation
- **Support:** Access Quantum Technical Support information
- **About:** Access detailed information about your version of StorNext and the system on which it is running. Also shows StorNext patent information.



Chapter 3

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. As you complete tasks, click **Next** to proceed to the next configuration task, or click **Back** to return to the previous task. Some tasks allow you to skip the task for configuration at a later time. These tasks have a **Next/Skip** button instead of a **Next** button.

You can display the Configuration Wizard at any time by selecting **Configuration Wizard** from the StorNext **Configuration** 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 **Configuration** menu.

Following are the setup and configuration tasks the Configuration Wizard allows you to complete:

- [Step 1: Welcome](#): View disks and libraries currently available for StorNext usage
- [Step 2: Licenses](#): Enter license information for StorNext features and components

- [Step 3: Name Servers](#): Specify and order the machines acting as StorNext name servers
- [Step 4: File Systems](#): Add a StorNext file system
- [Step 5: Storage Destinations](#): Add a library, storage disks, and other storage destinations
- [Step 6: Storage Policies](#): Add a Storage Manager or replication storage policy
- [Step 7: Email Server](#): Specify an email server to handle StorNext notifications
- [Step 8: Email Notification](#): Add email notifications recipients
- [Step 9: Done](#): Signify that you are finished using the Configuration Wizard. You can also convert to a high availability (HA) system.

This chapter provides an overview of the steps necessary to complete each of the Configuration Wizard's tasks.

High Availability Systems

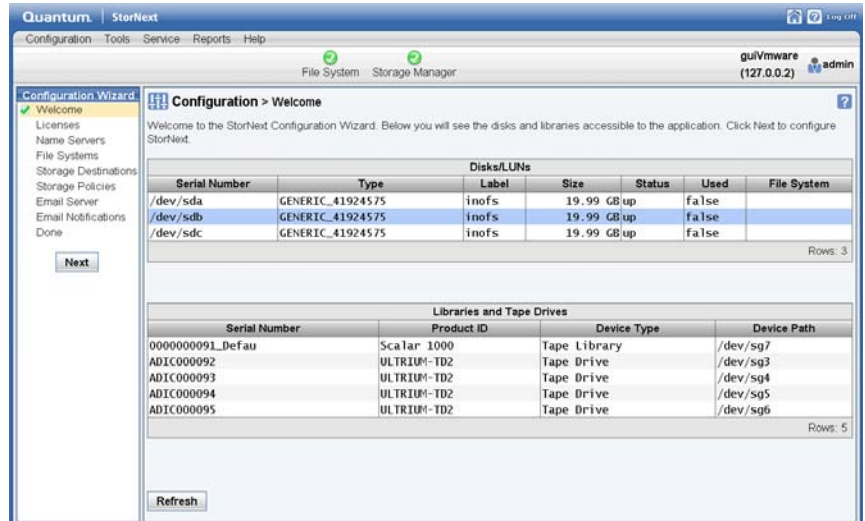
This chapter contains some instructions that pertain to high availability (HA) systems, but if you plan to convert to HA you should read [Chapter 9, Converting to HA](#). In particular, be sure to read and follow the [Pre-Conversion Steps](#) on page 220.

Step 1: Welcome

The first screen in the Configuration Wizard is the Welcome screen. This screen shows disks and libraries that are currently available for StorNext usage. As you add new disks and libraries, the information on this screen is updated.

If desired, you can manually update the screen by clicking **Refresh**. When you are ready to proceed to the next step, click **Next** in the left column.

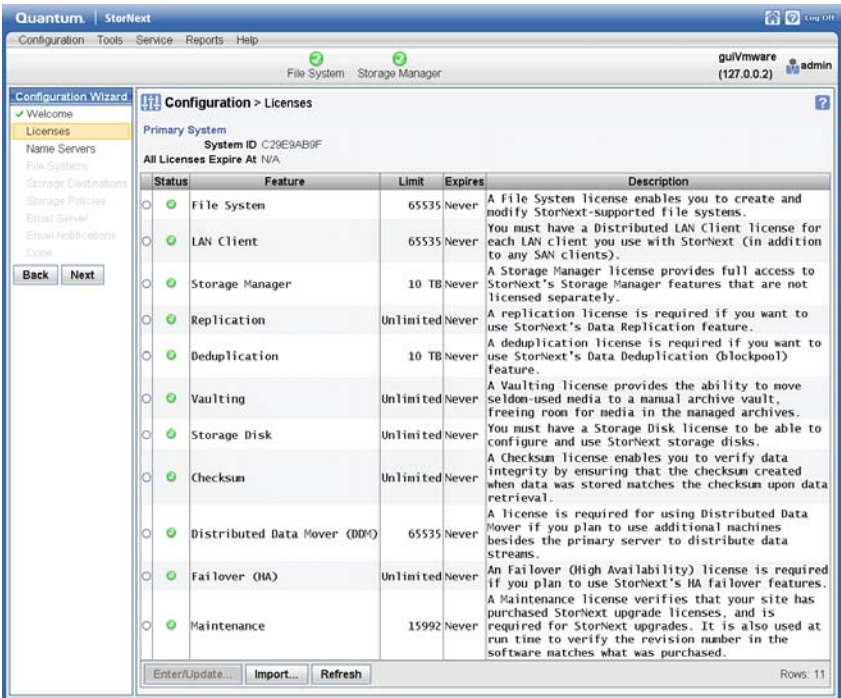
Figure 4 Configuration >
Configuration Wizard
Welcome Screen



Step 2: Licenses

The StorNext Licenses screen enables you to view your current licenses and enter new licenses (or update expired licenses) for StorNext features.

Figure 5 Configuration > Licenses Screen



For information about how to enter or update licenses, refer to the StorNext online help or the *StorNext Licensing Guide*.

Step 3: Name Servers

This screen enables you to manage machines acting as File System name servers. You may specify either a hostname or an IP addresses, but an IP address is preferable because it avoids problems associated with the lookup system (e.g., DNS or NIS).

The hostnames or IP addresses are copied into the StorNext `fsnameservers` file. This specifies both the machines serving as File System Name Server coordinator(s) and defines the metadata network(s) used to reach them. 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 the IP addresses of multiple machines in the `fsnameservers` file, one entry per line. The first IP address listed defines the path to the primary coordinator. A redundant path to this coordinator may then be specified. Any subsequent IP addresses listed serve as paths to backup coordinators. To create redundancy, Quantum recommends that you select two machines to act as coordinators. Typically, the selected systems are also configured for FSM services (MDC), but this is not a requirement.

If the `fsnameservers` file does not exist, is empty or contains the localhost IP address (127.0.0.1), 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 sharing the FSS over the SAN.

The addresses in the `fsnameservers` file define the metadata networks and therefore the addresses used to access file system services. When an MDC sends a heartbeat to a nameserver, the nameserver records the source IP address from the UDP packet and uses that as the address to advertise for FSMs local to that MDC.

If a nameserver receives multiple heartbeats on redundant metadata network interfaces, there will be a different source address for the same FSM and host. The name server will select only one of the metadata network addresses to use as the address of the FSM advertised to all hosts in the cluster. Thus all metadata traffic will use only one of the two or more redundant metadata networks.

If the network being advertised for file system services fails, a backup network will be selected for FSM services. Clients will not necessarily reconnect using the new address. If a client maintains TCP connectivity using the old address, no reconnect will be necessary. If the client needs to connect or re-connect, it will use the currently advertised IP address of the file system services.

Multiple `fsnameserver` Hosts and Redundant Metadata Networks

The addition of name server hosts to the configuration will increase the amount of name server traffic on the metadata network. Using a redundant metadata network with multi-homed name servers further increases the load.

To help you weigh the benefits versus disadvantages of having multiple name server hosts and redundant meta-data networks, here are some points to consider:

- Metadata controllers needn't be name servers.
- Name servers needn't be MDCs.
- Each additional `fsnameserver` entry adds additional heartbeats from every file system host.
- If multiple metadata networks service an individual file system, each network must have an `fsnameserver` interface. Each `fsnameserver` host must have network interface(s) on every metadata network, and each interface must be listed in the `fsnameserver` file.
- At maximum heartbeat rate, a host sends a heartbeat message to every `fsnameserver` entry twice per second. The maximum rate is in effect on a given host when StorNext services are first started, and during transition periods when an FSM is starting or failing over. Thirty seconds after services are started and when a cluster is stable, non-MDC hosts reduce their heartbeat rate to once every 5 seconds.
- Each heartbeat results in a heartbeat reply back to the sender.
- The size of the heartbeat and reply message depends on the number of file systems in the cluster.

Calculating Network Requirements

The following section may help you understand how to calculate computing requirements for name server traffic in a cluster. This example assumes a transition period when all hosts are sending heartbeat messages at twice a second.

- 1 Every host sends a heartbeat packet to every name server address, twice per second. If the host is an MDC, the heartbeat packet contains a list of FSMs running locally.
- 2 Each name server maintains the master list of FSMs in the cluster. The heartbeat reply contains the list of all FSMs in the cluster.

- 3 The NSS packet is 72 bytes, plus the file system entries. Each file system entry is 24 bytes plus the name of the file system (one byte per character), including a zero byte to terminate the string.

The file system name is always rounded up to the next 8-byte boundary. For example, a file system name of 7 characters or less would be rounded up to 8 bytes, and a file system name with 8-15 characters would be rounded up to 16 bytes. If there is room in the packet, a list of file systems which are mounted, or could be mounted, is also included.

- 4 The heartbeat message size from non-MDC clients is small because there are no locally running FSMs. The heartbeat reply message size is significant because it contains file system locations for all FSMs in the cluster.
- 5 The maximum name server packet size is 63KB(64512). This allows up to 1611 FSMs with names of 7 characters or less. With file system names of 8-15 characters, the maximum packet can hold entries for 1342 FSMs. In configurations where the maximum packet size is reached, each host would receive 129024 bytes per second from each address in the fsnameservers file. This is roughly 1MBit per second per host/address. In a configuration with dual multi-homed name servers, there would be 4 addresses in the fsnameserver file. Each host would then receive 4Mbits per second of heartbeat reply data at the maximum heartbeat rate (twice a second).
- 6 A large cluster with 500 hosts, 1600 FSMs and 4 fsnameserver addresses would produce an aggregate of about 500×4 or 2000 Mbits or 2Gbits of heartbeat reply messages per second. If the 4 fsnameserver addresses belonged to two nameservers, each server would be generating 1Gbit of heartbeat reply messages per second.

Note: During stable periods, the heartbeat rate for non-MDC hosts decreases to one tenth of this rate, reducing the heartbeat reply rate by an equivalent factor.

- 7 The metadata network carries more than just name server traffic. All metadata operations such as open, allocate space, and so on use the metadata network. File system data is often carried on the metadata network when distributed LAN clients and servers are configured. Network capacity must include all uses of these networks.

Configuring a Foreign Server

The StorNext name service supports the concept of a *foreign server*. StorNext client nodes can mount file systems that are not local to the client's home cluster. Additionally, a client may belong to no StorNext cluster by having an empty or non-existent `fsnameservers` file.

Clusters serving foreign clients address some scalability and topology issues present in the traditional client model. Depending on your needs, traditional clients, foreign clients or a mixture may result in the best performance and functionality.

Configuring foreign servers requires creating the `fsforeignservers` file on client nodes. (Configuring foreign servers through the StorNext GUI is not currently supported.) The `fsforeignservers` file must be edited using a program appropriate for the platform, such as `vi` on Linux or Wordpad on Windows.

Configuring foreign servers allows customers to better scale large numbers of clients. Since foreign clients do not participate in FSM elections, a lot of complexity and message exchange in the voting process is eliminated. In a typical StorNext HA environment, clients have equal access to both MDC candidates, making the choice of active FSM more of a load balancing decision rather than one of access.

Another benefit of foreign servers is that certain topology environments prevent all clients from having equal access to all file system MDCs and associated primary storage. By selecting the set of file system services for each client through the foreign servers configuration, the client sees only the relevant set of file systems.

The format for the `fsforeignservers` file is similar to the `fsnameservers` file in that it contains a list of IP addresses or hostnames, preferably IP addresses. One difference is that the addresses in the `fsforeignservers` file are MDC addresses, addresses of hosts that are running the FSM services. This is in contrast to the `fsnameservers` file, where the name server coordinators specified may or may not also be acting as MDCs.


In the case that HA is present, you would specify both the active and the standby MDCs that are hosting FSMs for the file system in the `fsforeignservers` file.

No additional configuration is needed on the MDCs which act as foreign servers. Foreign clients send heartbeat messages to the addresses in the `fsforeignservers` file. The heartbeat rate is once every 5 seconds. The MDC nodes reply to these heartbeat with a list of local, active FSMs and the address by which they may be reached.

After the `fsforeignservers` file has been created, services can be restarted and the file systems available through this service may be mounted. All the usual requirements of a file system client apply. The client must have access to the primary storage disks or use the distributed LAN client mount option.

Note: For the HA setup, the `ha_vip` address can be entered in the `fsforeignservers` file.

Figure 6 StorNext Foreign Servers

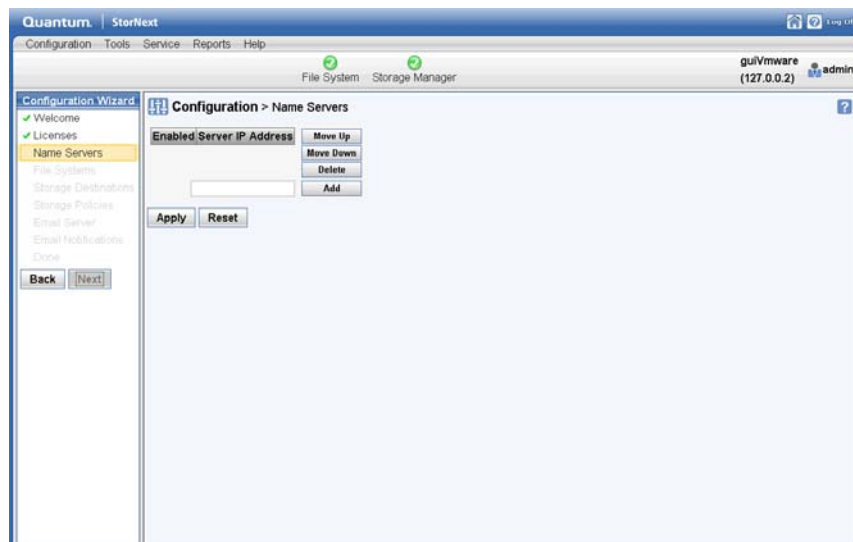


Entering Name Servers

Follow these steps to add a name server:

- 1 When the Configuration Wizard is displayed, choose **Name Servers** on the left side of the screen. (Alternatively, choose **Name Servers** from the **Configuration** menu.) The **Configuration > Name Servers** screen appears. (If name servers were previously created, a list of those IP addresses appears on the Name Servers screen.)

Figure 7 Name Servers Screen



- 2 To add a new name server, enter the IP address in the field to the left of the **Add** button. The new name server appears in the list of available name servers.
- 3 Click **Apply** to use the name server specified.
- 4 When the confirmation message warns you that changing the name server is a cluster-wide event, click **Yes** to continue or **No** to abort.
- 5 After you click **Yes**, a message informs you that Name Servers has been updated. Click **OK** to continue.
- 6 If there are previously configured name servers, you can specify the order in which name servers are used. To set the order, select a server and then click **Move Up** or **Move Down** until the selected server is in the correct order.

A green check mark icon under the **Enabled** column heading indicates that the server is currently enabled as a name server. A red X icon indicates that the server is not currently enabled.

Deleting a Name Server

To delete a name server, select the name server you want to delete and then click **Delete**. Finalize the deletion by clicking **Apply**.

Step 4: File Systems


After entering license information, the next task is to create a file system. When you reach this step, any previously created file systems are displayed.

The following procedure describes the easiest way to create a new file system quickly so you can complete the Configuration Wizard tasks and begin using StorNext immediately. For more detailed information about file systems such as editing, modifying or deleting existing file systems or performing additional file system actions, see [File System Tasks](#) on page 63.

- 1 When the Configuration Wizard is displayed, choose **File Systems** on the left side of the screen. (Alternatively, choose **File Systems** from the **Configuration** menu.) The **Configuration > File Systems** screen displays all currently configured file systems. (If you are running the Configuration Wizard for the first time, there will be no existing file systems.)

From this screen you can view, add, edit, or delete a file system. For information on these procedures, see the online help.

Figure 8 Configuration > File System Screen



- 2 Click **New** to add a new file system. The **Configuration > File Systems > New** Screen appears.

Figure 9 Configuration > File System > New Screen

Quantum StorNext

Configuration Tools Service Reports Help

File System Storage Manager

guiVmware (127.0.0.2) admin

Configuration Wizard

- ✓ Welcome
- ✓ Licenses
- ✓ Name Servers
- File Systems**
- Storage Destinations
- Storage Policies
- Email Server
- Email Notifications
- Done

Back Next

Configuration > File Systems > New

*File System Name

*Mount Point

Storage Manager ☐

Replication / Deduplication ☐

Stripe Group Configuration ☒ Generated ☐ Manual

Cancel Continue...

* Required Field

- 3 Enter the following fields:

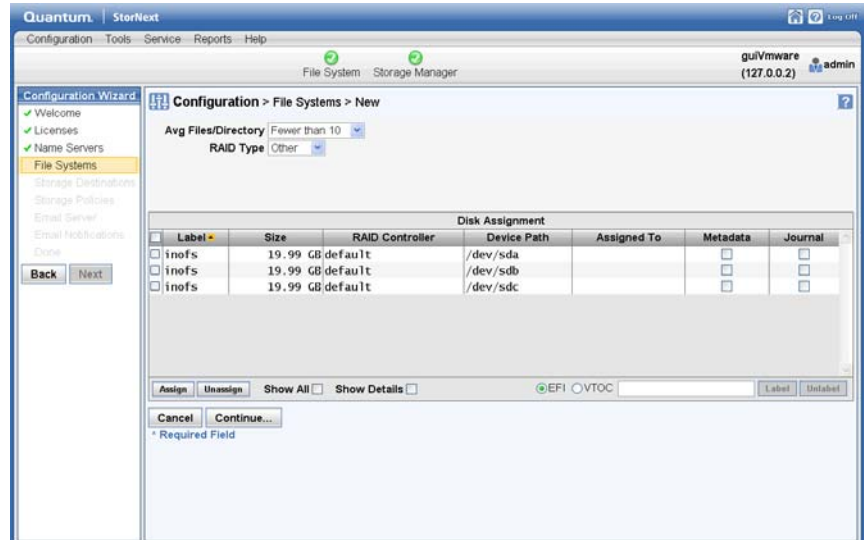
- **File System Name**
- **Mount Point**
- **Storage Manager**
- **Replication/Deduplication**
- **Stripe Group Configuration: Generated or Manual**

(For detailed information about what to enter on this screen, see the online help.)

If you chose Manual Configuration, skip to [Manual Configuration](#).

- 4 Click **Continue** to proceed to the second configuration screen.

Figure 10 Configuration > File System > New Screen 2



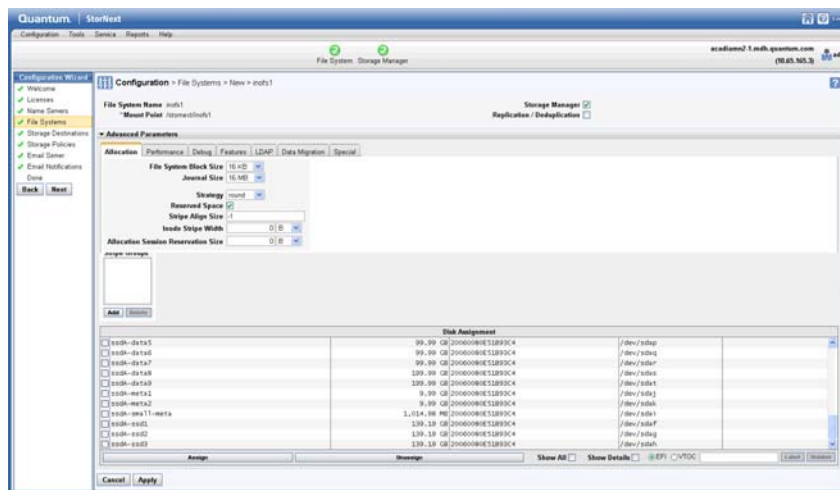
- 5 At the **Avg Files/Directory** field, select one of the options to indicate the approximate average number of files per directory you anticipate the new file system will contain. Options range from fewer than 10 files to more than 1000 files.

This value helps StorNext calculate capacity. You should select the option that most closely corresponds to your projected usage. You are not limited by the option you select. For example, if you indicate that there will typically be fewer than 10 files per directory, there is nothing to prevent the directory from containing more than 10 files.

- 6 At the **RAID Type** field, select the type of RAID the file system uses.
- 7 Select one or more disks to assign to the file system.
- 8 After selecting one or more disks, click the **Metadata** checkbox to designate the disk(s) as used for metadata, or click the **Journal** checkbox to use the disk(s) for journaling. A disk can be used for both metadata and journaling.
- 9 At the field to the left of the **Label** button, enter a disk label name. Click **Label** to apply the label name to the selected disks. (Clicking **Unlabel** removes label names from selected disks.) When asked to confirm the action (either Label or Unlabel), click **OK** to continue or **Cancel** to abort.

- 10 After you are finished entering label information, click **Assign** to assign the disk(s) to the file system. (Clicking **Unassign** removes any existing associations between disks and the file system. For example, if you assign disks erroneously, clicking **Unassign** is an easy way to remove associations and reassign disks.)
- 11 Click **Continue**.

Figure 11 Configuration > File System > New Screen 3

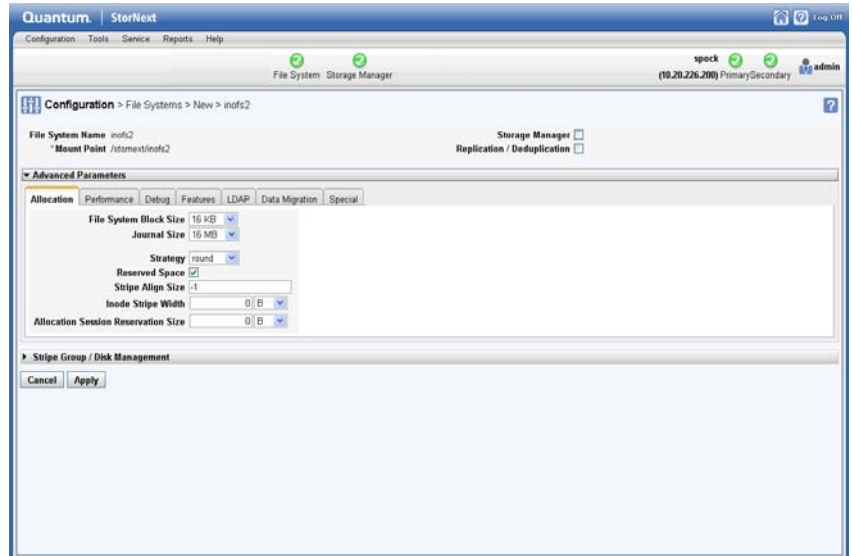


- 12 Click the arrows beside the headings **Advanced Parameters** and **Stripe Group/Disk Management** to display that information. If desired, make any changes in these areas.
- 13 When you are satisfied with the file system parameters, click **Apply**. StorNext automatically configures and mounts the file system based on the information you entered.

Manual Configuration

If you chose Manual Configuration, you must enter the fields on the **Advanced Parameters** tabs and the **Stripe Group/Disk Management** fields. (Click the arrow to the left of these headings to display the tabs and fields.)

Figure 12 Configuration > File System > New Screen 3



For information about entering these tabs and fields, see the online help.

- 1 When you are finished entering Advanced Parameter and Stripe Group/Disk Management information for the manually configured file system, click **Apply** to save your changes and create the file system.
- 2 When a message informs you that the file system was successfully created, click **OK**.

Allocation Session Reservation

The Advanced Parameters section of this screen contains a field called **Allocation Session Reservation Size**. The Allocation Session Reservation feature optimizes on-disk allocation behavior. Allocation requests occur whenever a file is written to an area that has no actual disk space allocated, and these requests are grouped into sessions. The amount you specify at this field determines the size of the chunk of space reserved for a session.

For more information about the Allocation Session Reservation feature, refer to the StorNext File System Tuning Guide.

Editing a File System
and Performing Other
Actions

Once you have created a file system, you have the option of editing the file system or performing one of the following actions:

Stop or start the file system	Update the file system
Start and activate the file system	Perform a metadata dump for the file system
Start and mount the file system	Check the file system
Activate the file system	Expand the file system Note: StorNext does not support expansion on stripe groups containing mixed-sized LUNs. For example, if you create a file system that has two different-sized disks in a userdata only stripe group and then attempt to add a new disk to that stripe group and expand it, the expansion will fail.
Mount or unmount the file system	Migrate the file system
Make the file system	

Refer to the StorNext online help for instruction on editing a file system or performing one of the available actions.

Step 5: Storage Destinations

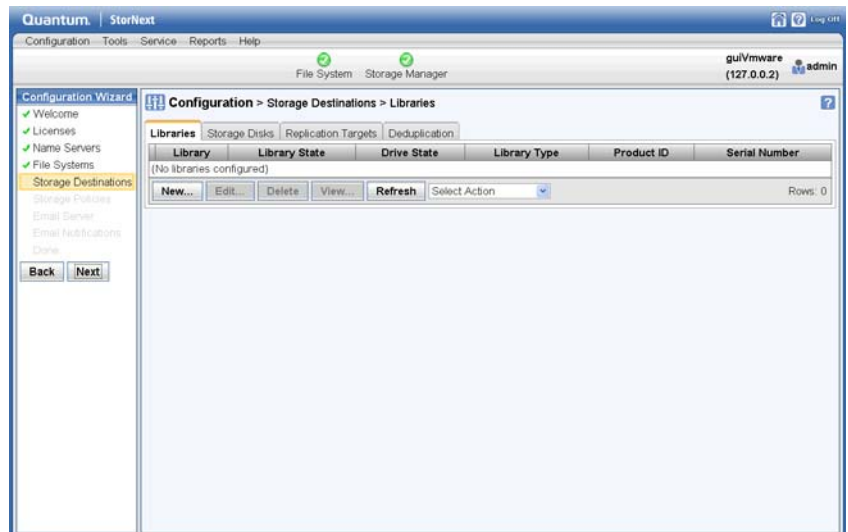
After you have created at least one file system, the Configuration menu's Storage Destinations option allows you to add, edit, or delete libraries and storage disks. You can also enter or edit targets for data replication, and specify a blockpool host file system for data deduplication.

Adding a New Library

Follow this procedure to add a new library:

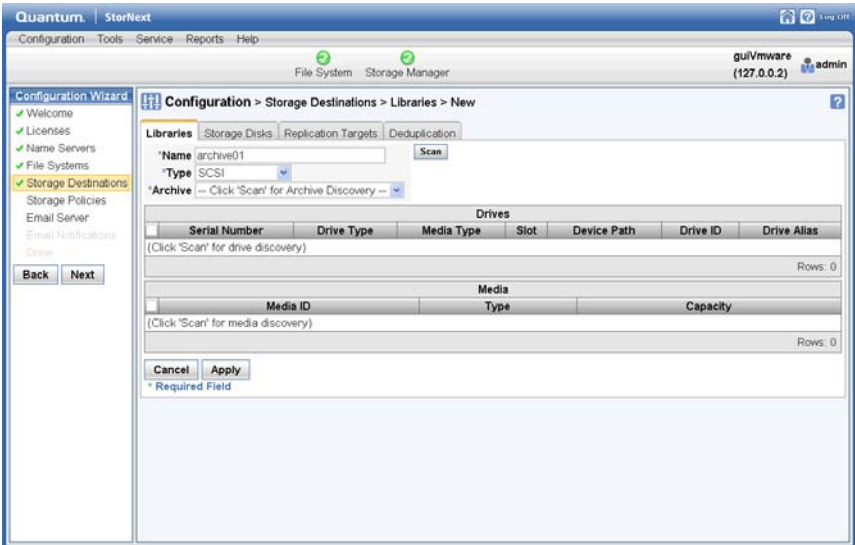
- 1 When the Configuration Wizard is displayed, choose **Storage Destinations** on the left side of the screen. (Alternatively, choose **Storage Destinations** from the **Configuration** menu.)
- 2 If necessary, click the **Library** tab. The **Configuration > Storage Destinations > Library** Screen appears.

Figure 13 Storage Destinations
> Library Screen



- 3 Click **New**. The **Configuration > Storage Destinations > Library > New** Screen appears.

Figure 14 Storage Destinations
> Library > New Screen



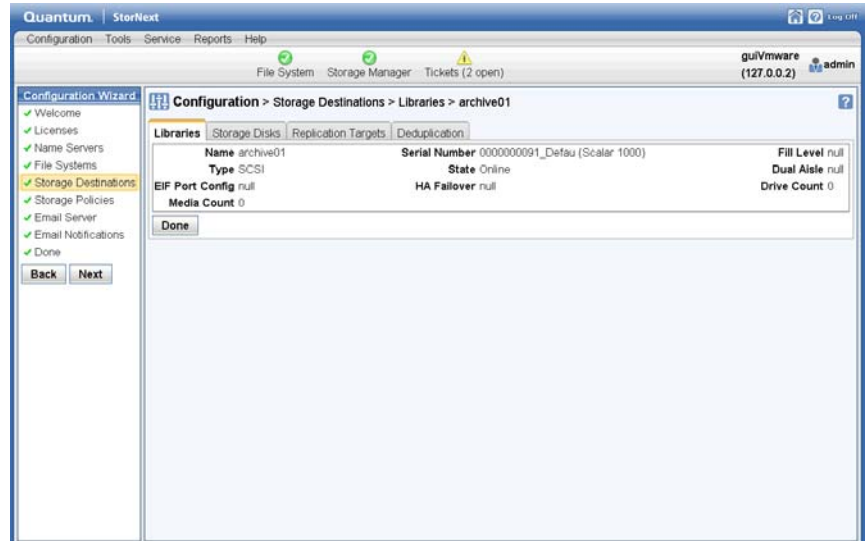
- 4 Enter the fields at the top of the screen. (For detailed information about what to enter on this screen, see the online help.)
- 5 Click **Scan** to have StorNext discover available drives and media.
- 6 Select a tape drive or drives to add to your new library.
- 7 In the **Media** section, view media available for use.
- 8 Click **Apply**.
- 9 After a message informs you that the library was successfully created, click **OK**.
- 10 Repeat steps 3- 9 to add additional tape drives and media to the new library.

Viewing an Existing Library

- Follow this procedure to view details for a previously created library:
- 1 Choose **Storage Destinations** from the **Configuration** menu. If necessary, click the **Library** tab. The **Configuration > Storage Destinations > Library** screen appears. (See [Figure 13.](#))

- 2 Select the library whose information you want to view.
- 3 Click **View**, or choose **View** from the actions dropdown list. The library detail screen appears.

Figure 15 Library Details Screen



- 4 The library information screen provides the following information:
 - **Name:** The name of the library
 - **Serial Number:** The library's serial number
 - **Fill Level:** The library's current fill level
 - **Type:** The type of library (e.g., SCSI, ACSLS, etc.)
 - **State:** The library's current state (e.g., online or offline)
 - **Dual Aisle:** Indicates whether the library has a dual aisle configuration
 - **EIF Port Config:** The current EIF port configuration. EIF stands for Enterprise Instrumentation Framework, and it helps StorNext process data from applications on the server.
 - **HA Failover:** Indicates whether HA failover is enabled for the library
 - **Drive Count:** The number of tape drives in the library

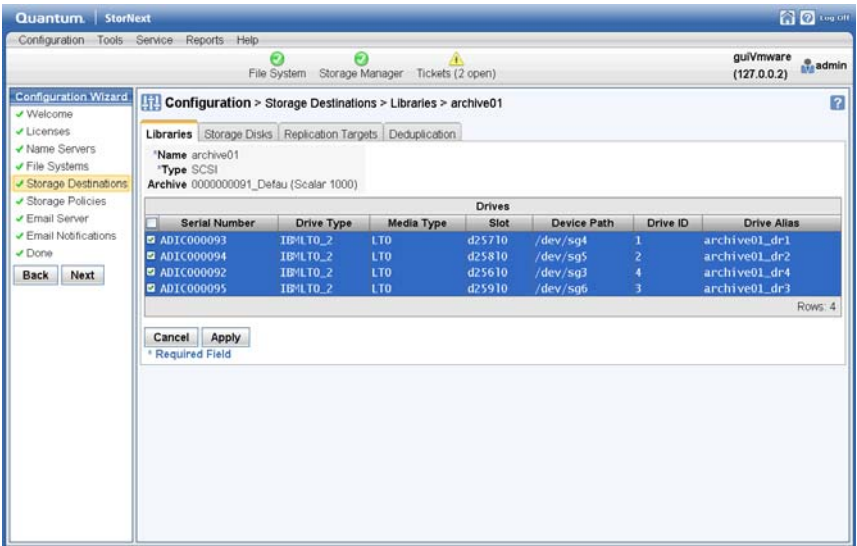
- **Media Count:** The number of media in the library
- 5 When you are finished viewing library information, click **Done**.

Editing a Library

Follow this procedure to edit parameters for an existing library:

- 1 If you have not already done so, choose **Storage Destinations** from the **Configuration** menu and then click the **Library** tab.
- 2 Select the library you want to edit.
- 3 Click **Edit**, or choose **Edit** from the actions dropdown list. After you select this option StorNext scans the library, which could take some time to complete depending on your configuration.

Figure 16 Edit Library Screen



- 4 If desired, click **Scan** if you want StorNext to scan the library for drives.
- 5 Select the tape drives you want included in your library, or click **All** to include all available tape drives. (To exclude all drives, click **None**.)
- 6 Click **Apply**.
- 7 When a confirmation message appears, click **Yes** to proceed, or **No** to abort.

- 8 After a message informs you that the library was successfully modified, click **OK**.

Deleting a Library

Follow this procedure to delete an existing library:

- 1 If you have not already done so, choose **Storage Destinations** from the **Configuration** menu and then click the **Library** tab.
- 2 Select the library you want to delete.
- 3 Click **Delete**, or choose **Delete** from the actions dropdown list.
- 4 When a confirmation message appears, click **Yes** to proceed with the deletion or **No** to abort.
- 5 After a message informs you that the library was successfully deleted, click **OK**.

Performing Other Library Actions

Towards the middle of the **Configuration > Storage Destinations > Library** screen is a dropdown list of actions you can perform for libraries.

Select the library for which you want to perform the action, and then choose one of these options from the **Select Action** dropdown list:

- **Audit:** Select this option to perform an audit on the selected 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 periodic audits on the library to ensure synchronization.
- **Remap-Audit:** This option synchronizes the StorNext databases with the library databases.
- **Online:** Select this option to set the library online.
- **Offline:** Select this option to take the library offline.
- **Drives Online:** Select this option to place the drives in the library online.
- **Drives Offline:** Select this option to take the drives in the library offline.
- **Add Media Bulkload:** Select this option to add media to the library via the bulk loading method.

- **Add Media Mailbox:** Select this option to add media to the library through the library's mailbox.

Storage Disk Overview

Storage disks are external devices on UNIX-based file systems that can be used for long term data storage. Storage disks function and operate the same way as physical tape media. You can add up to 16 storage disks.

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:

- 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.

The following procedures describe how to view, edit and delete storage disks. The procedure for adding a storage disk is identical to entering one through the StorNext Configuration Wizard as described in [Adding a New Storage Disk](#) on page 41.)

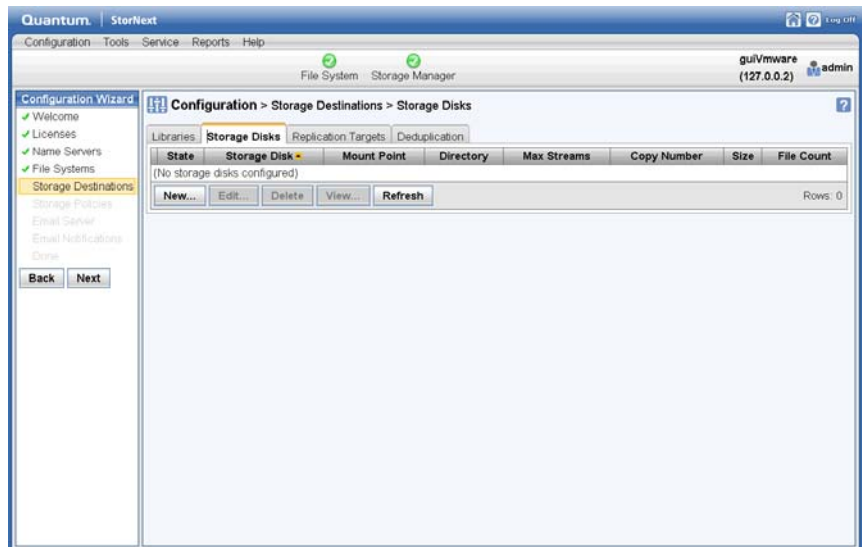
Caution: Storage disks can be an important and integral part of your system, but you should NEVER rely solely on storage disks as your only backup location.

Adding a New Storage Disk

Follow this procedure to add a new storage disk.

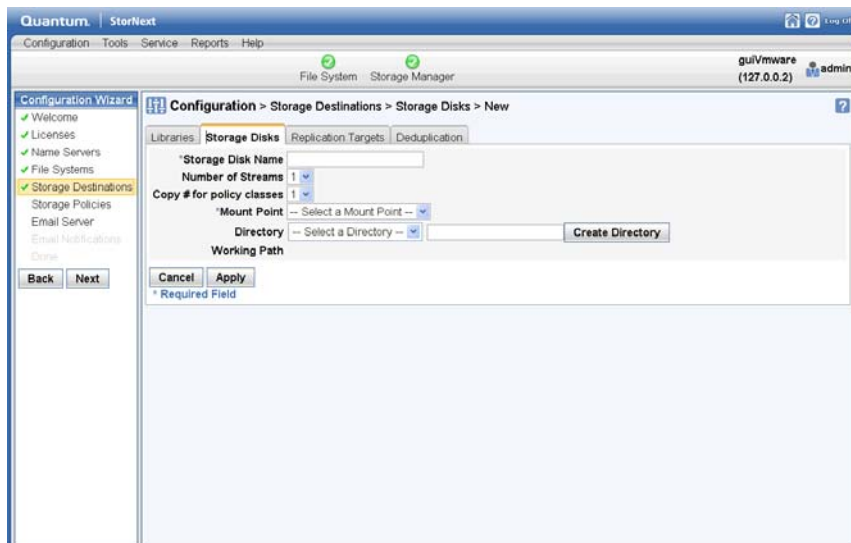
- 1 Click the **Storage Disk** tab. The **Configuration > Storage Destinations > Storage Disk** Screen appears.

Figure 17 Configuration > Storage Destinations > Storage Disk Screen



- 2 Click **New**. The **Storage Destinations > Storage Disk > New** Screen appears.

Figure 18 Storage Destinations
> Storage Disk > New Screen



- 3 Enter the fields on the screen. (For detailed information about what to enter on this screen, see the online help.)
- 4 Click **Apply**.
- 5 Repeat steps 2 - 4 to add additional storage disks.

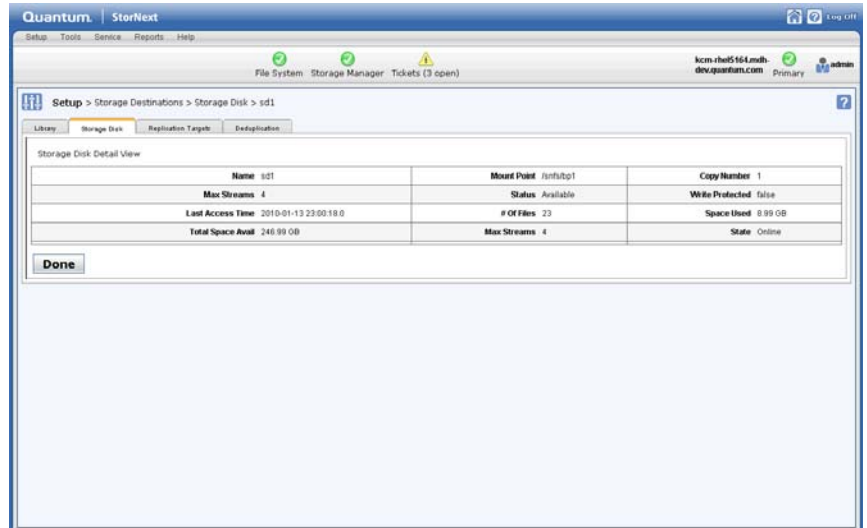
Viewing an Existing Storage Disks

Follow this procedure to view a list of previously configured storage disks.

- 1 Choose **Storage Destinations** from the **Configuration** menu.
- 2 Click the **Storage Disk** tab. Any previously configured storage disks are displayed. (See [Figure 17](#) on page 41.)
- 3 Select the storage disk whose information you want to view.

4 Click **View**.

Figure 19 View Storage Disk Screen



5 When you are finished viewing library information, click **Done**.

Editing a Storage Disk

Follow this procedure to edit a currently configured storage disk.

- 1 If you have not already done so, choose **Storage Destinations** from the **Configuration** menu and then click the **Storage Disk** tab.
- 2 Select the storage disk whose information you want to edit.
- 3 Click **Edit**.
- 4 Modify any of the fields you entered when creating the storage disk. (For field information, see the online help or the descriptions in [Adding a New Storage Disk](#) on page 41.)
- 5 Click **Apply**.
- 6 When a confirmation message appears, click **Yes** to proceed, or **No** to abort.
- 7 After a message informs you that the storage disk was successfully modified, click **OK**.

Deleting a Storage Disk

Follow this procedure to delete a currently configured storage disk.

- 1 If you have not already done so, choose **Storage Destinations** from the **Configuration** menu and then click the **Storage Disk** tab.
- 2 Select the storage disk you want to delete.
- 3 Click **Delete**.
- 4 When a confirmation message appears, click **Yes** to proceed with the deletion or **No** to abort.
- 5 After a message informs you that the storage disk was successfully deleted, click **OK**.

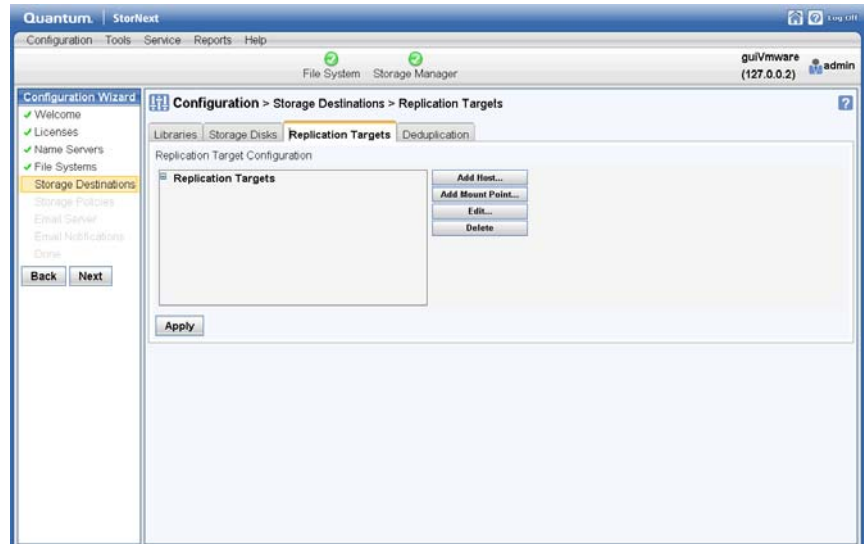
Caution: Before deleting a storage disk, make sure the policies that refer to the storage disk (as well as the file systems that use the storage disk for backup storage) have been safely backed up and configured or updated to not use the storage disk you want to delete.

Adding a New Data Replication Target

The following procedure describes how to add a new replication target host. (For more information about the replication feature, see [Chapter 6, Replication and Deduplication](#) and [Appendix B, Additional Replication and Deduplication Information](#).)

- 1 Click the **Replication Targets** tab. The **Configuration > Storage Destinations > Replication / Deduplication** Screen appears.

Figure 20 Configuration >
Storage Destinations >
Replication / Deduplication
Screen



- 2 Click **Add Host**.
- 3 Enter the fields in the **New Target** section. (For detailed information about what to enter on this screen, see the online help.)
- 4 Click **Add** to add the new replication target, or **Cancel** to abort without saving.
- 5 Click **Apply** to save your changes.
- 6 Repeat steps 3 - 6 to add more targets.
- 7 After a message informs you that the target was successfully added, click **OK**.

Editing a Data Replication Host

Follow this procedure to edit an existing data replication target.

- 1 If you have not already done so, click the **Replication Targets** tab.
- 2 If necessary, click the plus sign (+) beside the **Replication Targets** heading in the box titled **Replication Target Configuration**.
- 3 Select the replication target you want to edit.
- 4 Click **Edit**.

- 5 At the **Hostname or IP** field, modify either the host name or IP address for the replication target.
- 6 Click **Update** to save your changes, or **Cancel** to abort.

Deleting a Data Replication Target

Follow this procedure to delete a replication target.

- 1 If you have not already done so, choose click the **Replication Targets** tab.
- 2 If necessary, click the plus sign (+) beside the **Replication Targets** heading in the box titled **Replication Target Configuration**.
- 3 Select the replication target you want to delete.
- 4 Click **Delete**.

Caution: There is no confirmation message for this delete function, so make absolutely certain you want to delete the replication target before you click **Delete**.

Adding a New Mount Point

Follow this procedure to add a new mount point to a replication target.

- 1 If you have not already done so, choose **Storage Destinations** on the left side of the screen. (Alternatively, choose **Storage Destinations** from the **Configuration** menu.)
- 2 Click the **Replication Targets** tab.
- 3 Select the replication target (host) to which you would like to add a mount point. (You might need to click the dash to the left of the **Replication Targets** heading to display the available hosts.)
- 4 Click **Add Mount Point**.
- 5 Click **Scan Host** to identify available mount points on the selected host.
- 6 At the **Mount Point** field, select a mount point and then click **Add**.
- 7 Repeat steps 3 - 6 to add additional mount points.
- 8 Click **Apply** to save the changes.

- 9 After a message informs you that changes were successfully incorporated click **OK**.

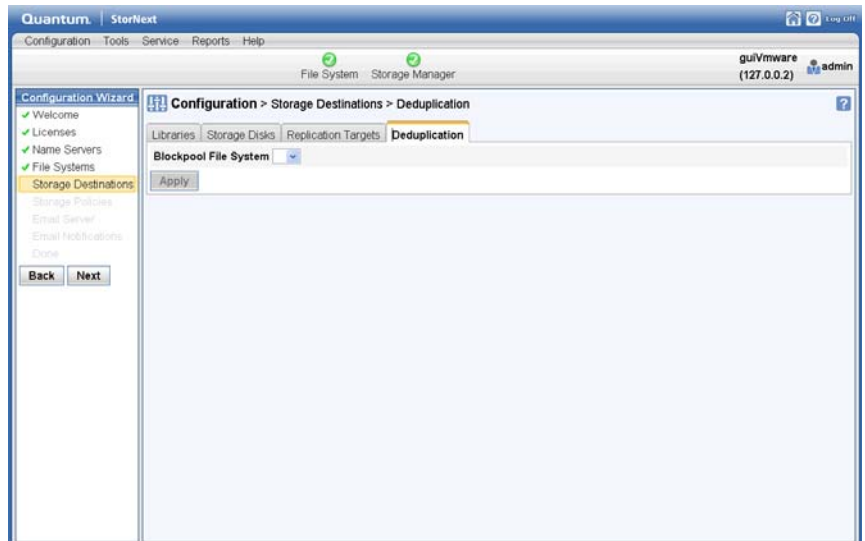
Enabling Data Deduplication

The **Deduplication** tab enables you to create a blockpool on a specified file system. (For more information about the deduplication feature, see [Chapter 6, Replication and Deduplication](#) and [Appendix B, Additional Replication and Deduplication Information](#).)

To create the blockpool, select the desired file system from the dropdown list next to the **Blockpool File System** label, and then click **Apply**.

Note: The blockpool should not be placed on a file system that will be used as the HA shared file system. This is a requirement even if you do not plan to use the StorNext Deduplication feature.

Figure 21 Configuration > Storage Destinations > Replication / Deduplication Screen (Blockpool)



Step 6: Storage Policies

There are two kinds of storage policies: Storage Manager storage policies and Replication/Deduplication storage policies. Replication/Deduplication storage policies control the way StorNext's replication and deduplication features behave. (For more information about the replication and deduplication features, see [Chapter 6, Replication and Deduplication](#) and [Appendix B, Additional Replication and Deduplication Information](#).)

A Storage Manager storage policy defines how files will be managed in a directory and subdirectories. Specifically, these are the available Storage Manager storage policy 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

Storage policies can be related to one or more directories. In this situation, all files in that directory and sub-directories are governed by the storage policy.

Note: The connection between a storage policy and a directory is called the relation point.

Here are some examples of storage policy usage:

- A directory in which to store backups every night is created. This directory is seldom accessed after the files are copied over. A storage policy could be set up to create two tape copies of the files, store one copy of the files to LTO 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.

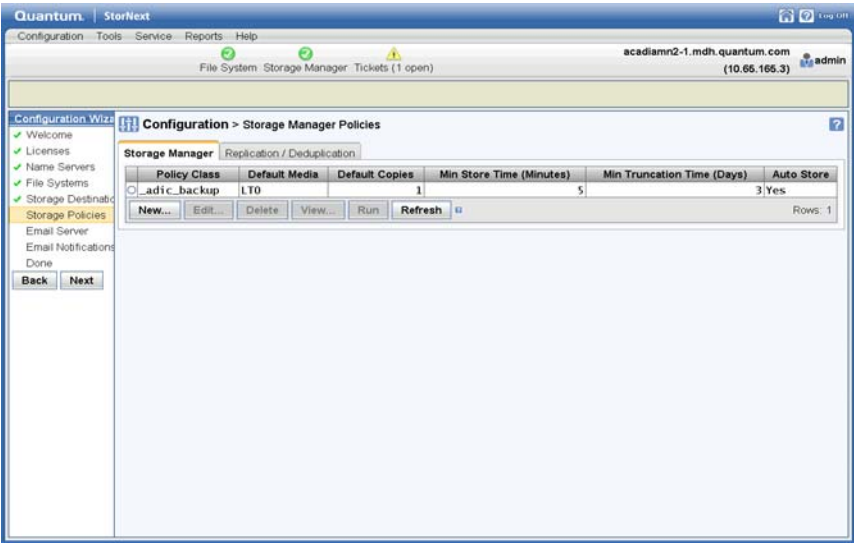
The in this case could be set up to create a single tape copy, store the files to LTO media 15 minutes after being on disk, and then truncate after 60 days of non-use. This policy can be associated with a directory such as: /sandsm/dsm1/docs.

Adding a Storage Manager Storage Policy

Follow this procedure to add a new Storage Manager storage policy:

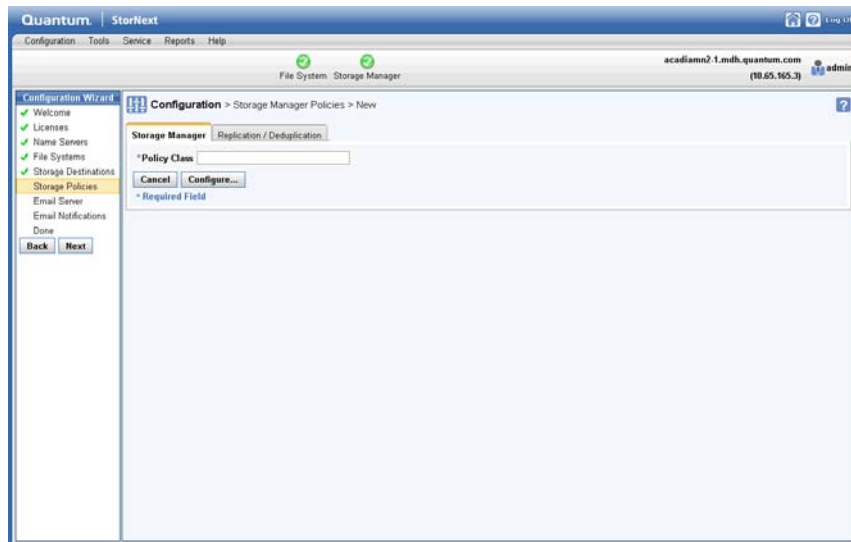
- 1 When the Configuration Wizard is displayed, choose **Storage Policies** on the left side of the screen. (Alternatively, choose **Storage Policies** from the **Configuration** menu.) The **Configuration > Storage Policies** Screen appears.

Figure 22 Configuration > Storage Policies Screen



- 2 Click **New**. The **Storage Policies > New** Screen appears.

Figure 23 Storage Policies >
New Screen



3 Enter the following fields:

- **Policy Class:** The name of the new policy you are creating.

Note: The policy class name must be unique. You cannot enter the name of an existing policy class.

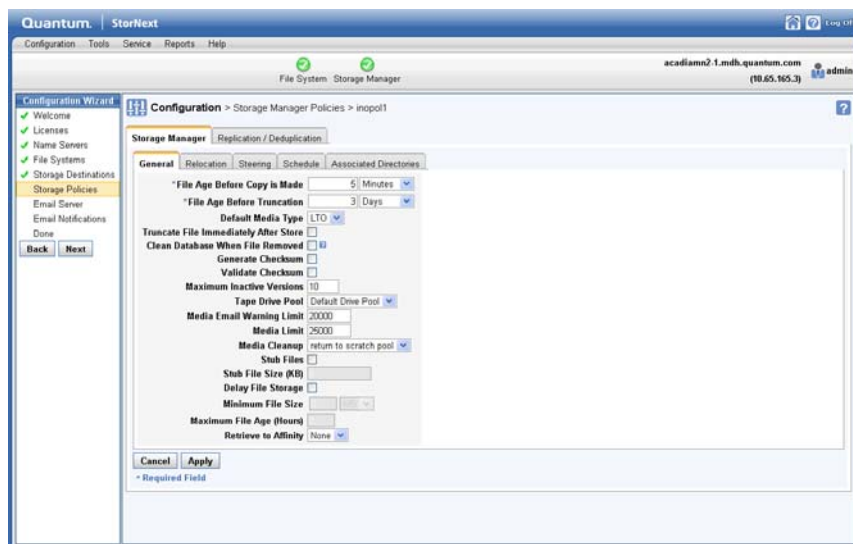
Also, if you use upper-class characters when entering the policy class name, the name will be converted to lower-case characters when the policy class is created.

- **Policy Type:** click the **Storage Manager** tab to create a policy for StorNext Storage Manager
 - Click **Configure** to continue.
- 4 Enter information on the **General**, **Relocation**, **Steering**, **Schedule** and **Associated Directories** tabs. (See the sections following for more information about these tabs.)
- 5 When you are finished entering information about the new policy, click **Apply**, or click **Cancel** to exit without saving.
- 6 After the Status screen informs you that the policy was created successfully, click **OK**.

The General Tab

The General tab contains parameters that apply to all storage policies. Fields marked with an asterisk are required. Enter additional fields as desired, or accept the displayed default values.

Figure 24 Storage Policies > New > General Tab



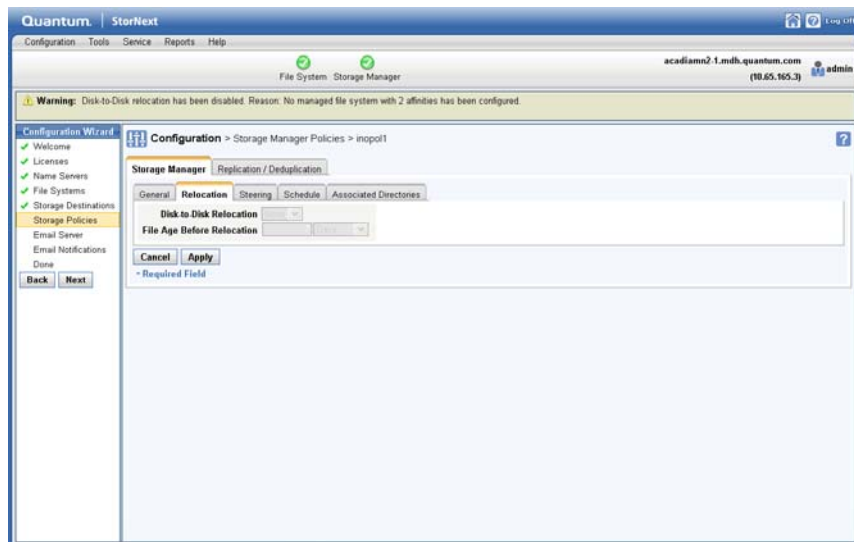
For instructions on what to enter on this screen, see the online help.

The Relocation Tab

The Relocation tab enables you to configure the Disk-to-Disk relocation feature.

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. In order to use this feature you must have a managed file system with at least two affinities configured.

Figure 25 Storage Policies >
New > Relocation Tab

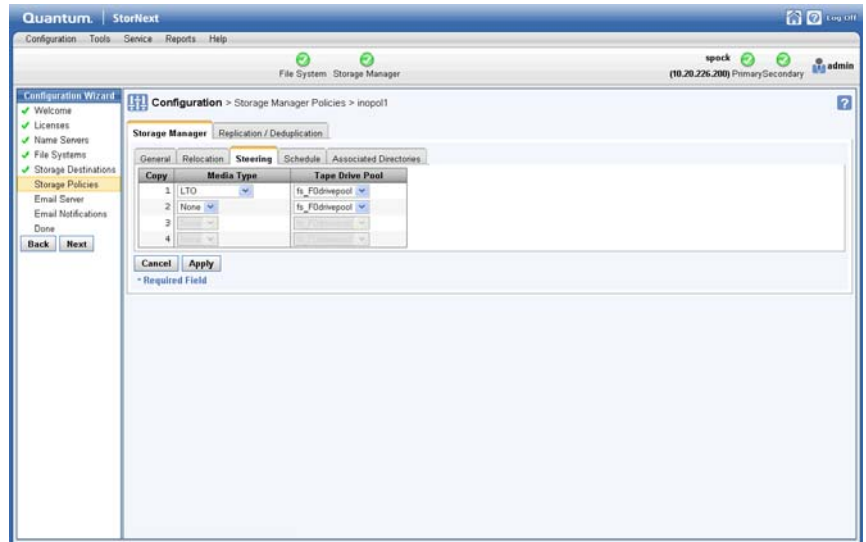


For instructions on what to enter on this screen, see the online help.

The Steering Tab

The Steering tab enables you to configure *file steering*, which allows you to direct a copy of a file to a designated drive pool. This is normally used when you want to direct two or more copies of a file to different archives by putting the tape drive in separate pools and then setting the copy number of the file to go to that pool. You can also use this feature to route your copies of the file to different media types, including storage disks.

Figure 26 Storage Policies >
New > Steering Tab



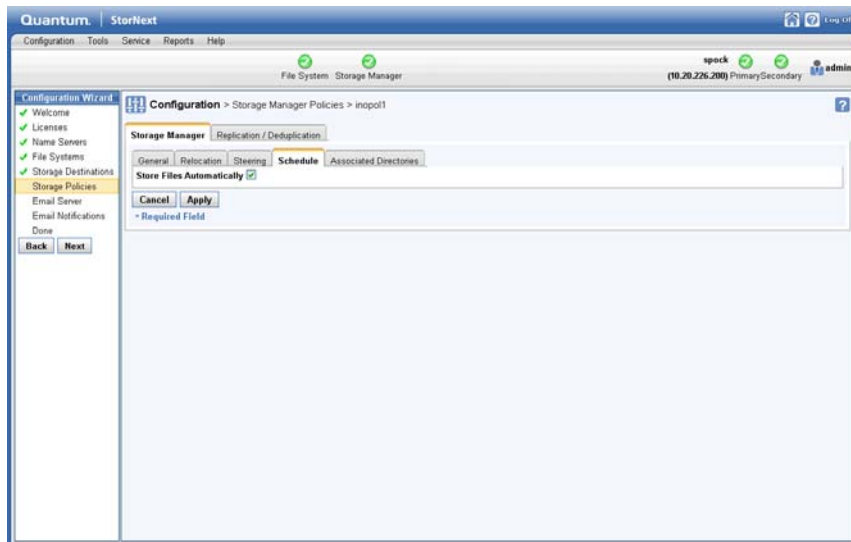
For instructions on what to enter on this screen, see the online help.

The Schedule Tab

The Schedule tab allows you to enable or disable the Store Files Automatically feature.

When this feature is enabled, StorNext automatically stores files for the current storage policy. If this feature is disabled, Quantum recommends that the files for the policy class be stored by scheduled events. (Scheduled events are certain activities which you can set up to run at specified times using StorNext's schedule. For more information, see [Scheduler](#) on page 114.)

Figure 27 Storage Policies >
New > Schedule Tab

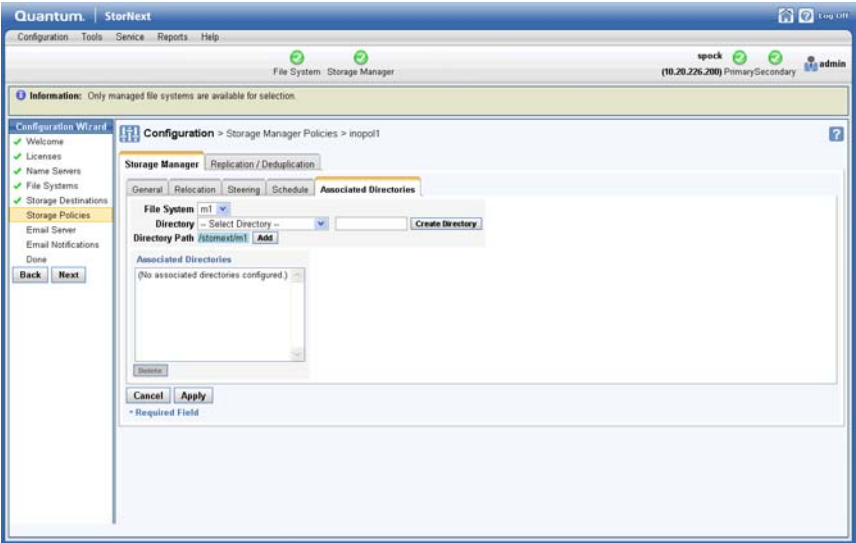


For instructions on what to enter on this screen, see the online help.

The Associated Directories Tab

The Associated Directories tab enables you to view or delete any existing associated directories in the file system for the policy, and to add new directories.

Figure 28 Storage Policies > New > Associated Directories Tab



For instructions on what to enter on this screen, see the online help.

Adding a Replication Storage Policy

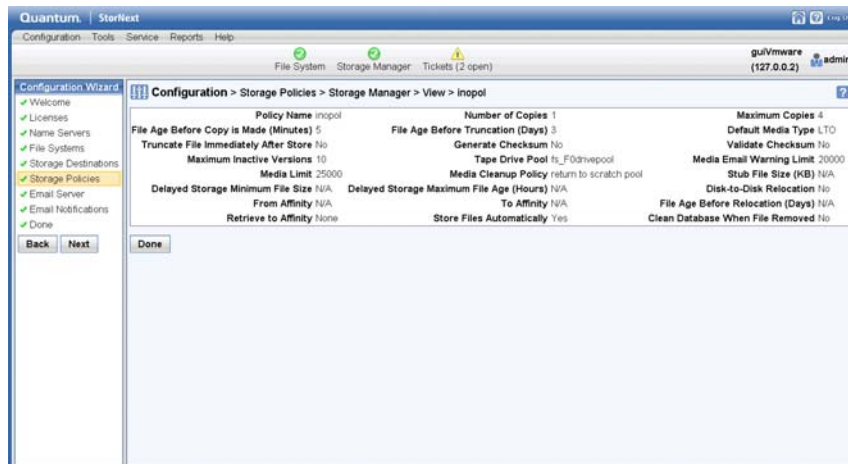
The steps for creating a replication storage policy are described in [Step 4: Create a Replication Storage Policy](#) on page 152.

Viewing a Storage Policy

To view storage policy details To view storage policy details for a Storage Manager or Replication policy, do the following:

- 1 From the **Configuration > Storage Policies** screen, select the storage policy you wish to view.
- 2 Click **View**.

Figure 29 View Storage Policies Screen



- 3 Click **Done** to return to the **Configuration > Storage Policies** screen.

Running a Storage Policy

Follow this procedure to run an existing storage policy.

- 1 If you have not already done so, choose **Storage Policies** from the **Configuration** menu.
- 2 Select the policy you want to run, and then click **Run**.
- 3 When a message informs you that the job was successfully initiated, click **OK** to continue.
- 4 To view job progress, select **Jobs** from the **Reports** menu.

Editing a Storage Policy

To edit an existing storage policy:

- 1 From the **Configuration > Storage Policies** screen, select the policy you wish to edit.
- 2 Click **Edit**.
- 3 Modify policy information as desired by clicking the tabs and editing or adding information. The process is the same as when you first created the policy.

If you are editing a Storage Manager policy, you can edit fields on the **General**, **Relocation**, **Steering**, **Schedule** and **Associated Directories** tabs. For more information about fields on these tabs, see the online help.

If you are editing a Replication global policy, you can edit fields on **Deduplication**, **Truncation**, **Outbound Replication**, **Inbound Replication**, and **Blackout** tabs. If you are editing a Replication target policy, you can modify field on only the **Inbound Replication** tab. For more information about fields on these tabs, see the online help.

- 4 Click **Apply** to save changes and return to the **Configuration > Storage Policies** screen, or **Cancel** to abort.

Deleting a Storage Policy

To delete an existing storage policy:

- 1 From the **Configuration > Storage Policies** screen, select the policy you wish to delete.
- 2 Click **Delete**.
- 3 Click **Yes** to confirm the deletion, or **No** to cancel.

Step 7: Email Server

The Email Server option allows you to specify the email server used for processing StorNext notification email messages. On this screen you will enter basic information such as the email server name and sending entity. You also have the option of sending a test message so you can verify that StorNext recognizes the email server whose information you entered.

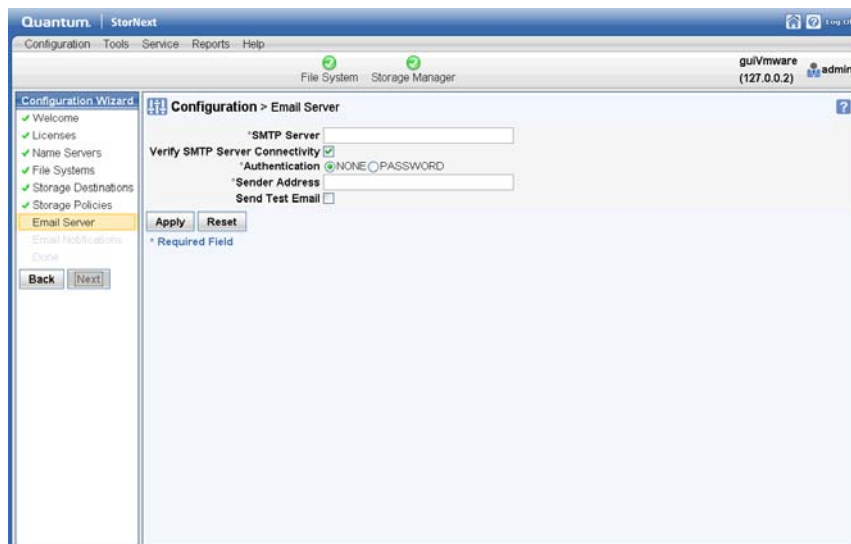
Note: The Email Server option does not configure your email server. Instead, it allows you to specify a previously configured email server so StorNext knows which server is responsible for processing notification messages. Before you use the Email Server option, make sure your email SMTP server is already configured.

Adding an Email Server

Follow this procedure to add a new email server.

- 1 When the Configuration Wizard is displayed, choose **Email** on the left side of the screen. (Alternatively, choose **Email Server** from the **Configuration** menu.) The **Configuration > Email Server** Screen appears.

Figure 30 Configuration > Email Server Screen



- 2 Complete the fields related to your email system configuration on the **Configuration > Email Server** screen. (For detailed information about what to enter on this screen, see the online help.)
- 3 Click **Apply** to save your changes.

Step 8: Email Notification

The Email Notification feature allows you to specify parties who should receive StorNext email messages about backup statuses, service tickets, admin alerts, and policy class messages.

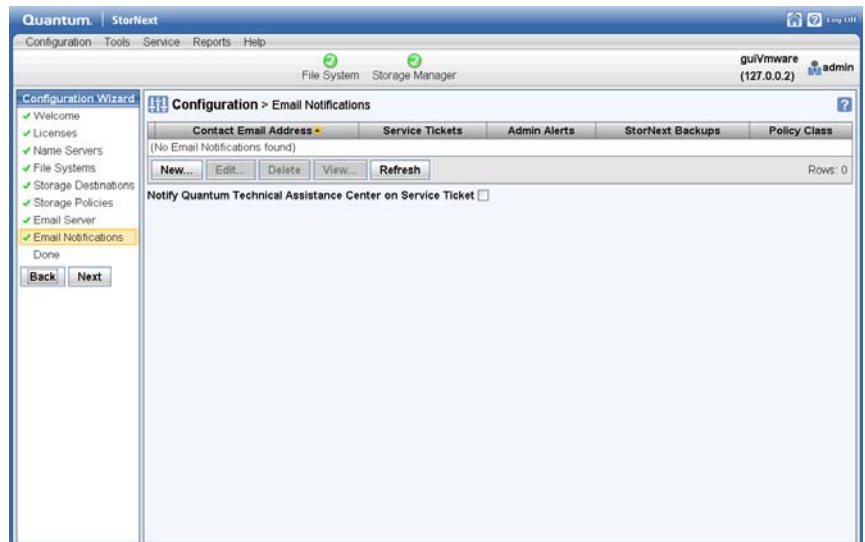
Note: In order for this feature to work properly, make sure you have specified a configured email server as described in [Adding an Email Server](#) on page 58.

Adding an Email Recipient

Follow this procedure to add a new email recipient.

- 1 Choose **Email Notifications** from the **Configuration** menu.
- 2 When the Configuration Wizard is displayed, choose **Email Notifications** on the left side of the screen. (Alternatively, choose **Email Notifications** from the **Configuration** menu.) The **Configuration > Email Notifications** Screen appears.

Figure 31 Configuration > Email Notifications Screen



- 3 Click **New**. The **Configuration > Email Notifications > New** screen appears.

Figure 32 Configuration > Email Notifications New Screen

Quantum StorNext

Configuration Tools Service Reports Help

File System Storage Manager

guiVmware (127.0.0.2) admin

Configuration Wizard

- ✓ Welcome
- ✓ Licenses
- ✓ Name Servers
- ✓ File Systems
- ✓ Storage Destinations
- ✓ Storage Policies
- ✓ Email Server
- ✓ **Email Notifications**
- Done

Back Next

Configuration > Email Notifications > New

* Contact Email Address

Contact Name

Admin Alerts ☐

StorNext Backups ☐

Service Tickets Disabled

Policy Class

Cancel Apply Reset

* Required Field

- 4 Complete the fields for the new email recipient. (For detailed information about what to enter on this screen, see the online help.)
- 5 Click **Apply** to save your changes.
- 6 When the confirmation message appears, click **Yes** to proceed or **No** to abort.
- 7 When a message informs you that the email notification recipient was successfully added, click **OK** to return to the **Configuration > Email Notifications** screen.

Viewing Email Recipient Information

Follow this procedure to view details for an existing email recipient.

- 1 If you have not already done so, when the Configuration Wizard is displayed, choose **Email Notifications** on the left side of the screen. (Alternatively, choose **Email Notifications** from the **Configuration** menu.)

- 2 On the **Configuration > Email Notifications** screen, review the list of current email recipients.
- 3 Select the recipient whose information you want to view, and then click **View**.
- 4 When you are finished viewing recipient information, click **Cancel** to return to the **Configuration > Email Notifications** screen.

Editing an Email Recipient

Follow this procedure to edit information for a previously entered email recipient.

- 1 If you have not already done so, when the Configuration Wizard is displayed, choose **Email Notifications** on the left side of the screen. (Alternatively, choose **Email Notifications** from the **Configuration** menu.)
- 2 On the **Configuration > Email Notifications** screen, select the recipient whose information you want to edit and then click **Edit**.
- 3 Modify any of the fields on the screen. (For detailed information about what to enter on this screen, see the online help.)
- 4 When you are finished making modifications, click **Apply** to save your changes and return to the **Configuration > Email Notifications** screen. (To exit without saving, click **Cancel**.)

Deleting an Email Recipient

Follow this procedure to delete a previously entered email recipient.

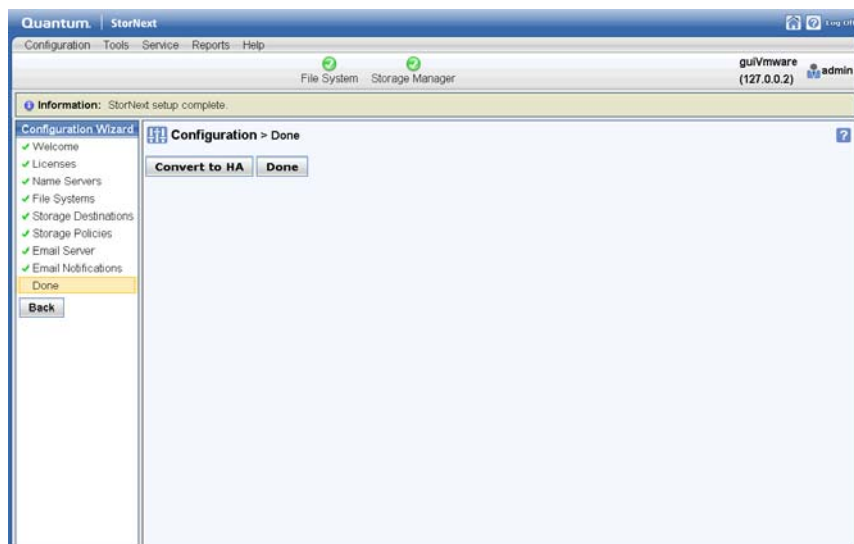
- 1 If you have not already done so, when the Configuration Wizard is displayed, choose **Email Notifications** on the left side of the screen. (Alternatively, choose **Email Notifications** from the **Configuration** menu.)
- 2 On the **Configuration > Email Notifications** screen, review the list of current email recipients.
- 3 Select the recipient you want to delete and then click **Delete**.
- 4 When the confirmation message appears, click **Yes** to proceed or **No** to abort the deletion.
- 5 When a message informs you that the email notification recipient was successfully deleted, click **OK** return to the **Configuration > Email Notifications** screen.


Step 9: Done

The last step in the Configuration Wizard is to click **Done** to indicate that you have completed all configuration steps.

On this screen you can also convert to a high availability (HA) configuration by clicking **Convert to HA**. Clicking this button is the same as choosing **High Availability > Convert** from the **Tools** menu. For information about entering the fields on this screen and converting to an HA system, see [Converting to HA](#) on page 221.

Figure 33 Configuration > Configuration Wizard Done Screen





Chapter 4

File System Tasks

In addition to the basic file system tasks described for the Configuration Wizard in [Step 4: File Systems](#) on page 29, the **Tools > File Systems** menu contains additional options that enable you to perform the following file system-related tasks:

- [Label Disks](#): Apply EFI or VTOC label names for disk devices in your StorNext libraries
- [Check File System](#): Run a check on StorNext files systems prior to expanding or migrating the file system
- [Affinities](#): Allocate additional storage to a file system by creating a new stripe group in the file system configuration file, and assigning new disks to the stripe group
- [Migrate Data](#): Move data files from a source file system to a destination stripe group, freeing stripe groups so they can be removed from an existing StorNext file system
- [Truncation Parameters](#): Enter truncation parameters for your file systems in order to free up file storage that isn't being actively used

Label Disks

Each drive used by StorNext 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.

There are two types of label:

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

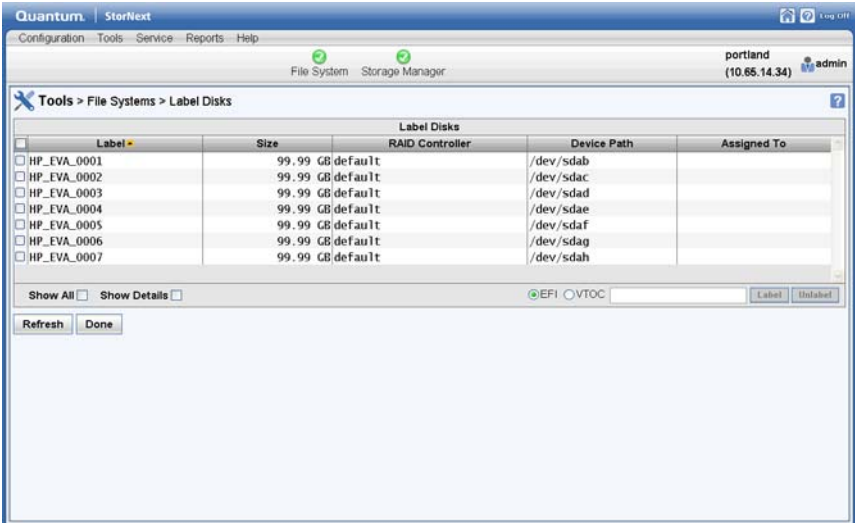
Labeling a Device

Follow this procedure to label any new or unused devices, or relabel a device that has been unlabeled.

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

- 1 Choose **Label Disks** from the **Tools > File Systems** menu. The **Tools > Label Disks** screen appears.

Figure 34 Label Disks Screen



- 2 Select the disk devices to which you want to apply labels. (Click **All** to select all available disks.) If a disk device already has a label, continuing with this procedure overwrites the existing label.

Caution: Overwriting or renaming a disk device label may result in a complete loss of data on that disk device.

- 3 Specify the label type by choosing **EFI** or **VTOC**.
- 4 Enter a label name in the text field to the right of the **EFI** and **VTOC** buttons.
- 5 Click **Label**.
- 6 When the confirmation message appears, verify that the disk you are labeling is empty, and then click **OK** to proceed. (Click **Cancel** to abort without labelling the disk.)

Note: If you later unlabel a device and then decide to make the unlabeled device usable by the StorNext File System, you must first relabel the device. The relabeling process is identical to labeling initially.

Unlabeling a Device

Follow this procedure to remove a label from a previously labeled device. If you unlabel a device and then decide later to make the 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](#).

Note: You cannot remove the label from a disk device that has been previously assigned to a file system. You can identify these devices by the file system name under the Filesystem heading.

- 1 If you have not already done so, choose **Label Disks** from the **Tools > File Systems** menu.
- 2 Select the disk devices from which you want to remove labels. (Click **All** to select all available disks.)
- 3 Click **Unlabel**.
- 4 When the confirmation message appears, click **OK** to verify that you want to unlabel the selected disk(s). (Click **Cancel** to abort without unlabelling the disk.)

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

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.

Note: After expansion you must perform a metadata dump. StorNext provides an option that will do this for you automatically, but the process can take longer than if you do the metadump manually. If you do not create a new metadump, subsequent StorNext backups will not succeed and a RAS message will be generated at backup time.

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 **Migrate Data** from the **Tools** menu.

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 "read-only," 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 read-only 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. StorNext can support up to 512 stripe groups.

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.

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 [Check File System](#) on page 73.)
- 2 Expand the file system. (See [Step 4: File Systems](#) on page 29.)
- 3 Move data stripe groups or metadata/journal stripe groups. (See [Migrate Data](#) on page 80.)
- 4 Mark source stripe groups as read-only.
- 5 Reboot all clients after the expansion.

Using Resource Allocation From the Command Line

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 interface (CLI):

- [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 snfsdefrag 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 `sdiskmove` 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 `sdiskmove` 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 `sdiskmove`.

Caution: The metadata/journal stripe group you want to move cannot contain data.

`Sdiskmove` 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.

Check File System

Before you perform either File System Expansion or Migration, 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.

Also, this operation could consume a significant amount of space on the local file system. For example, for large file systems you should allow at least 20GB of free space on the local file system for temporary files.

For more information about file system expansion, refer to the StorNext online help.

There are two ways to check file systems:

- Checking while the file system is offline
- Checking while the file system is active

When the file system is offline, you can run the check in either traditional mode or read-only mode. Read-only mode typically completes faster, but is not as thorough.

When the file system is active, you must run the check in read-only mode. The advantage of this method is that you don't have to take the file system offline to run the check.

Note: Running a check on an active file system could result in false errors which occur because you are running the check while the file system is still running.

Whenever you run the check in read-only mode, Quantum strongly recommends also running the Recover Journal step before you check the file system. Running Recover Journal ensures that all operations have been committed to disk, and that the metadata state is up to date.

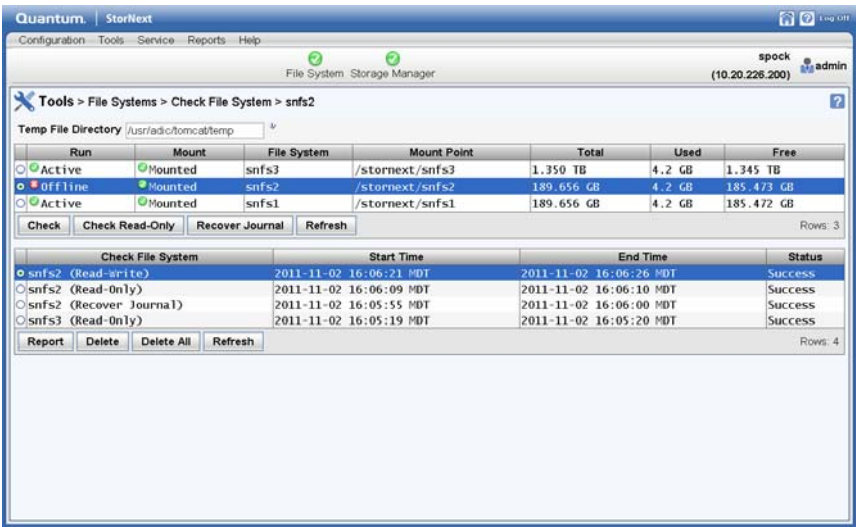
Regardless of which method you choose to check the file system, you should plan carefully when to run a file system check and plan accordingly.

Use the following procedure to perform a file system check.

Note: If you plan to run the check while the file system is offline, before you begin the following procedure you should first stop that file system as described in the StorNext online help.

- 1 Choose **Check File System** from the **Tools > File Systems** menu. The **Tools > Check > [file system name]** screen appears.

Figure 35 Check File System Screen



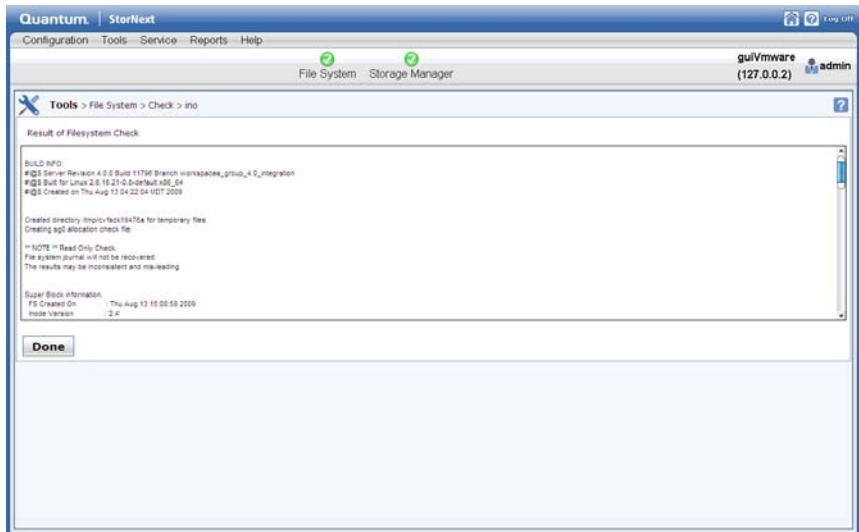
- 2 At the **Temp File Directory** field, enter a new directory if the specified directory does not have enough space to perform the check. (The checking process on large file systems can take hundreds of megabytes or more of local system disk space for working files.)
- 3 Select the file system you want to check.
- 4 If you plan to run the check in read-only mode, Quantum recommends running Recover Journal by clicking **Recover Journal**. When a message asks you to confirm that you want to run Recover Journal, click **Yes** to proceed or **No** to abort.
- 5 Do one of the following:
- a If the file system you want to check is active, click **Check Read-Only** to check the file system in read-only mode.

- b If the file system you want to check is offline, click **Check** to check the file system in “regular” mode, or **Check Read-Only** to check in read-only mode.

Viewing and Deleting a
Check Report

After you have run at least one file system check, information about the process appears at the bottom of the screen: file system name, the time the check was initiated and completed, and the status of the check. To view details about a specific check, select the desired check at the bottom of the screen and then click **Report**. When you are finished viewing the report, click **Done** to return to the previous screen.

Figure 36 Check File System
Report



To delete a check report from the list, select the check you want to delete and then click **Delete**. To delete all previously run checks listed, click **Delete All**.

File System Check Output Files

If you do not want to use StorNext to view output from the file system check, you can view output in two files:

- `/usr/cvfs/data/<fsname>/trace/cvfsck-<timestamp>`
For example: `/usr/cvfs/data/snfs1/trace/cvfsck-02_22_2010-12_15_19`
- `/usr/adic/gui/logs/jobs/CHECK_FS-<timestamp>-<jobid>`
For example: `/usr/adic/gui/logs/jobs/CHECK_FS-20100222_121519-77`

Affinities

This section describes StorNext's "stripe group affinity" feature, and also provides some common use cases.

A *stripe group* is a collection of LUNs (typically disks or arrays) across which data is striped. Each stripe group also has a number of associated attributes, including affinity and exclusivity.

An *affinity* is used to steer the allocation of a file's data onto a set of stripe groups. Affinities are referenced by their name, which may be up to eight characters long. An affinity may be assigned to a set of stripe groups, representing a named pool of space, and to a file or directory, representing the logical point in the file system and directing the storage to use the designated pool.

Exclusivity means a stripe group has both an affinity and the exclusive attribute, and can have its space allocated only by files with that affinity. Files without a matching affinity cannot allocate space from an exclusive stripe group. Files with an affinity that is exclusive cannot be stored on other stripe groups without that affinity. If the exclusive stripe group(s) become filled, no more files with that affinity can be stored.

Affinities for stripe groups are defined in the file system configuration file. A stripe group may have multiple affinities, and an affinity may be assigned to multiple stripe groups.

Note: A managed file system can have only two affinities, but a non-managed file system can have multiple affinities.

Allocation Strategy

- StorNext has multiple allocation strategies which can be set at the file system level. These strategies control where a new file's first blocks will be allocated. Affinities modify this behavior in two ways:
- A file with an affinity will be allocated only on a stripe group with matching affinity.
- A stripe group with an affinity and the exclusive attribute will be used only for allocations by files with matching affinity.

Once a file has been created, StorNext attempts to keep all of its data on the same stripe group. If there is no more space on that stripe group, data may be allocated from another stripe group. If the file has an affinity, only stripe groups with that affinity will be considered; if all stripe groups with that affinity are full, new space may not be allocated for the file, even if other stripe groups are available.

Example Use Cases

Affinities can be used to segregate audio and video files onto their own stripe groups. For example:

- Create one or more stripe groups with an AUDIO affinity and the exclusive attribute.
- Create one or more stripe groups with a VIDEO affinity and the exclusive attribute.
- Create one or more stripe groups with no affinity (for non-audio, non-video files).
- Create a directory for audio using 'cvmkdir -k AUDIO audio'.
- Create a directory for video using 'cvmkdir -k VIDEO video'.

Files created within the audio directory will reside only on the AUDIO stripe group. (If this stripe group fills, no more audio files can be created.)

Files created within the video directory will reside only on the VIDEO stripe group. (If this stripe group fills, no more video files can be created.)

To reserve high-speed disk for critical files:

- Create a stripe group with a FAST affinity and the exclusive attribute.
- Label the critical files or directories with the FAST affinity.

The disadvantage here is that the critical files are restricted to only using the fast disk. If the fast disk fills up, the files will not have space allocated on slow disks.

To reserve high-speed disk for critical files, but allow them to grow onto slow disks:

- Create a stripe group with a FAST affinity and the exclusive attribute.
- Create all of the critical files, pre allocating at least one block of space, with the FAST affinity. (Or move them using `snfsdefrag`, after ensuring they are non-empty.)
- Remove the FAST affinity from the critical files.

Because files will allocate from their existing stripe group, even if they no longer have a matching affinity, the critical files will continue to grow on the FAST stripe group. Once this stripe group is full, they can allocate space from other stripe groups, since they do not have an affinity.

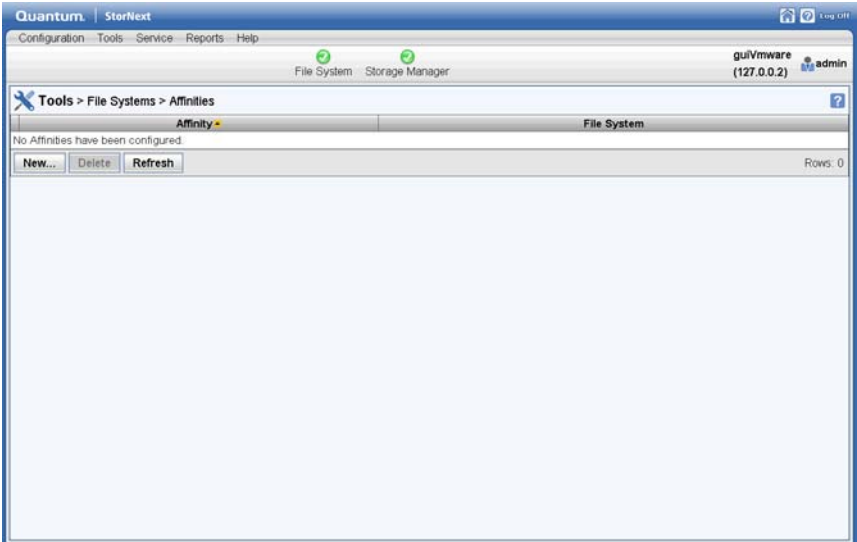
This will not work if critical files may be created later, unless there is a process to move them to the FAST stripe group, or an affinity is set on the critical files by inheritance but removed after their first allocation (to allow them to grow onto non-FAST groups).

Adding a New Affinity

Follow this procedure to add affinities:

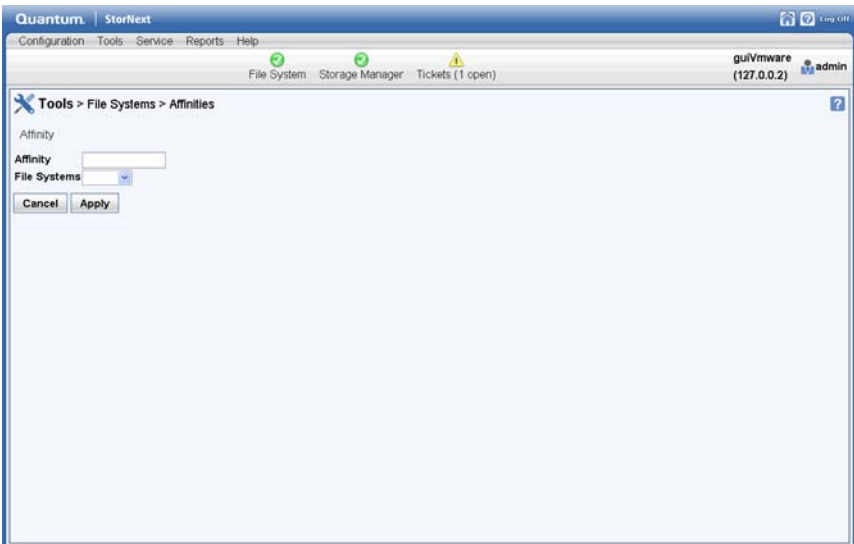
- 1 Choose **Affinities** from the **Tools > File Systems** menu. The **Tools > File Systems > Affinities** screen appears.

Figure 37 Affinities Screen



2 Click **New**. The **New Affinity** screen appears.

Figure 38 New Affinity Screen



3 At the **Affinity** field, enter the name of the new affinity.

- 4 At the **File Systems** field, select the file system to which you want to associate the new affinity.
- 5 Click **Apply** to create the affinity.
- 6 When a message notifies you that the affinity was successfully created, click **OK** to continue.

Deleting an Affinity

Follow this procedure to delete affinities:

- 1 If you have not already done so, choose **Affinities** from the **Tools > File Systems** menu. The **Tools > File Systems > Affinities** screen appears.
- 2 Select the affinity you want to delete.
- 3 Click **Delete**.
- 4 When asked to confirm the deletion, click **Yes** to proceed or **No** to abort.
- 5 When a message notifies you that the affinity was successfully deleted, click **OK** to continue.

Migrate Data

Migrating file system data refers to moving data files from a file system's source stripe group to all the other stripe groups on the same file system, and then freeing the source stripe group so it can be removed from the file system. You can select the source stripe group only, not the destination stripe group(s). Files will be moved randomly to new stripe groups while respecting their affinity rules (if any). When migrating, make sure the source stripe group is completely empty when the process completes, because source files that are updated while the file system is running may be left behind, requiring a second iteration of the migration.

During file system migration, you indicate a file system from which to move data. This operation can be performed only on data-only stripe groups. Stripe groups which contain metadata and/or journal files

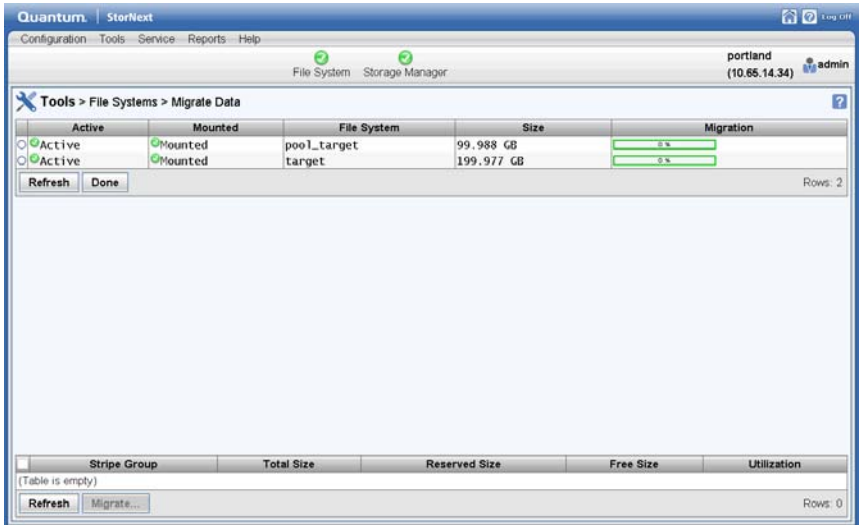
cannot be migrated by this operation. During movement the file system is left online and read/write operations occur normally.

The time it takes to complete the migration process depends on the amount of data being moved between source file system and destination stripe groups. When moving a data stripe group, the file system continues to run during the move. StorNext does not block any new read/write requests, or block updates to existing files on the source file system. 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.

Use the following procedure to perform file system migration.

- 1 Choose **Migrate Data** from the **Tools > File Systems** menu. The **Tools > File systems > Migrate** screen appears.

Figure 39 Migrate Screen



- 2 Select the file system from which files will be migrated.
- 3 Select the stripe group to which files will be migrated.
- 4 Click **Migrate**.

Caution: This particular function does not provide a confirmation message, so be absolutely sure you want to migrate data from the selected file system to the selected stripe group before you click **Migrate**.

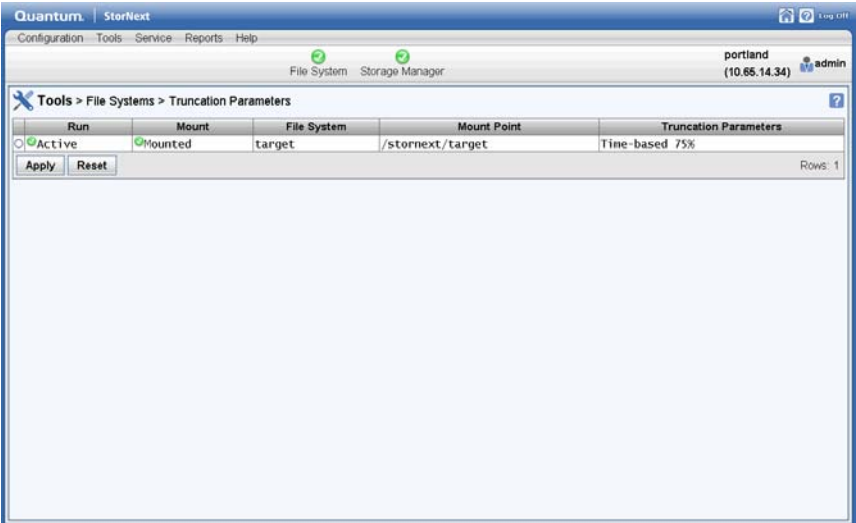
Truncation Parameters

The Truncation Parameters screen enables you to view or change the following information pertinent to the truncation feature as it pertains to StorNext Storage Manager:

- **Run:** Indicates the current status of the truncation feature: Online or Offline.
- **Mount:** Indicates whether the file system is currently mounted.
- **File System:** Displays the name of the truncation-enabled file system.
- **Mount Point:** Shows the mount point for the truncation-enabled file system
- **Truncation Parameters:** Shows the current truncation setting, such as Time-based 75%.

Note: This screen pertains **ONLY** to truncation for StorNext Storage Manager users. It does not apply to deduplication or other StorNext features.

Figure 40 Truncation
Parameters Screen



To change truncation parameters:

- 1 Click the line containing the file system whose truncation parameters you want to change. Parameters appear at the bottom of the screen.
- 2 As desired, modify any of the following fields. (See the online help for information about what to enter at each field.)
 - **Enable Truncation**
 - **Truncation Mode**
 - **Minimum Usage (%)**
 - **Low Water (%)**
 - **High Water (%)**
- 3 Click **Apply** to save your changes.
- 4 When a confirmation message appears, click **Yes** to continue or **No** to abort without saving.

Note: When you save changes to truncation parameters, the StorNext Policy Manager must be restarted. This process could take several minutes, so plan accordingly.

- 5 Click **Done** when you are finished viewing or changing truncation parameters.



Chapter 5

Storage Manager Tasks

The **Tools > Storage Manager** menu contains options that enable you to perform the following Storage Manager-related tasks:

- [Storage Components](#): View your system's libraries, storage disks, and tape drives, and place those devices online or offline
- [Drive Pools](#): View, add, edit, or delete drive pools (groups of tape drives allocated for various administrator-defined storage tasks)
- [Media Actions](#): Perform various actions on the storage media in your library
- [Storage Exclusions](#): Specify types of file names to exclude from StorNext Storage Manager
- [Truncation Exclusions](#): Specify files or directories to exclude from the truncation process
- [Tape Consolidation](#): Enter parameters for automatically consolidating space on tape media
- [Library Operator Interface](#): The StorNext Library Operator Interface allows you to perform media-related actions remotely from the library
- [Software Requests](#): View current software requests in progress or cancel a request
- [Scheduler](#): Schedule tasks to run automatically based on a specified schedule

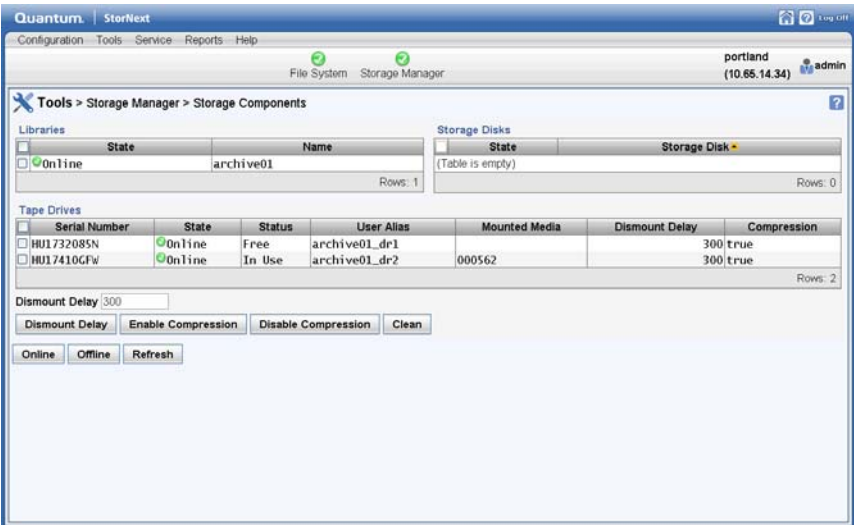
- [Alternate Retrieval Location](#): Specify a remote retrieval location to use in situations where files stored on tape or a storage disk cannot be accessed.
- [Distributed Data Mover \(DDM\)](#): Spread the distribution of data across several machines rather than the primary server.

Storage Components

The Tools menu's Storage Components option enables you to view your system's libraries, storage disks, and tape drives, and place those devices online or offline. The **Tools > Storage Manager > Storage Components** screen is divided into three sections corresponding to libraries, storage disks and tape drives.

To access the **Tools > Storage Manager > Storage Components** screen, choose **Storage Components** from the **Tools > Storage Manager** menu.

Figure 41 Storage Components Screen



Setting Devices Online and Offline

The process for setting devices online or offline is identical regardless of device type. Select the library, storage disk or tape drive you want to place online or offline. You can select multiple devices in each category, or select all available devices in each category by clicking **All**. After you are satisfied with your selections, click either **Online** to place selected devices online, or **Offline** to take selected devices offline.

Additional Options for Tape Drives

There are four additional options available for tape drives:

- **Dismount Delay:** This option enables you to specify the time, in seconds, that a tape drive remains idle before the media in that drive is dismounted. Select the tape drives for which you want the delay, enter the desired time interval at the Dismount Delay field, and then click **Dismount Delay**.
- **Enable Compression:** Compression is a feature supported by some tape drives which maximizes the amount of available storage space. To enable compression, select the tape drives for which you want to enable compression and then click **Enable Compression**.
- **Disable Compression:** If compression was previously enabled and you want to disable it, select the tape drives for which you want to disable compression and then click **Disable Compression**.
- **Clean:** This option allows you to request that a drive be cleaned. Before choosing this option, make sure the library contains a cleaning cartridge. When you are ready to proceed, click **Clean**.

Note: Although not recommended, if your library does not contain any cleaning cartridges you can disable drive cleaning. Refer to parameter `FS_CLEAN_DRIVES` in `/usr/adic/TSM/config/fs_sysparm.README` to disable drive cleaning.

Drive Pools

Drive pools are groups of tape drives allocated for various administrator-defined storage tasks, and enable you to delimit storage processes

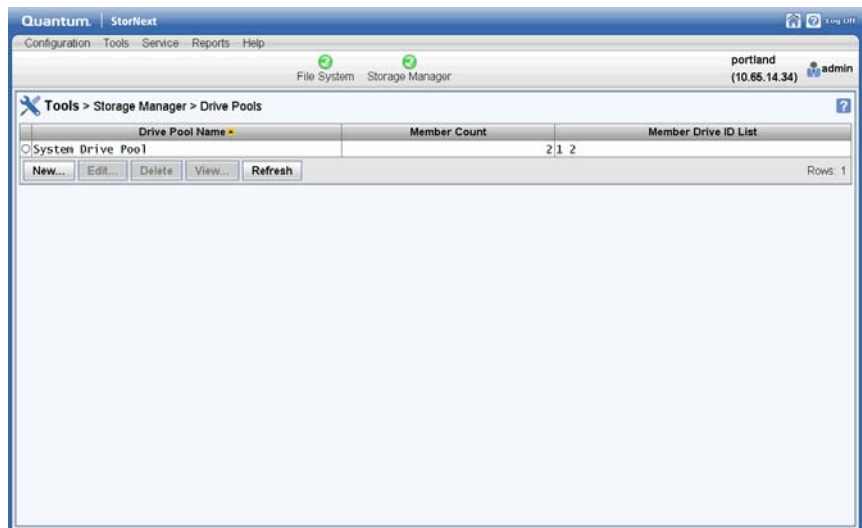
based on data type, performance, security, location, or all of these variables. Drive pools can reside in a single tape library or span multiple tape libraries.

Viewing Drive Pool Information

Follow this procedure to view drive pool information.

- 1 Choose **Storage Manager > Drive Pools** from the **Tools** menu. The **Drive Pools** screen appears.

Figure 42 Drive Pools Screen



- 2 Select the drive pool whose information you want to see, and then click **View**.
- 3 The following information appears:
 - **Serial Number:** The serial numbers of all tape drives in the drive pool
 - **Drive Alias:** The corresponding alias number for each drive
 - **Media Type:** The type of tape drive media for each drive (e.g., LTO)
 - **Library:** The name of the library to which each drive belongs
 - **Pool Name:** The name of the drive pool to which each drive belongs

- 4 When you are finished viewing drive pool information, click **Done**.

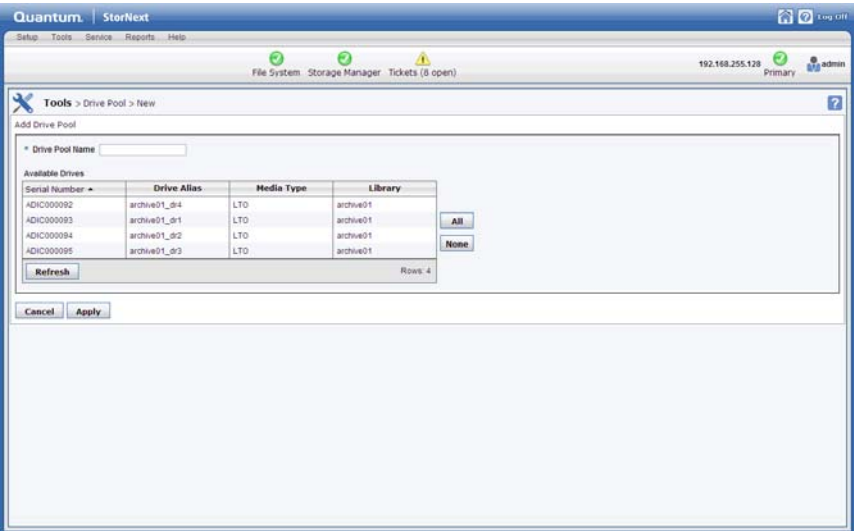
Adding a Drive Pool

Follow this procedure to add a drive pool.

Note: This procedure requires restarting the StorNext Storage Manager component.

- 1 If you have not already done so, choose **Storage Manager > Drive Pools** from the **Tools** menu.
- 2 Click **New** to add a new drive pool. The **Drive Pools > New** screen appears.

Figure 43 New Drive Pool Screen



- 3 Enter the following fields. (For information about what to enter at each field, see the online help.)
 - **Drive Pool Name**
 - **Available Drives**
- 4 Click **Apply**.

- 5 When the confirmation message appears, click **Yes** to proceed or **No** to abort. If you click **Yes**, StorNext Storage Manager will be restarted as part of the creation process.
- 6 After a message informs you that the drive pool was successfully created, click **OK** to continue.

Editing a Drive Pool

Follow this procedure to edit a drive pool.

Note: This procedure requires restarting the StorNext Storage Manager component.

- 1 If you have not already done so, choose **Storage Manager > Drive Pools** from the **Tools** menu.
- 2 Select the drive pool you want to modify, and then click **Edit**.
- 3 Select or deselect available drives for the drive pool. (You cannot change the drive pool name.)
- 4 Click **Apply**.
- 5 When the confirmation message appears, click **Yes** to proceed or **No** to abort.
- 6 After a message informs you that the drive pool was successfully modified, click **OK** to continue.

Deleting a Drive Pool

Follow this procedure to delete a drive pool. Before you begin, you must first remove all drives in the pool you want to delete.

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

- 1 If you have not already done so, choose **Storage Manager > Drive Pools** from the **Tools** menu.
- 2 Select the drive pool you want to delete, and then click **Delete**.
- 3 When a confirmation message appears, click **Yes** to proceed with the deletion or **No** to abort.

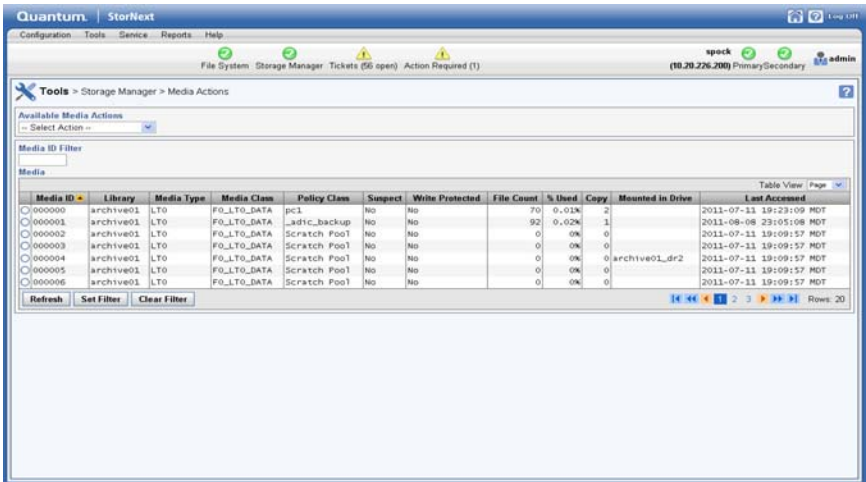
- 4 After a message informs you that the drive pool was successfully deleted, click **OK**.

Media Actions

The Tools menu's **Media Actions** option enables you to perform various actions on the storage media in your library.

To access the **Tools > Storage Manager > Media Actions** screen, choose **Media Actions** from the **Tools > Storage Manager** menu.

Figure 44 Media Actions Screen



Viewing Media Information

After you choose the **Media Actions** option, the following information about all of the storage media appears:

- **Media ID:** The unique identifier for the media.
- **Library:** The name of the library in which the media resides.
- **Media Type and Class:** The media type and class of media. (For example, LTO, F0_LTO_DATA)

- **Policy Class:** The name of the policy class, if any, associated with the media.
- **Suspect:** Indicates whether the media is “suspect” or potentially defective.
- **Write Protected:** Indicates whether write protection is enabled on the media.
- **File Count:** The current number of files currently on the media.
- **% Used:** Indicates the percentage of the media which is currently used.
- **Copy:** Indicates the policy class copy number on the media.
- **Mounted in Drive:** Indicates whether the media is mounted.
- **Last Accessed:** The time the media was last accessed.

Filtering Media

Most **Media Actions** screens contain a filtering feature that allows you to restrict the available media to those whose media ID contains the string you specify. Follow these steps to filter media:

- 1 At the **Media ID Filter** field, enter the string you want all available media IDs to include.
- 2 Click **Set Filter**.
- 3 Click Refresh to update the list of available media. Only media whose IDs contain the string you entered will be shown.
- 4 To reset the filter string, click **Clear Filter**. If desired, repeat steps 1 - 3 to use a new filter string.

Performing Media Actions

At the top of the screen is a dropdown list of actions you can perform for selected media. Select the media for which you want to perform the action, and then choose one of these options from the Available Actions list:

Mount Media

Select this option to mount the storage media.

- 1 After you select this option, select from the **Library** dropdown list the library containing the media you want to mount.
- 2 Select the media to mount.
- 3 At the **Mount Media Parameters > Drive** field, select the drive on which to mount the media.
- 4 Click **Apply**.
- 5 When the confirmation message appears, click **Yes** to mount the media, or **No** to abort.

Dismount Media

Select this option to dismount previously mounted media.

- 1 After you select this option, a list of mounted media appears.
- 2 Select the media you want to dismount, and then click **Apply**.
- 3 When the confirmation message appears, click **Yes** to dismount the media, or **No** to abort.

Move Media

Select this option to move media from one library to another. This media action will retain all information about the data contained on the media being moved.

- 1 After you select this option, select from the **Library** dropdown list the library containing the media you want to move.
- 2 Select one or more media to move, or check the box to the left of the **Media ID** heading to select all media.
- 3 At the **Move Media Parameters > Destination Library** field, select the destination library to which you want to move the selected media.
- 4 Click **Apply**.
- 5 When the confirmation message appears, click **Yes** to move the selected media, or **No** to abort.
- 6 Use the Library Operator Interface feature to complete the actual physical move of the media. (See [Library Operator Interface](#) on page 111 for more information.)

Manual Move Media

Select this option to manually move media from one library to another. This media action is typically used to move media to a new archive from a dead or offline archive.

- 1 After you select this option, select from the **Library** dropdown list the library containing the media you want to move.
- 2 Select one or more media to move, or check the box to the left of the **Media ID** heading to select all media.
- 3 At the **Move Media Parameters > Destination Library** field, select the destination library to which you want to manually move the selected media.
- 4 Click **Apply**.
- 5 Complete the process by manually moving the media you specified to the destination library.

Remove Media

Select this option to remove media from the StorNext Storage Manager. Only media with no active files on a media can be selected for removal. The media is removed from the system and is physically ejected from the library.

- 1 After you select this option, select from the **Library** dropdown list the library containing the media you want to remove.
- 2 Select one or more media to remove, or check the box to the left of the **Media ID** heading to select all media.
- 3 Click **Apply**.
- 4 When the confirmation message appears, click **Yes** to remove the selected media, or **No** to abort.

Purge Media

Select this option to purge media from the StorNext Storage Manager. All files are removed from the selected media, and then the media is removed from the StorNext Storage Manager and is physically ejected from the library.

- 1 After you select this option, select from the **Library** dropdown list the library containing the media you want to purge.
- 2 Select one or more media to purge, or check the box to the left of the **Media ID** heading to select all media.
- 3 Click **Apply**.
- 4 When the confirmation message appears, click **Yes** to purge the selected media, or **No** to abort.

Reclassify Media

Select this option to change the media type classification for selected media.

- 1 After you select this option, select from the **Media Class** dropdown list the current media class designation you want to change.
- 2 Select one or more media to reclassify, or check the box to the left of the **Media ID** heading to select all media.
- 3 At the **Reclassify Media Parameters > Destination Media Class** field, select the new media type designation for the selected media. Select one of these options:
 - **DATA**: This media class means that media are candidates for read/write operations. Most media residing in the library have this classification unless they are full.
 - **ADDBLANK**: This is the default class with which media are associated when they are added to StorNext MSM. (Running the `Fsmedin` command pulls media from this class and changes the classification to **DATA**.)
 - **IMPORT**: Before running the `fsmedin` command on TSM-exported media, the classification should be changed to **IMPORT**.
 - **CHECKIN**: This classification is used for re-entering media which have been checked out. Media must be reclassified with **CHECKIN** prior to TSM performing `fsmedin` with the `checkin` option.
 - **MIGRATE**: TSM reclassifies media to this classification when the media becomes full according to the `FS_PERCENT_FULL` system parameter. Media with this classification can still be read.

- **CLEAN:** Media in the class are cleaning media. If the barcode of a media ends with CLN, MSM imports the media into this class instead of ADDBLANK.
- **REMOVE:** Media get reclassified to REMOVE when fsmedout is used.
- **BACKUP:** Media with this classification were used for backups before backups were managed by StorNext storage policies. Consequently, this classification is rarely used.

4 Click **Apply**.

5 When the confirmation message appears, click **Yes** to reclassify the selected media, or **No** to abort.

Assign Media to Policy Class

Select this option to assign media to a previously created policy class.

- 1 Select one or more media to assign, or check the box to the left of the **Media ID** heading to assign all media.
- 2 At the **Assign Media to Policy Class Parameters > Destination Policy Class** field, select the policy class to which you want to assign selected media.
- 3 Click **Apply**.
- 4 When the confirmation message appears, click **Yes** to assign the selected media, or **No** to abort.

Transcribe Media

Transcribe (copy) the contents of one media type to another media type, or reclaim (defragment) media. During the transcription or reclamation process, StorNext uses two drives to transcribe one media to another media, file by file.

Caution: For StorNext 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 Select one or more media to transcribe, or check the box to the left of the **Media ID** heading to select all media.

- 2 Click **Apply**.
- 3 When the confirmation message appears, click **Yes** to dismount the media, or **No** to abort.
- 4 Repeat steps 1 - 3 to transcribe additional media.

If transcription or reclamation starts and all the drives are in use, StorNext prioritizes and queues the job. When two drives become available, the queued job starts and all active files are transcribed. When transcription is complete, StorNext 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. 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.

Media Attributes

Select this option to view the attributes currently assigned to your media, or to change attributes.

- 1 If desired, filter the displayed list of media by selecting one or more of the following media attribute filters: **Suspect**, **Marked**, **Full**, **Unavailable**, or **Write Protected**. The list refreshes each time you select a media attribute filter.

"Suspect" means the media might not be physically sound, and could be in a potentially damaged or unusable condition.

"Marked" means the media should be made inaccessible.

"Full" means the media has reached capacity and should not be available for further writing.

"Unavailable" means the media is not available for writing or reading.

"Write Protected" means the media is protected against further writing and cannot be overwritten or have data added.

- 2 Select from the list one or more media whose attributes you want to change, or check the box to the left of the **Media ID** heading to select all media.
- 3 At the **Media Attributes Parameters > New Media State** field, select the new attribute you want to apply to the selected media.
- 4 Click **Apply**.
- 5 When the confirmation message appears, click **Yes** to move the selected media, or **No** to abort.
- 6 Repeat steps 3 - 5 to apply additional attributes to the selected media.

Clean Media by Media ID

Select this option if you want to select media for cleaning based on media ID. 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 Select one or more media you want to clean, or check the box to the left of the **Media ID** heading to select all media.
- 2 At the **Clean Media by Media ID Parameters > End Time** field, enter the time when you want the cleaning process to stop. (You can also use the displayed default end time.)
- 3 Use the format **yyyy:MM:dd:HH:mm:ss** when entering an end time.
- 4 Click **Apply**.
- 5 When the confirmation message appears, click **Yes** to begin cleaning media, or **No** to abort.

Clean Media by File System

Select this option if you want to select media for cleaning based on the file system with which media are associated. Periodic cleaning helps prevent inactive information from growing to an unmanageable size.

When you select this option all media on the selected file system are cleaned. 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 At the **Clean Media by File System Parameters > Managed and Mounted File Systems** field, select the file system whose media you want to clean.
- 2 At the **Clean Media by File System Parameters > End Time** field, enter the time when you want the cleaning process to stop. (You can also use the displayed default end time.) Use the format **yyyy:MM:dd:HH:mm:ss** when entering an end time.
- 3 Click **Apply**.
- 4 When the confirmation message appears, click **Yes** to begin cleaning media, or **No** to abort.

Clean Media by Policy Class

Select this option if you want to select media for cleaning based on the policy class with which media are associated. When you select this option all media associated with the selected policy class are cleaned. Periodic cleaning helps prevent inactive information from growing to an unmanageable size. When you select this option all media on the selected file system are cleaned. 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 At the **Clean Media by Policy Class Parameters > Policy Classes** field, select the policy class whose media you want to clean.
- 2 At the **Clean Media by Policy Class Parameters > End Time** field, enter the time when you want the cleaning process to stop. (You can also use the displayed default end time.) Use the format **yyyy:MM:dd:HH:mm:ss** when entering an end time.

- 3 Click **Apply**.
- 4 When the confirmation message appears, click **Yes** to begin cleaning media, or **No** to abort.

Add Media Mailbox

Select this option to add media through a library mailbox.

- 1 At the **Add Media Mailbox Parameters > Library** field, select the library with the mailbox through which you want to add media.
- 2 Click **Apply**.
- 3 When the confirmation message appears, click **Yes** to proceed, or **No** to abort.
- 4 When a message informs you that the operation was successful, click **OK**. After you see this message you are ready to load media through the library mailbox.

Add Media Bulk Load

Select this option to add media to a library via bulk loading.

- 1 At the **Add Media Bulk Load Parameters > Library** field, select the library into which you want to bulk load media.
- 2 Click **Apply**.
- 3 When the confirmation message appears, click **Yes** to proceed, or **No** to abort.
- 4 When a message informs you that the operation was successful, click **OK**. After you see this message you are ready to bulk load media into the library.

Storage Exclusions

The Tools menu's **Storage Exclusions** option enables you to specify types of files you want excluded from storage and StorNext processing. For example, you may want to exclude certain confidential files or files containing sensitive material.

The process involves specifying file names, as well as criteria so StorNext knows how to identify the files to exclude. You can create as many exclusion patterns as you like, and as your needs change you can delete unwanted exclusion patterns. These exclusions are stored in a file which StorNext checks before executing store operations.


Note: If you exclude a file and then later rename that file, the renamed file will not be excluded from store operations unless:

- there is an existing exclusion that covers the renamed file
- OR
- you create a new exclusion for the renamed file.

Accessing Storage Exclusions

To access the **Tools > Storage Manager > Storage Exclusions** screen, choose **Storage Exclusions** from the **Tools > Storage Manager** menu. Any previously saved exclusions are displayed.

Figure 45 Storage Exclusions Screen

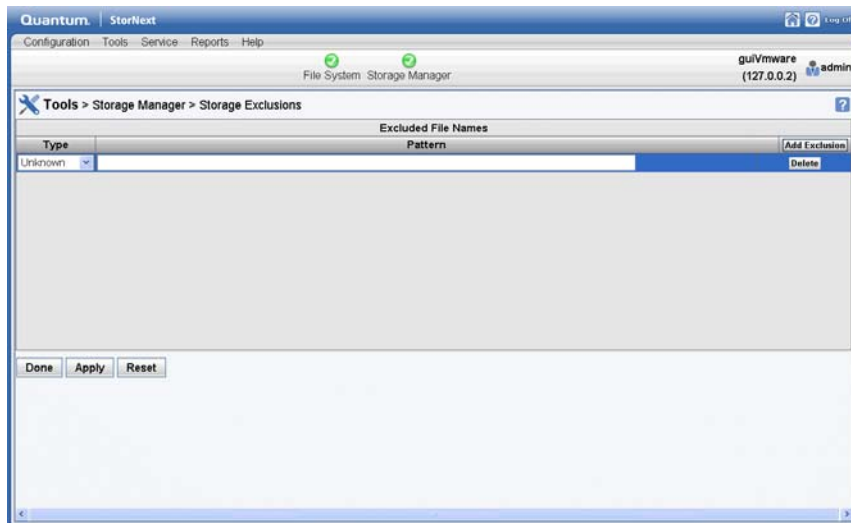


Adding an Exclusion Pattern

To add a new exclusion:

- 1 Click **Add Exclusion**. The Exclusion screen appears.

Figure 46 Exclusion Screen



- 2 Choose from the **Type** pulldown list one of the following types of exclusions:

- **Unknown**: This type appears until you select a different type, and remains if you select no type.
- **Match**: Files which match the string in the Pattern field are excluded. Wildcards can be used for this type.
- **Match Path**: Files in the path which match the string in the Pattern field are excluded. Wildcards can be used for this type.

The difference between **Match** and **Match Path** is that slashes must be explicitly entered for the Match Path type. When using Match Path you cannot substitute a wildcard for a slash. However, for storage exclusions only file names can be specified (no paths) so these two types have the same effect.

- **Begins With**: Files beginning with the string in the Pattern field are excluded.
- **Ends With**: Files ending with the string in the Pattern field are excluded.

- **Contains:** File names containing the string in the Pattern field are excluded.
- **Exact:** Only files that *exactly* match the string in the Pattern field are excluded.

- 3 At the **Pattern** field, enter the search string for the file name. Depending on the exclusion type you selected, this could be a whole or partial file name, or a string.

If you selected the Match or Match Path type, you can use the following wildcards at the Pattern field:

- ? (question mark): Substitute any single character. For example, if you enter **t?p**, it will match "top" "tip" and "tap".
- * (asterisk): Substitute one or more characters. For example, if you enter **f*I**, it will match "ful" "fail" "foil" "fall" and "frail".
- [] (brackets): When you use this wildcard, it allows a set of characters to be specified. For example, if you enter **[abc]***, any string beginning with the letters "a" "b" or "c" will be matched. When using brackets you can also specify a range by using the - (dash) character. For example, if you enter **file[1-4]**, the strings "file1" "file2" "file3" and "file4" are matched. You can also specify a complement of characters to *not* match. For example, if you enter **[!abc]***, any string that does not begin with the letters "a", "b" or "c" will be matched.

Examples

To give you an example how exclusions work with wildcards, following are some exclusion examples, including the Type selected and the string entered in the **Pattern** field to create the exclusion:

- To exclude files that have "confidential" in their name:
Type=**Contains** or **Exact**; Pattern=**confidential**
 - To exclude files beginning with the letter x or y: Type=**Match**;
Pattern=**[xy]***
- 4 Once saved, exclusions are applied automatically, and all saved exclusions apply at the same time. To save and apply the exclusion, click **Apply**. When asked to confirm updating the exclusions list, click **Yes**.
- 5 To delete an exclusion, click the **Delete** button to the right of the exclusion you want to delete.

Note: This particular delete function does not have a confirmation message, so be sure you want to delete an exclusion before you click the **Delete** button.

- 6 When you are finished, click **Done** to return to the StorNext home page.

Truncation Exclusions

The Tools menu's **Truncation Exclusions** option enables you to specify files you want to exclude from the truncation process. The file path must be included when specifying the exclusion criteria.

Since paths are included in the criteria, this allows you to specify criteria which will affect all files within a directory. This basically allows an exclusion to be specified for a directory.

For example, you may want to exclude directories containing system files or files used during system login. When you create an exclusion for a directory, none of the files in that directory are truncated.

The process involves specifying directory paths as part of the criteria so StorNext knows how to locate and identify the files to exclude. You can create as many exclusion patterns as you like, and as your needs change you can delete unwanted exclusion patterns. These exclusions are stored in a file which StorNext checks when storing but *before truncating*.

Note: If you exclude a file and then later rename that file after it has been stored, the renamed file will continue to be excluded from truncation unless:

- the renamed file does not match exclusion criteria and is modified so that it gets stored again


OR

- you use the `fschfiat -t c` command to remove the exclusion from the file
-

Accessing Truncation Exclusions

To access the **Tools > Storage Manager > Truncation Exclusions** screen, choose **Truncation Exclusions** from the **Tools > Storage Manager** menu. Any previously saved exclusions are displayed.

Figure 47 Truncation Exclusions Screen

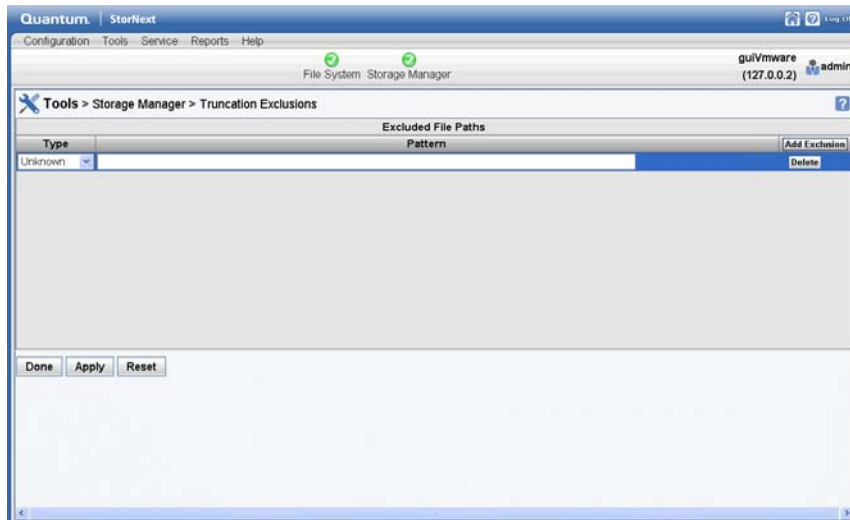


Adding an Exclusion Pattern

To add a new exclusion:

- 1 Click **Add Exclusion**. The Exclusion screen appears.

Figure 48 Exclusion Screen



- 2 Choose from the **Type** pulldown list one of the following types of exclusions:
 - **Unknown:** This type appears until you select a different type, and remains if you select no type.
 - **Match:** File paths which match the string in the Pattern field are excluded. Wildcards can be used for this type. To enter an exclusion which includes all files in a directory, the directory name and a wildcard must be specified. For example, enter `/sn/foodir/*` to exclude all files in the directory `/sn/foodir`.
 - **Match Path:** File paths which match the string in the Pattern field are excluded. Wildcards can be used for this type.

The difference between **Match** and **Match Path** is that slashes must be explicitly entered for the Match Path type. When using Match Path you cannot substitute a wildcard for a slash.
 - **Begins With:** File paths beginning with the string in the Pattern field are excluded.
 - **Ends With:** File paths ending with the string in the Pattern field are excluded.
 - **Contains:** File paths containing the string in the Pattern field are excluded.

- **Exact:** Only file paths that *exactly* match the string in the Pattern field are excluded.

- 3 At the **Pattern** field, enter the search string for the file path. Depending on the exclusion type you selected, this could be a whole or partial path name for a file.

If you selected the Match or Match Path type, you can use the following wildcards at the Pattern field:

- **?** (question mark): Substitute any single character. For example, if you enter **t?p**, it will match "top" "tip" and "tap".
- ***** (asterisk): Substitute one or more characters. For example, if you enter **f*I**, it will match "ful" "fail" "foil" "fall" and "frail".
- **[]** (brackets): When you use this wildcard, it allows a set of characters to be specified. For example, if you enter **[abc]***, any string beginning with the letters "a" "b" or "c" will be matched. When using brackets you can also specify a range by using the - (dash) character. For example, if you enter **file[1-4]**, the strings "file1" "file2" "file3" and "file4" are matched. You can also specify a complement of characters to *not* match. For example, if you enter **[!abc]***, any string that does not begin with the letters "a", "b" or "c" will be matched.

Examples

To give you an example how exclusions work with wildcards, following are some exclusion examples, including the Type selected and the string entered in the **Pattern** field to create the exclusion:

- Exclude files in directories that have "confidential" in their name: Type=**Contains** or **Exact**; Pattern=***confidential*/***
 - Exclude files in directories beginning with the letter x or y: Type=**Match**; Pattern=***[xy]*/***
 - Exclude the files in a directory named "/private/confidential": Type=**Match Path**; Pattern=**/private/confidential/***
- 4 Once saved, exclusions are applied automatically, and all saved exclusions apply at the same time. To save and apply the exclusion, click **Apply**. When asked to confirm updating the exclusions list, click **Yes**.
 - 5 To delete an exclusion, click the **Delete** button to the right of the exclusion you want to delete.

Note: This particular delete function does not have a confirmation message, so be sure you want to delete an exclusion before you click the **Delete** button.

- 6 When you are finished, click **Done** to return to the StorNext home page.

Tape Consolidation

The Tape Consolidation feature provides a way to automatically consolidate tape volumes which contain unused space that is no longer tracked in the Storage Manager database.

Releases prior to StorNext 4.1 permitted you to consolidate tape space only by manually running the `fsdefrag` (defragment tape media) command, but this functionality can now be scheduled to run automatically at specified times.

The Tape Consolidation process consists of three steps:

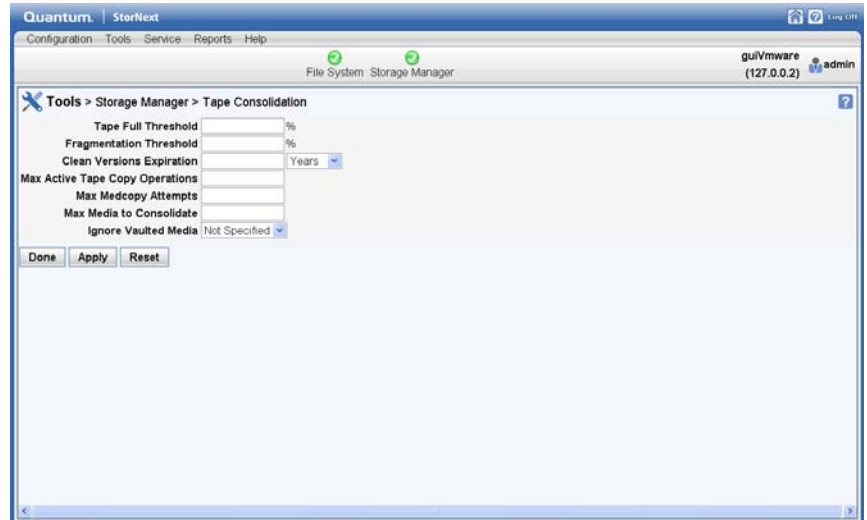
- 1 Setting configuration parameters by specifying criteria for determining which tapes are fragmented
- 2 Creating a schedule to clean old inactive versions of files on tapes
- 3 Creating a schedule for tape defragmentation

Setting Tape Consolidation Parameters

The first step in the tape consolidation process is to configure the feature so StorNext knows which tapes in your system qualify for consolidation.

- 1 Choose **Tape Consolidation** from the **Tools > Storage Manager** menu. The Tape Consolidation screen appears.

Figure 49 Tape Consolidation
Screen



2 Enter the following fields:

- **Tape Full Threshold:** Specify the percentage at which tapes become candidates for consolidation. For example, enter **85** if you want tapes flagged for consolidation when they are 85% full.
- **Fragmentation Threshold:** Specify the percentage of fragmentation at which tapes become candidates for consolidation. For example, enter **15** if you want tapes flagged for consolidation when 15% of the tape becomes fragmented.

Note: This percentage indicates the amount of data written to the tape, not overall tape capacity. For example, suppose a tape has been written to the halfway point. Of that amount, only half the data is still tracked by Storage Manager. Therefore that tape has a fragmentation percentage of 50%, not 25%.

- **Clean Versions Expiration:** Specify the number of **Days**, **Weeks**, **Months** or **Years** after which you want to clean up versions for deleted or modified files.
- **Max Active Tape Copy Operations:** Specify the number of allowable concurrent active tape copy operations. Fragmented

media are defragmented by copying the media to new media. Therefore, each copy operation uses two tape drives.

- **Max Medcopy Attempts:** Specify the maximum number of attempts before StorNext stops trying to copy a media experiencing copy failures.
- **Max Media to Consolidate:** Specify the maximum number of tape media which will be consolidated during one fsdefrag process.
- **Ignore Vaulted Media:** Enter **True** to ignore tape media in the vault, or **False** to include media in the vault.

Note: To be candidates for defragmentation, media must pass both the Tape Full Threshold AND Fragmentation Threshold percentages. If a media passes only one or the other threshold it will be ignored for consolidation.

- 3 Click **Apply** to save and apply the parameters you just entered.
- 4 When asked to confirm, click **Yes** to proceed or **No** to abort and return to the **Tape Consolidation** screen.
- 5 If you clicked Yes, a message informs you that the Tape Consolidation configuration was updated. Click **OK** to continue.
- 6 When you are finished configuring Tape Consolidation, click **Done** to return to the StorNext home page.

Scheduling Tape Cleaning and Defragmentation

The next steps in the Tape Consolidation process are to schedule version cleaning and defragmentation. The process for scheduling these operations is identical to scheduling any other process.

The defragmentation schedule item should normally be scheduled a few hours after the versions cleaning item. The two schedule items work together in managing out-of-date media contents. The clean versions item cleans up the database information for old inactive file segments, and the defragmentation item is used to replace media which have become fragmented due to these segments being cleaned.

Note: There is no default schedule for defragmentation, and the feature is off unless manually scheduled.

For more information about scheduling, see [Scheduler](#) on page 114.

Library Operator Interface

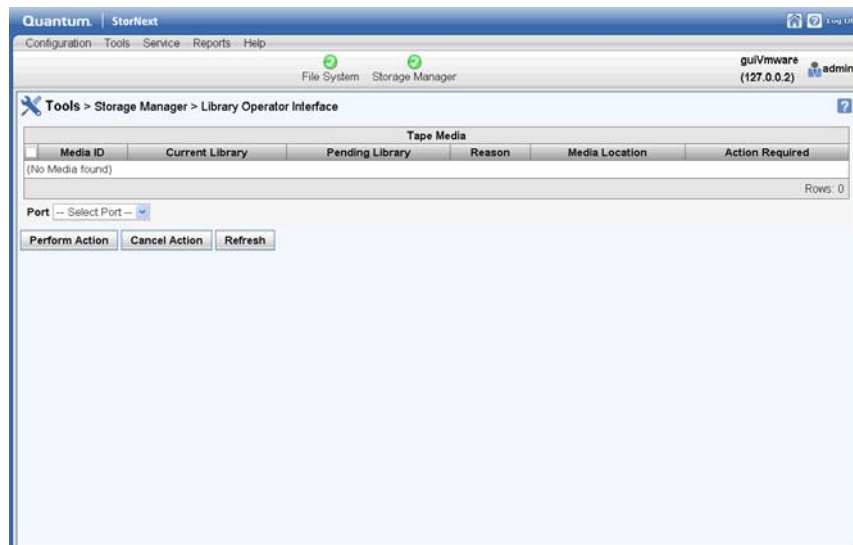
The StorNext Library Operator Interface allows you to perform media-related actions remotely from the library.

On the **Tools > Storage Manager > Library Operator Interface** screen you can view the following information about media in the library:

- **Media ID:** The unique identifier for each piece of media
- **Current Library:** The name of the library where media currently reside
- **Pending Library:** The name of the destination library to which the media action will be carried out
- **Reason:** The reason for performing the media action
- **Media Location:** The current physical location of the media
- **Action Required:** The action to be performed on selected media

- 1 Choose **Library Operator Interface** from the **Tools > Storage Manager** menu. The **Library Operator Interface** Screen appears.

Figure 50 Library Operator Interface Screen



- 2 Select one or more media on which to perform the action indicated in the **Action Required** column, or click **All** to select all media.
- 3 Select from the **Port** dropdown list the mailbox port ID for the library on which the action will be performed.
 - If the action is to enter media into a SCSI-attached library, open the mailbox and enter the media at this time.
- 4 Click **Perform Action** to initiate the action shown in the **Action Required** column, or click **Cancel Action** to cancel the action on the media.
 - If the action is to enter media into an ACSLS attached library, open the cap and enter the media at this time.
 - If the action is to eject media from an ACSLS or SCSI-attached library, open the cap/mailbox and remove the media at this time.
- 5 When the confirmation message appears, click **Yes** to proceed or **No** to abort.

Note: If you do not agree that a required action is necessary, you can select the line containing the media action and click **Cancel Action**.

Software Requests

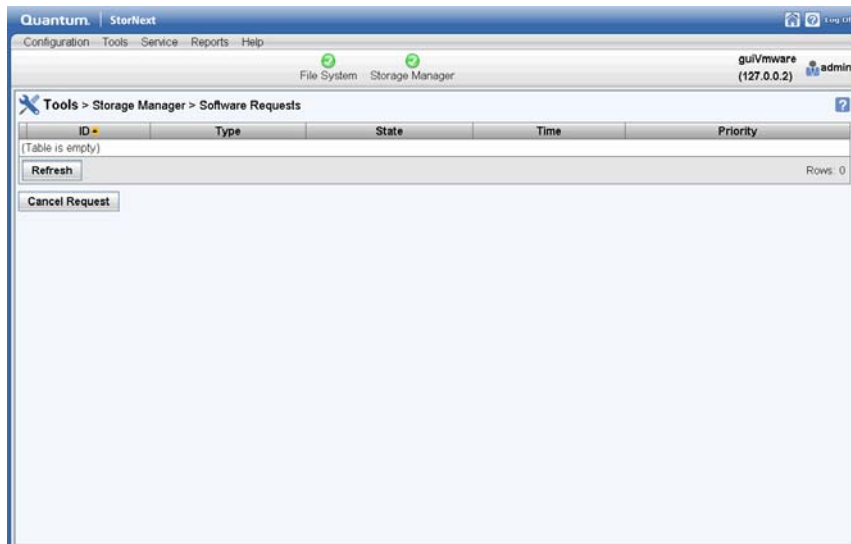
The Software Requests menu option enables you to view software requests currently in process, and to cancel requests.

On the **Tools > Software Requests** screen you can view the following information about pending and currently running software requests:

- **ID:** The software request's identifier
- **Type:** The type of software request currently running
- **State:** The current state of the request
- **Time:** The time the software request was initiated
- **Priority:** The priority assigned to the software request

- 1 Choose **Software Requests** from the **Tools > Storage Manager** menu. The **Software Requests** Screen appears.

Figure 51 Software Requests
Screen



- 2 If desired, click Refresh to update the list of software requests.
- 3 To cancel a software request, select the request you want to cancel and then click **Cancel Request**.
- 4 When the confirmation message appears, click **Yes** to proceed or **No** to abort.

Scheduler

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

- **Clean Versions:** This scheduled event cleans old inactive versions of files.
- **Clean Info:** This scheduled background operation removes from StorNext knowledge of media.
- **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.

- **Partial Backup:** By default, a partial backup is run on all days of the week the full backup is not run. Partial backups include database journals, configuration files, and file system journal 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.
- **Tape Defragmentation:** This scheduled event defragments tapes to consolidate and free up space. You should schedule the clean versions event before the defragmentation event. Only tapes that meet the parameters entered on the **Tools > Storage Manager > Tape Consolidation** screen are included in the defragmentation process. (For more information, see [Tape Consolidation](#) on page 108.)

Each of these events (with the exception of Tape Consolidation) has a default schedule, but you can configure the schedules to suit your system needs. To change the schedule or add Tape Consolidation, see [Adding a Schedule](#) or [Editing an Existing Schedule](#).

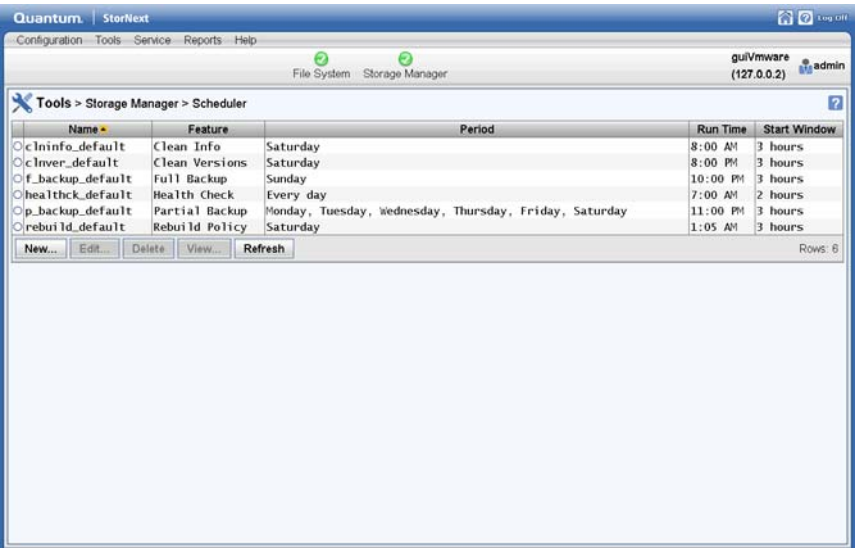
Note: To ensure that scheduled StorNext tasks are started at the correct time, StorNext servers should be rebooted whenever changes are made to the system time.

Viewing a Schedule

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

- 1 Choose **Scheduler** from the **Tools > Storage Manager** menu. A list of currently scheduled events appears.

Figure 52 Scheduler Screen



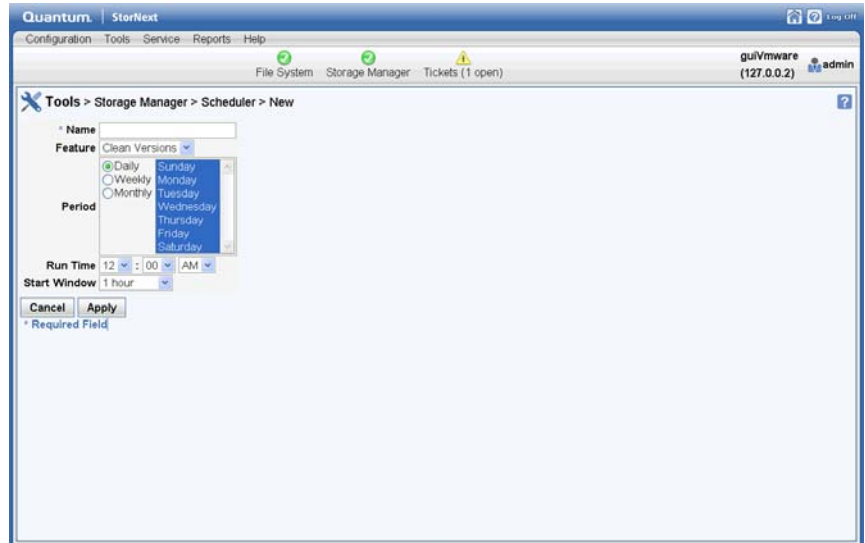
- 2 Select the event you want to view, and then click **View**.
- 3 When you are finished viewing event details, click **Done**.
- 4 Repeat steps 2 - 3 to view additional events.

Adding a Schedule

Follow this procedure to add a new schedule.

- 1 If you have not already done so, choose **Scheduler** from the **Tools > Storage Manager** menu.
- 2 Click **New**. The **Scheduler > New** screen appears.

Figure 53 Scheduler > New
Screen



- 3 At the **Name** field, enter the name you want to assign to the new schedule.
- 4 Select one of the following schedulable event types from the Feature dropdown list:
 - **Clean Versions**
 - **Clean Info**
 - **Rebuild Policy**
 - **Partial Backup**
 - **Full Backup**
 - **Health Check**
 - **Tape Defragmentation**
- 5 At the **Period** field, select the execution interval for the new schedule: **Daily**, **Weekly** or **Monthly**. You can also select multiple days by holding down the **Control** key as you click the day.
- 6 At the **Run Time** field, specify when you want the schedule to start. Enter the **hour**, **minute**, and **a.m.** or **p.m.**
- 7 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.

- 8 Click **Apply** to save the new schedule, or **Cancel** to exit without saving.
- 9 When a message informs you that the new schedule was successfully created, click **OK** to continue.

Editing an Existing Schedule

Follow this procedure to edit an existing schedule. The procedure for modifying an existing schedule is the same regardless of the event type.

- 1 If you have not already done so, choose **Scheduler** from the **Tools > Storage Manager** menu.
- 2 Select the schedule you want to modify, and then click **Edit**.
- 3 Change the schedule **Period** interval, **Run Time**, or **Start Window** as desired. You cannot change the schedule name or select a different feature (schedule type).
- 4 Click **Apply** to save your changes, or **Cancel** to exit without saving.
- 5 When the confirmation message appears, click **Yes** to proceed or **No** to abort.
- 6 When a message informs you that the new schedule was successfully modified, click **OK** to continue.

Deleting an Existing Schedule

Follow this procedure to delete an existing schedule. The procedure for deleting an existing schedule for an event is the same regardless of the event type. Each event type has a default schedule. You must have at least one schedule, so you will not be allowed to delete a solitary schedule.

- 1 If you have not already done so, choose **Scheduler** from the **Tools > Storage Manager** menu.
- 2 Select the schedule you want to delete, and then click **Delete**.
- 3 When the confirmation message appears, click **Yes** to proceed or **No** to abort.
- 4 When a message informs you that the new schedule was successfully deleted, click **OK** to continue

Alternate Retrieval Location

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. (Both machines must be using the same operating system.)

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. Because the file already exists in the StorNext file system, it retains the permissions it already has. No permssions are changed based on the file on the alternate machine.

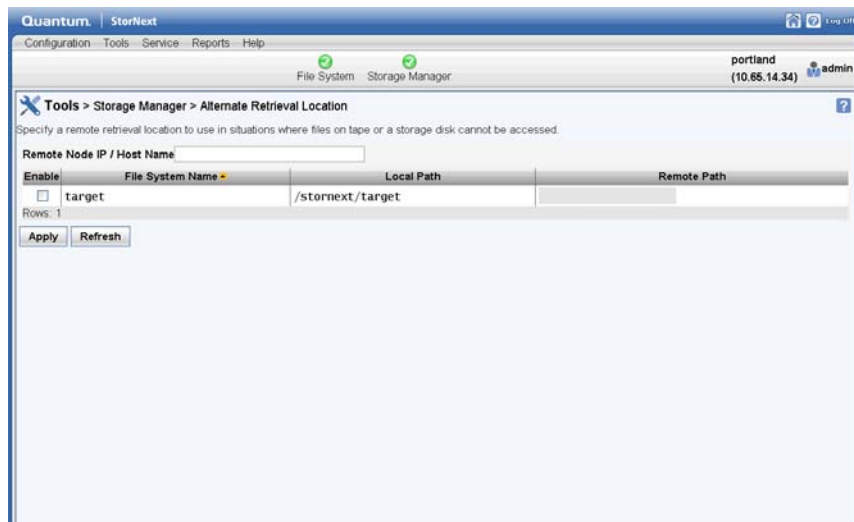
This feature applies only to managed file systems that have at least one configured policy class.

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.

Follow this procedure to configure an alternate retrieval location.

- 1 Choose **Alternate Retrieval Location** from the **Tools > Storage Manager** menu. The **Alternate Retrieval Location** screen appears.

Figure 54 Alternate Retrieval
Location Screen



- 2 At the **Remote Node IP / Host Name** field, enter either the IP address or the host name of the remote server from which you would like to retrieve data.
- 3 Select **Enable** to activate the Alternate Retrieval Location feature.
- 4 At the field under the **Remote Path** heading, enter the directory path for the remote node (server). This directory is the path that corresponds to the mount point. The remainder of the file path from the mount point downwards must be identical on the alternate host. In other words, you don't assign a relocation policy to a directory; it is enabled for the entire file system.
- 5 Click **Apply** to save your changes, or **Cancel** to exit without saving.
- 6 When the confirmation message appears, click **Yes** to proceed or **No** to abort.
- 7 After a message informs you that the alternate retrieval location was successfully added, click **OK**.

Distributed Data Mover (DDM)

StorNext contains support for a feature called Storage Manager Distributed Data Mover (DDM).

This section contains the following main topics related to DDM:

- [Distributed Data Mover Overview](#)
- [Installing the DDM Feature on Clients](#)
- [Accessing Distributed Data Mover](#)
- [Enabling DDM](#)
- [Managing DDM Hosts](#)
- [Host Priority](#)
- [Distributed Data Mover Reporting](#)

Distributed Data Mover Overview

Quantum developed the Distributed Data Mover feature to enhance the data movement scalability of its StorNext Storage Manager software. With this feature the data movement operations are distributed to client machines from the metadata controller, which can improve the overall throughput of data movement to archive tiers of storage.


Previously, the data mover process, `fs_fmover`, ran only on the metadata controller (MDC), allowing up to one `fs_fmover` process per tape drive or storage disk (SDisk) stream to run at one time.

Note: The DDM feature supports only storage disks on StorNext file systems, not on NFS.

The new feature expands data moving flexibility by transferring the mover process to clients that have access to the drives and managed file systems. The actual data moving process remains the same, with the added benefit that the load on the metadata controller is alleviated by moving those processes to clients.

The following diagram illustrates the data moving process when the Distributed Data Mover feature is enabled:

Distributed Data Mover Enabled:



Legend:

- `fs_fcopyman`: Manages copy requests and invokes a mover proxy when copy resources have become available
- `fs_fmover`: The process that performs copy operations, either on the MDC or a client
- `fs_fmoverp`: The proxy for a mover that always runs on the MDC. This proxy kicks off and waits for an `fs_fmover` process, and then makes any needed database updates etc. when the mover completes.

Note: Proxies are used even in local-only operations.

Feature Highlights

The Distributed Data Mover feature provides the following benefits:

- Concurrent utilization of shared StorNext tape and disk tier resources by multiple Distributed Data Mover hosts
- Scalable data streaming
- Flexible, centralized configuration of data movement
- Dynamic Distributed Data Mover host reconfiguration
- Support for StorNext File System storage disks (SDisks)
- Works on HA systems without additional configuration

Distributed Data Mover Terms

Following are definitions for some terms as they pertain to the Distributed Data Mover feature:

- **Mover:** A process that copies data from one device/file to another. The process can run locally on the metadata controller or on a remote client. (See definitions for these terms below.)
- **Host:** Any server/client on the SAN. Any host can serve as a location for a mover to run as long as it meets the specifications listed in the Supported Operating Systems section below.
- **Metadata Controller (MDC):** The server on which the StorNext Storage Manager software is running. (The metadata controller host.) Also known as the local host, or the primary server on HA systems.
- **Remote Client:** A host other than the MDC.

Tape Devices and Persistent SCSI Reserve

The Distributed Data Mover feature uses persistent SCSI reserve, so all tape devices used with this feature must support the PERSISTENT RESERVE IN/OUT functionality as described in SCSI Primary Commands-3 standard (SPC-3). One implication is that LTO-1 drives cannot be used with the DDM feature.

SCSI-3 persistent reservation commands attempt to prevent unintended access to tape drives that are connected by using a shared-access technology such as Fibre Channel. Access to a tape drive is granted based on the host system that reserved the device. SCSI-3 persistent reservation enables access for multiple nodes to a tape device and simultaneously blocks access for other nodes. (StorNext support for SCSI-3 persistent reservations for tape devices was incorporated in the 4.1.1 release.)

Quantum recommends enabling SCSI-3 persistent reservations. Refer to parameter `FS_SCSI_RESERVE` in `/usr/adic/TSM/config/fs_sysparm.README` to direct the StorNext Manager to use SCSI-3 persistent reservations.

One implication of using SCSI-3 reservations is that all tape devices used must support the PERSISTENT RESERVE IN/OUT functionality as described in SCSI Primary Commands-3 standard (SPC-3).

The StorNext Distributed Data Mover feature requires that SCSI-3 persistent reservations be used.

Verifying SCSI 3 Tape Drive Compatibility

A third-party utility is available to help you determine whether your tape devices are or are not compatible with SCSI-3 persistent reservations. This utility is called `sg3_utils`, and is available for download from many sites. This package contains low level utilities for devices that use a SCSI command set. The package targets the Linux SCSI subsystem.

You must download and install the `sg3_utils` package before running the following commands. In the following example, there are two SAN-attached Linux systems (`sfx13` and `sfx14` in this example) zoned to see a tape drive.

- **Step 1.** Register the reservation keys by running these commands:

```
[root@sfx13]# sg_persist -n -d /dev/sg81 -o -I -S 0x123456
```



```
[root@sfx14]# sg_persist -n -d /dev/sg78 -o -I -S 0xabcdef
```
- **Step 2:** List the reservation key by running this command:

```
[root@sfx13]# sg_persist -n -k /dev/sg81
```
- **Step 3.** Create reservation by running this command:

```
[root@sfx13]# sg_persist -n -d /dev/sg81 -o -R -T 3 -K 0x123456
```
- **Step 4.** Read reservation by running this command:

```
[root@sfx14]# sg_persist -n -d /dev/sg78 -r
```
- **Step 5.** Preempt reservation by running this command:

```
[root@sfx14]# sg_persist -n -d /dev/sg78 -o -P -T 3 -S 0x123456 -K 0xabcdef
```
- **Step 6.** Release reservation by running this command:

```
[root@sfx14]# sg_persist -n -d /dev/sg78 -o -L -T 3 -K 0xabcdef
```

- **Step 7.** Delete key by running these commands:

```
[root@sfx13]# sg_persist -n -d /dev/sg81 -o -o -K  
0x123456
```

```
[root@sfx14]# sg_persist -n -d /dev/sg78 -o -o -K  
0xabcdef
```

Limitations

Quantum does not currently support using multiple paths to tape drives. Also, VTL does not support SCSI-3 persistent reservations.

Installing the DDM Feature on Clients

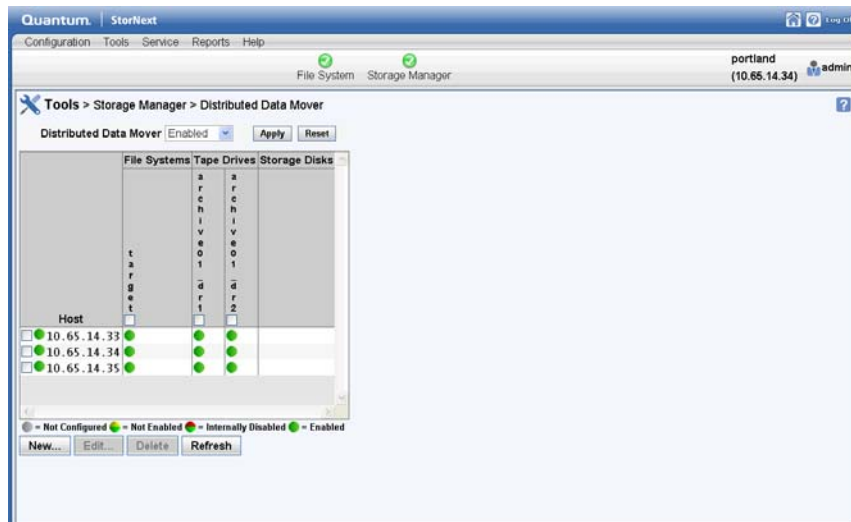
You must install the snfs-mover rpm on each client you want to act as a distributed data mover. Follow these installation steps for each client:

- 1 Log in as root.
- 2 Install the appropriate client with DDM package from the MDC.
- 3 Install the rpms in the .tar archive.
 - For a new client installation, run either the command `rpm -ivh *.rpm` or `rpm -Uvh *.rpm`
 - For a client upgrade, run the command `rpm -Uvh *.rpm`

Accessing Distributed Data Mover

To enter DDM settings and manage hosts and devices, choose **Distributed Data Mover** from the **Tools > Storage Manager** menu. The **Configuration > Distributed Data Movers** screen appears. This screen shows any previously configured DDM-enabled hosts, managed file systems, tape drives and storage disks, as well as the current status: **Enabled**, **Not Configured**, **Not Enabled** or **Internally Disabled**.

Figure 55 Configuration > Distributed Data Mover Screen



"Configured" versus "Enabled"

“Configured” means a host or device has been added to the list of hosts and devices to be used for DDM operations. DDM does not recognize a host or device until it has been configured.

“Enabled” means a host or device has been configured and is ready to be used for DDM operations. A host or device cannot be enabled until it is first configured, but a configured host or device may be either enabled or disabled.

Enabling DDM

The DDM screen's **Use Distributed Movers** field allows you to enable or disable the DDM feature, or to enable a mode called "Threshold."

When DDM is disabled, data moving responsibilities rest solely on the primary server. When DDM is enabled, data moving responsibilities are distributed among the hosts you specify as described in [Managing DDM Hosts](#) on page 127.

When you select the Threshold option, the local host (i.e., the metadata controller) is given a preference over the remote clients. The characteristics of this option are:

- Mover processes will not be assigned to a remote client until a threshold of local movers are already running.
- After reaching the threshold of local running movers, the “all” option is used for allocating new mover requests.
- If not specified, the default value for the threshold is zero. This means if a value is not set for the threshold via fsddmconfig the system will effectively run in “all” mode.

You should use the Threshold option only if you want most data moving operations to run locally on the MDC.

After you choose **Disabled**, **Enabled**, or **Threshold**, click **Update** to save your selection.

Managing DDM Hosts

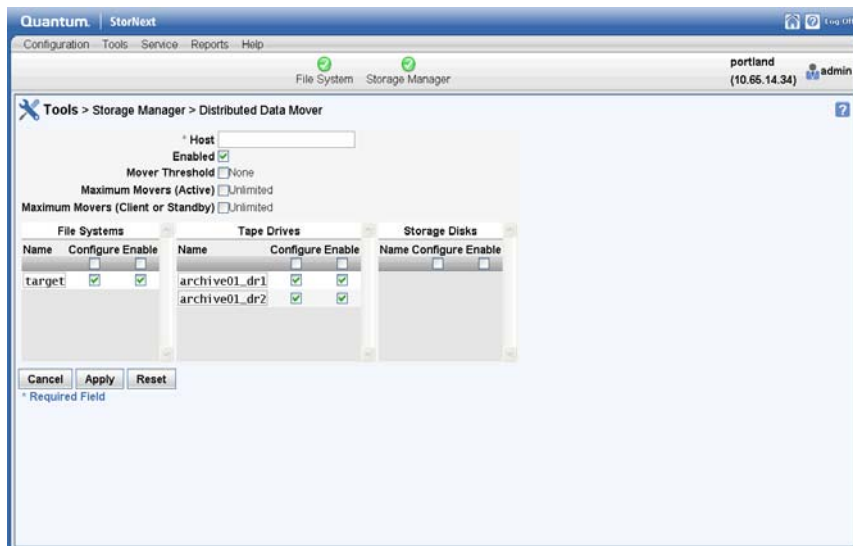
The **Distributed Data Mover** screen enables you to add and configure a new host for DDM, or change settings for a previously configured host. You can also delete a host from the list of DDM-enabled hosts.

Adding a Host

- 1 If you have not already done so, choose **Distributed Data Mover** from the **Tools > Storage Manager** menu.

- 2 Click **New**. Fields appear where you can enter host information

Figure 56 DDM Screen New Host



- 3 At the **Host** field, enter the name or IP address of the host you are adding.

Note: If you use the IP address instead of the host name you must *continue* using the IP address and cannot switch between using IP addresses and host names. For example, you can't configure a host using the IP address and then try and configure a device on the host using the host name.

- 4 Enter the remaining fields in the upper portion section of the screen:
 - **Enabled**
 - **MDC Mover Threshold**
 - **Max Movers (active MDC)**
 - **Max Movers (client or standby MDC)**

For information about what to enter at these fields, see the online help.

- 5 Under the corresponding headings, select **Configure** and/or **Enable** for the **Managed File Systems**, **Tape Drives** and **Storage Disks** you want to include in DDM processing.
- 6 To add your selections to the new host, click **Apply**. (To exit without saving, click **Cancel**. To remain on the screen but clear your entries, click **Reset**.)
- 7 When the confirmation message appears, click **Yes** to proceed or **No** to abort.
- 8 After a message informs you that your changes were successfully saved, click **OK** to continue.

Editing a Host or Devices

- 1 If you have not already done so, choose **DDM** from the **Tools > Storage Manager** menu.
- 2 Select from the **Hosts** list the host you want to edit.
- 3 Click **Edit**.
- 4 Modify the host configuration as desired.
- 5 If desired, select or remove managed file systems, tape drives and storage disks.
- 6 Click **Apply** to save your changes.
- 7 When the confirmation message appears, click **Yes** to proceed or **No** to abort.
- 8 After a message informs you that your changes were successfully saved, click **OK** to continue.

Deleting a Host

- 1 If you have not already done so, choose **DDM** from the **Tools > Storage Manager** menu.
- 2 Select from the **Hosts** list the host you want to delete.
- 3 Click **Delete** to exclude the host from DDM operation. (Clicking **Delete** does not actually delete the host from your system, but rather excludes it from the list of DDM hosts.)
- 4 When the confirmation message appears, click **Yes** to proceed or **No** to abort.

- 5 After a message informs you that the host was successfully deleted, click **OK** to continue.

Host Priority


When all hosts are chosen, no special preference is given to the local host. Following are the operating characteristics when all hosts are chosen:

- Internally within the StorNext Copy Manager (fs_fcopyman) there will be a list of hosts to use (that is, the local host and the remote clients). At this time there is no way to specify the order in which hosts appear in the list.
- For each host in the list, the following information is tracked:
 - The number of movers currently running on the host
 - The time of the last assignment
 - Whether the host is currently disabled

Note: If a host has fewer drives and all other hosts have multiple drives (for example, two drives versus ten,) the host with the fewer drives will almost always be chosen for operations on those two drives because it is likely to have the fewest running movers.

Distributed Data Mover Reporting

A DDM Report which shows current configuration information and activity is available from the **Reports** menu. For more information about the DDM report, see [The Distributed Data Mover Report](#) on page 249.



Chapter 6

Replication and Deduplication

StorNext incorporates replication and deduplication technologies which can dramatically improve storage efficiency and streamline processing.

This chapter provides the following topics pertaining to these two technologies:

- [Replication Overview](#)
- [Replication Terms and Concepts](#)
- [Some Replication Scenarios](#)
- [Configuring Replication](#)
- [Running Replication Manually \(Optional\)](#)
- [Replication Statuses and Reporting](#)
- [Troubleshooting Replication](#)
- [Data Deduplication Overview](#)
- [Setting Up Deduplication](#)
- [Data Deduplication Functions](#)

Replication Overview

This section provides some background information that will help you understand how replication is configured and how processing occurs.

Replication Configuration Overview

StorNext Replication makes a copy of a source directory and sends the information to one or more target directories. The target directories may be on other host machines, or may be on the same host as the source directory.

Replication behavior is defined by a *Replication/Deduplication Policy*. (The other type of StorNext policy is a *Storage Manager Policy*, which governs how StorNext Storage Manager works).

Here are some important facts about StorNext Replication/Deduplication policies.

- A replication/deduplication policy exists on only one SNFS file system. For example, a policy in a file system called `/stornext/sn1` can be used only to replicate directories in that file system. A separate policy would be needed to replicate directories from file system `/stornext/sn2`.
- If a replication/deduplication policy will be used in any file system on a machine, you must configure a *blockpool* for that machine. The blockpool for a machine contains data (called blocklets) if the Deduplication feature is used, but the blockpool must be configured for replication use even if you do not use deduplication.
- A policy may be applied to one directory or more than one directory in the file system.
- A single policy can define behavior for replication sources and targets, as well as for deduplication. This single policy can also define the directories affected by the policy.
- However, it is often convenient to configure a policy that does primarily one thing. For example, you could create a policy that controls replication source behavior. Such a policy might be called a "replication source policy" or a "source policy," even though the policy could be configured to define other actions.

When configuring replication you must configure a policy for the replication source, and another policy for the replication target. You typically configure the replication source by creating a new policy for that file system. You typically configure the replication target by *editing the policy named "target"* for the file system on the target host machine.

This is an important distinction:

- Configure the replication source by *creating a new policy*
- Configure the replication target by *editing the "target" policy*

Note: When the replication source is on a different machine than the replication target (which is a typical situation,) you must use two instances of the StorNext GUI: one instance connected to the source machine, and another instance connected to the target machine.

Configuring replication is discussed in more detail in the section [Configuring Replication](#) on page 146.

Replication Process Overview

The actual replication process occurs in two stages:

- 1 **Data Movement Stage:** In this stage StorNext moves the data for files from the source file system to the target file system. Data movement occurs continuously as files are created or modified in a source directory.

Note: A configuration option allows this "continuous" data movement to be disabled during periods when the host machine or the network may be busy.

Data movement occurs in one of the following two ways.

- **Deduplicated Data:** If deduplication has been enabled for the policy that controls the source directory, deduplicated data moves from the source host machine to the target host. With deduplication enabled there may be less data moved than if the entire file were copied. This is because for deduplicated replication, only the unique deduplicated segments need to be copied.
- **Non-deduplicated Data:** If deduplication is not enabled for the policy that controls the source directory, the entire file is copied


from the source directory to the target host. The entire file is copied whenever a file is created or modified.

When data movement is in progress or even after it has just completed, the replicated files may not be visible yet in the target file system's directories. Replicated files become visible in stage 2.

- 2 File System Namespace Realization Stage:** In this stage StorNext enumerates all the files in the source directory and recreates the file name and subdirectory structure (the *namespace*) on the target file system. Unlike in the Data Movement Stage, this stage happens only at scheduled times, or when namespace realization is manually initiated by the administrator.

The following illustration shows in simple terms how replicated data flows from the one replication source to one replication target.

Figure 57 Replication Process



Files Excluded From Replication

Certain files may not be included in the replication process for various reasons. For example, a file that is open for read-only would be replicated, but a file that is open for write (including all of the various varieties of "write"), would not be replicated.

To determine which specific files were not included in the replication process, see the **Replication/Deduplication Completion Report**, which is accessible from the **Replication/Deduplication Policy Summary Report**. For more information about Replication reports, see [Replication Deduplication Reports](#) on page 250.

Here are some situations in which a file may be excluded from the replication process:

- Files that were truncated by Storage Manager before a replication policy was set up on a directory are not replicated. If you have an existing directory on which Storage Manager has been running and files are truncated, the files will not replicate from the truncated state. They must be retrieved from tape first. Once they are retrieved they will become candidates for replication and will not be truncated again until they have been either deduplicated or replicated (in the case of non-deduplication replication).
- Named pipes and device special files are not replicated.
- In both deduplication and non-deduplication replication, the completion report would mention if the file contents changed during namespace replication. This means that the replicated file on the target may represent an intermediate state taken during replication.

Replication Terms and Concepts

This section contains terms and concepts related to replication. Some terms have already been mentioned in the context of explaining replication and how it works. For these terms that have already been mentioned, this section contains a more complete, expanded definition.

Namespace Realization

Namespace refers to the directory structure which contains replicated data. Replicated data is always transferred separately from namespace data (although some small file data is transferred along with the namespace).

Namespace realization refers to the process in which the replicated directory structure (the namespace) appears on the replication target.

Because file data and namespace data is transferred separately, in some situations it might take longer for replicated data to complete transferring than for the namespace realization to complete. This is especially likely to happen if there is a backlog of file data waiting to be transferred at the time when namespace is either scheduled to run or is manually initiated.

Blockpool

The *Blockpool* is a data repository on the target. A blockpool is required on each machine used for replication or deduplication. If you use only replication, the blockpool file system can be small. If you configure deduplication as well as replication, the blockpool file system must be larger: at least large enough to hold the pool of deduplicated data segments.

When you configure StorNext for the first time, the Configuration Wizard enables you to specify the name of the file system you want to use for the blockpool.

Note: Once you specify the file system on which the blockpool resides, you cannot later choose a different blockpool file system. Use care when specifying the blockpool file system.

Blackout Period

A *Blackout* is a period during which replication does not occur. You may schedule replication blackouts for periods in which you do not want replication data transfers to occur on a busy network. For example, an administrator may create a blackout during periods when WAN usage and traffic is the heaviest. In this situation replication might run after hours when WAN usage and traffic would be much lower.

Replication Source Policy and Replication Source Directory

A *replication source policy* is a replication/deduplication policy that has "Outbound Replication" turned On via the policy's Outbound Replication tab.

The policy also has a Source Directories tab. The directories specified on this tab will be replicated, and these directories are called *replication source directories*.

Replication Target Directory

A *replication target directory* is the location to which replicated data is sent. The replication target may be a directory on a separate host machine, or it may be a directory on the source host machine. Regardless of where the target directory resides, it is very important that you use the replication target directory *only* for replicated data. Also, *do not allow users to modify the files in the replication target directories*.

When creating replication target directories, remember that the target directory must be *at least* as large as the sum of all replication source directories from which replicated data is sent. For example, if you have two source directories that are both 100GB, your replication target directory must be at least 200GB.

Replication Schedule

You can specify a *replication schedule* to define when the file system namespace realization should occur for an outbound replication schedule. For example, you might specify that you want namespace realization to occur at 6am and 6pm every day.

If you do not specify a replication schedule, you must manually run the replication policy whenever you want the realization to occur.

Replication Copies

Replication Copies is the number of copies of replicated data saved on the target. StorNext currently supports 1 to 16 replication copies per target. The number of replication copies is entered or modified in replication policies.

Bandwidth Throttling

Bandwidth Throttling refers to limiting the receive rate and transmit rate for replicated data (Replication Stage 1). This StorNext feature allows network administrators to specify (in bytes per second) a ceiling for incoming and outgoing replicated data. When bandwidth throttling is enabled, replicated data will not be transmitted or received at a rate higher than the specified maximum. Bandwidth throttling is a useful tool for controlling network traffic.

Multilink

StorNext provides support for Multilink configurations, which means you can have multiple connections on one network interface card (NIC), or even multiple connections on multiple NICs. StorNext provides a tool that shows you the NICs on your system, and allows you to specify the number of channels per NIC. On this same screen you can specify the NICs you want enabled for replication.

One advantage of using multiple NICs (or multiple channels on one NIC) is higher aggregate bandwidth because you can have multiple parallel paths. For this reason, multilink is valuable for load balancing.

When configuring multilink, be aware that the routing table of the host operating system determines the interface used for a connection to a specific endpoint. If there are multiple NICs in the same subnet as the destination IP endpoint, the host OS selects what it considers the best route and uses the interface associated with that route.

Note: An alternative to StorNext Multilink is to use the Linux bonding driver to enslave multiple Ethernet interfaces into a single bond interface. This is a Linux feature, not a StorNext feature, but is supported by the StorNext software.

Virtual IP (vIP)

Virtual IP or *vIP* is similar to an alias machine name. StorNext uses virtual IP addresses to communicate with machines rather than using the physical machine name. Virtual IPs are required in HA (high availability) environments, and are also used for multilink NICs.

Your network administrator can provide you with the virtual IP addresses and virtual netmasks you need for your system.

Note: If your replication source policy or target policy is an HA system, you must specify the vIP address in the field labeled "Address for Replication and Deduplication" on the **Outbound Replication** tab for the policy named "global" on each file system for which you will use replication. The default value for this field is "localhost".

Remember that each StorNext file system will have a policy named "global," and you should edit this field for each of those policies named "global."

Some Replication Scenarios

StorNext provides replication support to meet a variety of needs. This section describes some common replication scenarios.


Scenario 1: Simplest Replication

In this simple replication scenario, the host machine host1 contains a StorNext file system called `/stornext/fs1/`. Host machine host2 has a StorNext file system called `/stornext/fs2/`.

In this scenario we can replicate directory `/stornext/fs1/video` on host1 to file system `/stornext/fs2` on host2. Replicated files will appear in the directory `/stornext/fs2/video`, which is the default location on host2.

The following graphic illustrates replication scenario 1.

Figure 58 Replication scenario 1



Scenario 2: Replicating Multiple Copies in the Same Target File System

In this scenario the directory `/stornext/fs1/video` on host1 is again replicated to file system `/stornext/fs2` on host2. However, when the namespace realization occurs we want to retain the previous replicated target directories.

For this scenario assume that we want to keep four copies of the replication target directory. So, in file system `/stornext/fs2` on host2 we will find these four directories:


- `/stornext/fs2/video` (Contains the most recent realization)
- `/stornext/fs2/video.1` (Contains the second-most recent realization)

- /stornext/fs2/video.2 (Contains the third-most recent realization)
- /stornext/fs2/video.3 (Contains the fourth-most recent realization)

When using replication according to this scenario, in the StorNext GUI use the "Copies to Keep on Target" box on the **Outbound Replication** tab to enable multiple copies.

The following graphic illustrates replication scenario 2.

Figure 59 Replication Scenario 2



Question: Why would we want to keep multiple directories containing replicated data?

Answer: To save previous versions of the replicated directory. If you maintain only a single directory, the directory is overwritten each time replication occurs. For example, if replications happen daily at midnight, each of the replicated target directories will contain the contents of the source directory from that day's midnight replication.

You may keep from 1 to 16 copies on the target for each source directory.

Question: Will keeping extra copies use a lot of extra disk space on the target?

Answer: Not necessarily. For example, if file `video/ myTVshow.mov` has not changed for the last 4 replications, then the four files would be:

- `/stornext/fs2/video/myTVshow.mov`
- `/stornext/fs2/video.1/myTVshow.mov`
- `/stornext/fs2/video.2/myTVshow.mov`
- `/stornext/fs2/video.3/myTVshow.mov`

All of these files share the same data extents in the file system. An even greater space saving can be realized by enabling deduplication for replication, but using this feature is outside of the scope of the current scenario discussion.

Scenario 3: Replicating to Multiple Target Hosts / File Systems

In this scenario we replicate directory `/stornext/fs1/video` on `host1` to file system `/stornext/fs2` on `host2` and to file system `/stornext/fs3` on machine `host3`. Replicated files will appear in the target directories `/stornext/fs2/video` on `host2` and in `/stornext/fs3/video` on `host3`.


In this scenario we can also use the "Copies to Keep on Target" option. When "Copies to Keep on Target" is specified in a replication source policy, multiple copies are retained in each of the target file systems.

A replication source policy may specify up to three target hosts.

A target host may received replicated data from up to 5 source hosts.

The following graphic illustrates replication scenario 3.

Figure 60 Replication Scenario
3

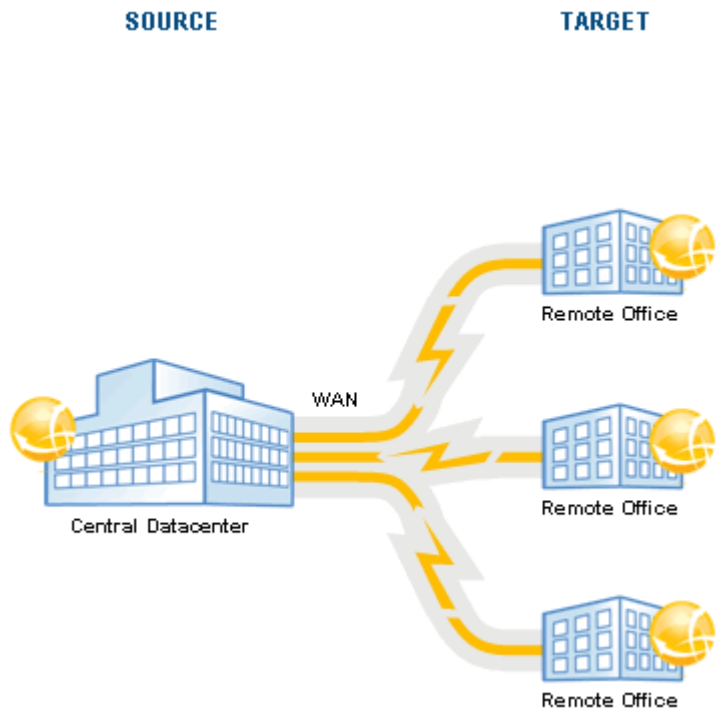


Additional Replication Possibilities

Here are some other possible replication combinations StorNext supports:


- Replication of a directory from one source host to multiple target hosts and/or multiple target file systems.

Figure 61 Replicating From
One Source to Multiple Targets



- Replication from multiple sources hosts or file systems to a single target host and file system.

Figure 62 Replicating From
Multiple Sources to One Target



- Replication on HA systems - the source host and/or the target host can be an HA pair.
- Replication with Storage Manager, where replicated data is moved to tape from either the source directory or the target host/file systems.
- Replication plus deduplication (in combination with any of the three source-to-target setups), with or without Storage Manager.

When you are first using replication, Quantum recommends beginning with simple one-to-one replication (Scenario1).

Non-Supported Replication Between Source and Target

Replicating simultaneously between a replication source and target is not currently supported by StorNext. When configuring replication, be sure to avoid this particular scenario.

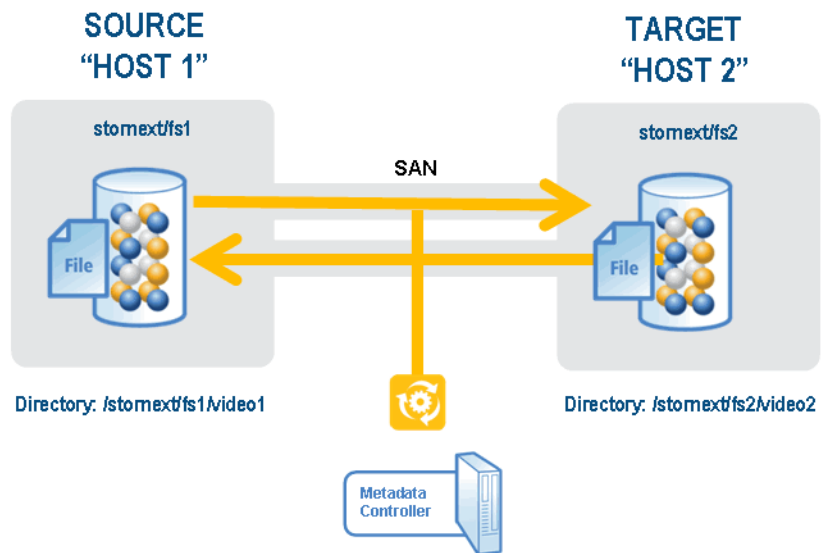
Example

In this non-supported configuration, Machine host1, file system fs1, directory video1 replicates to Machine host2, file system fs2, directory video2

While at the same time

Machine host2, file system fs2, directory video2 replicates to
Machine host1, file system fs1, directory video1

Figure 63 Non-Supported
Replication From Source to
Target



"Chained" Replication

"Chained replication" is replicating from source to a target and then to another target, and then to another target, and so on. This type of replication is not currently supported by StorNext. For example, you cannot replicate from A to B to C. When configuring replication, be sure to avoid this particular scenario.

Chained replication should not be confused with replicating from one source to multiple targets, which *is* supported. (See [Additional Replication Possibilities](#) on page 142.)

Configuring Replication

This section describes how to configure simple one-to-one replication from one source directory on one file system to one target file system. The source and target StorNext server computers can be the same machine, standalone servers, or High Availability (HA) redundant servers. When replication-target file systems are on an HA Cluster, it is best to convert the cluster to HA before configuring replication source policies that point to them. This allows the use of the virtual IP (vIP), which is required for HA configurations.

Additional configuration options for StorNext features such as HA or Replication with Storage Manager are also covered.

Before you begin configuring, make sure you have the Replication and/or Deduplication licenses required for these features. If you are using an HA configuration, basic StorNext single-server or HA Clusters should already be set up. (For more information, see [Chapter 3, The Configuration Wizard](#)).

These instructions assume you are using the StorNext Configuration Wizard and have already completed the first three steps: **Welcome**, **Licenses**, and **Name Servers**.

Note: To ensure that the policy is created properly, you **MUST** perform the following steps in the order indicated. For example, if you create the file systems without also creating the blockpool and then load data, the policy will not be created and applied.

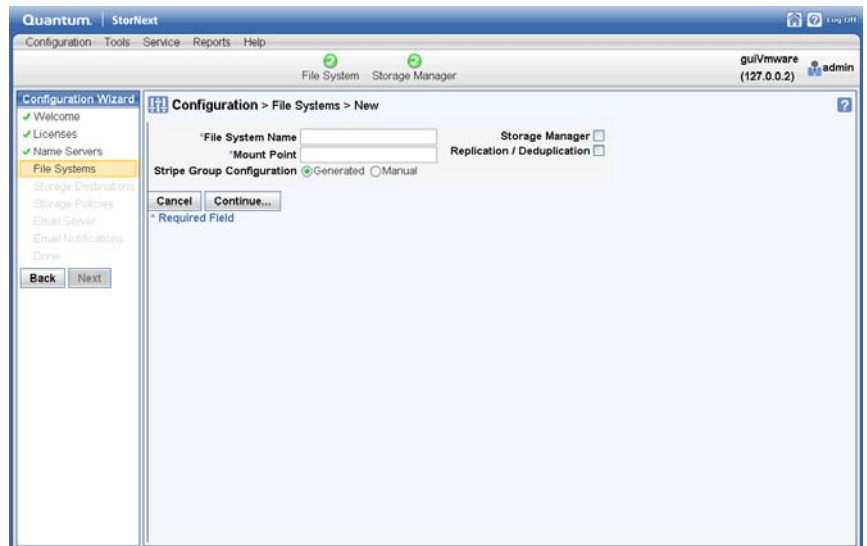
Step 1: Create Source and Target File Systems

After you complete the first three Configuration Wizard steps, the first replication step is to create file systems: the blockpool file system(s), and the source and target file systems you plan to use.

Note: Although StorNext supports replicating from multiple source hosts and file systems to multiple target hosts and file systems, for simplicity this procedure describes how to replicate between one source and one target file system on the same host.

- 1 If you have not already done so, launch the StorNext Configuration Wizard and proceed through **Welcome**, **License** and **Name Servers** steps to the **File System** step.
- 2 The **Configuration > File System** screen appears.
- 3 On the Configuration > File System screen, click **New**. The **Configuration > File System > New Screen** appears.

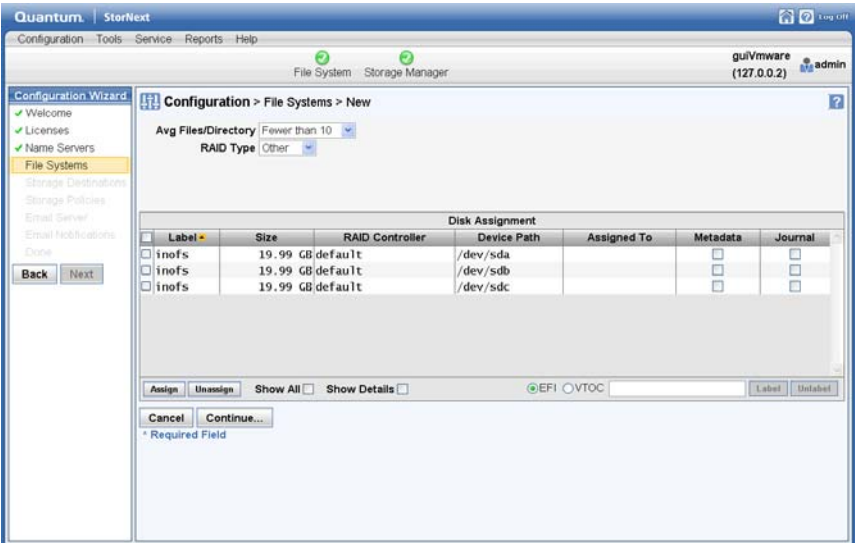
Figure 64 Configuration > File System > New Screen



- 4 At the **File System Name** field enter the name of a file system to be used as a replication source. A default mount-point path automatically appears but you can change this mount point if you wish.
- 5 Choose the **Replication/Deduplication** option. A warning message alerts you that "A blockpool has not been created." Disregard this message for now because you will create the blockpool file system in the [Step 2: Setting up the Blockpool](#).

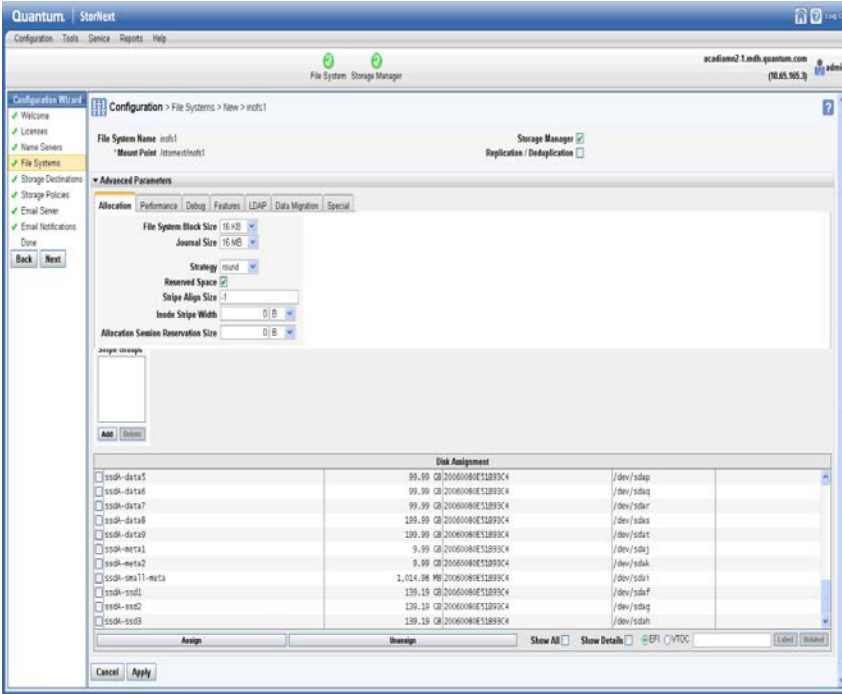
- 6 Make sure the **Generate** option is selected, and then click **Continue** to proceed.

Figure 65 Configuration > File System > New Screen 2



- 7 Select a set of LUNs for the file system, and then click **Assign**.
- 8 Click **Continue**.

Figure 66 Configuration > File System > New Screen 3



- 9 If desired, click the arrows beside the **Advanced Parameters** and **Stripe Group/Disk Management** headings to view information.
- 10 Click **Apply** to save the new file system. (For more information about creating file systems, see [Step 4: File Systems](#) on page 29.)

Creating a Target File System and Blockpool File System

- 1 Repeat the process (steps 1 - 8) and create the file system you intend to use as a target for replication on this same server.
- 2 Configure another file system for the Blockpool that has neither Data Migration Space nor Replication/Deduplication enabled.

Step 2: Setting up the Blockpool


In this step you will set up the blockpool on the blockpool file system you just created in the previous step.

- 1 Choose the StorNext Configuration Wizard's **Storage Destinations** task. The **Configuration > Storage Destinations** screen appears.

There are four tabs on this screen: **Library**, **Storage Disk**, **Replication Targets**, and **Deduplication**. When configuring replication we are concerned with the **Replication Targets** and **Deduplication** tabs. (The deduplication infrastructure is used to handle the transfer of file data for the replication feature, so it must be configured even when the deduplication feature is not used.)

- 2 Click the **Deduplication** tab. The **Configuration > Storage Destinations > Deduplication** Screen appears.

Figure 67 Configuration > Storage Destinations > Deduplication Screen (Blockpool)



- 3 Click the **Deduplication** tab. This tab has only one field called **Blockpool Host File System**. At this field select from the dropdown list the file system to use for the blockpool. (This is the file system you created in the previous step.)

Note: Once applied, the blockpool location cannot be moved to another file system. Be certain of the blockpool location before you continue.

- 4 After you select the blockpool file system, click **Apply**. A background job is started to create the blockpool.

Step 3: Creating Replication Targets

In this step you will specify the actual targets to which you want replicated data sent. (Namespace realization will also occur on these targets.)

- 1 Click the **Replication Targets** tab. The **Configuration > Storage Destinations > Replication Targets** Screen appears.
- 2 Click **Add Host**.

Figure 68 Storage Destinations
> Replication Targets Screen



- 3 At the **Hostname or IP** field, enter the host name or its IP address. If the target is an HA cluster, the address should be the vIP for that cluster.
- 4 Click **Scan Host** to populate the **Mount Point** box with appropriate file systems which are configured for replication/deduplication.
- 5 Select the file system you created for use as the target in [Step 1: Create Source and Target File Systems](#), and then click **Add**.
- 6 Click **Apply**. At this point you should see your file system listed as a replication target.

Note: If you were adding additional replication targets, you would repeat steps 3 - 6 to add additional hosts and file systems.

(Optional) Configuring Replication for an HA System

If you are planning to use replication on a high availability (HA) system, this is the point in the configuration process when you should configure HA. If you do not configure HA here, misconfiguration could result and you could be prevented from using replication on your HA system.

If you are using replication on an HA system, proceed to [Optional HA and Multilink Configuration](#) on page 161, and then return to Step 4: Create a Replication Storage Policy.

If you are *not* using replication on an HA system, proceed to Step 4: Create a Replication Storage Policy.

Step 4: Create a Replication Storage Policy

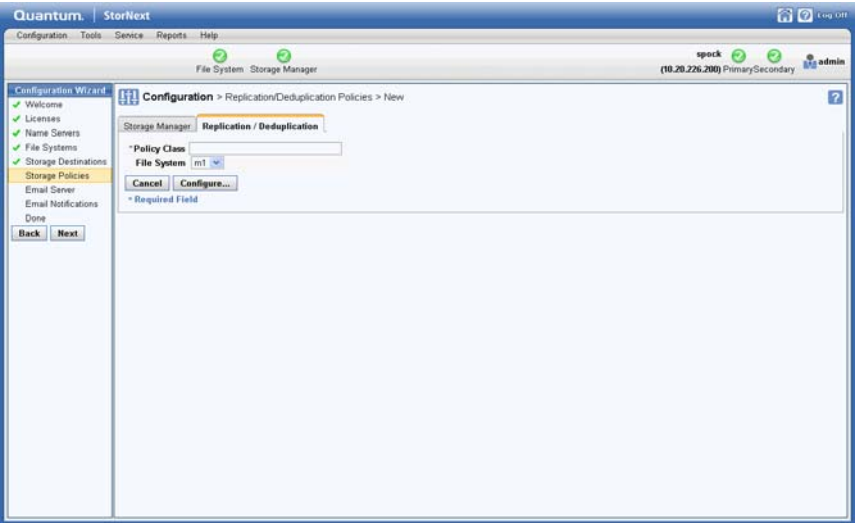
The next step in configuring replication is to create a replication storage policy. This policy contains the replication "rules" specific to your replication source and target file systems. You must create a replication policy for the source directory and enable inbound replication for the target file system.

Creating the Source Directory Replication Policy

- 1 Choose the StorNext Configuration Wizard's **Storage Policies** task. The **Configuration > Storage Policies** Screen appears.

- 2 Click **New**. The **Configuration > Storage Policies > New** Screen appears.

Figure 69 Configuration > Storage Policies > New Screen



- 3 Enter the following fields:
- **Policy Class:** The name of the new policy you are creating
 - **Policy Type:** Click the **Replication/Deduplication** tab to create a replication storage policy.

Note: The **Replication/Deduplication** tab remains unavailable (grayed out) until the blockpool directory has been completely created. Creating the blockpool directory is started on the **Storage Destinations** page's **Deduplication** tab, and proceeds asynchronously as a background job. This tab becomes enabled once that background job completes.

If there is no deduplication license (for example, if you intend to use replication but not deduplication,) creating the blockpool is still required but the background job will finish within a few seconds.

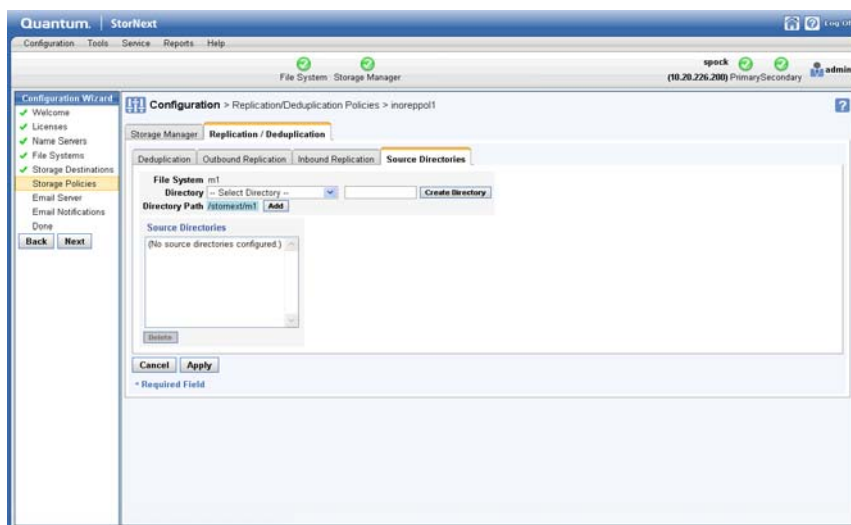
- **File System:** Choose the source file system from the dropdown list.

- Click **Configure**.

Choose the Source File System

- 1 After you click **Configure**, the screen for configuring a replication/deduplication storage policy appears.
- 2 Click the **Source Directories** tab.

Figure 70 Configuration > Storage Policies > New / Source Directories Screen

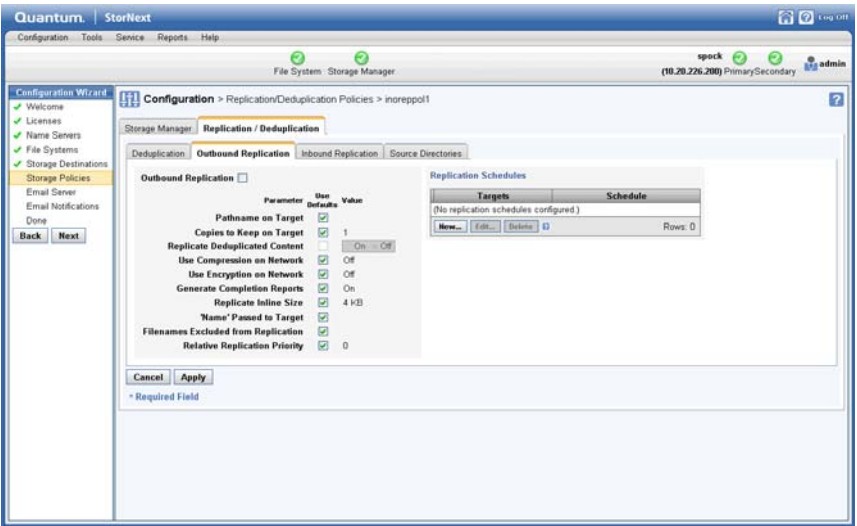


- 3 At the **Directory** field, select from the pulldown list an existing directory you want to use for the policy.
- 4 To create a new directory, enter a directory location at the field to the left of the **Create Directory** button, and then click **Create Directory** to create the specified directory.
- 5 After either selecting a directory from the pulldown list or creating a new directory, click **Add** to add the directory as the one used by the storage policy.

Enter OutBound Replication Information

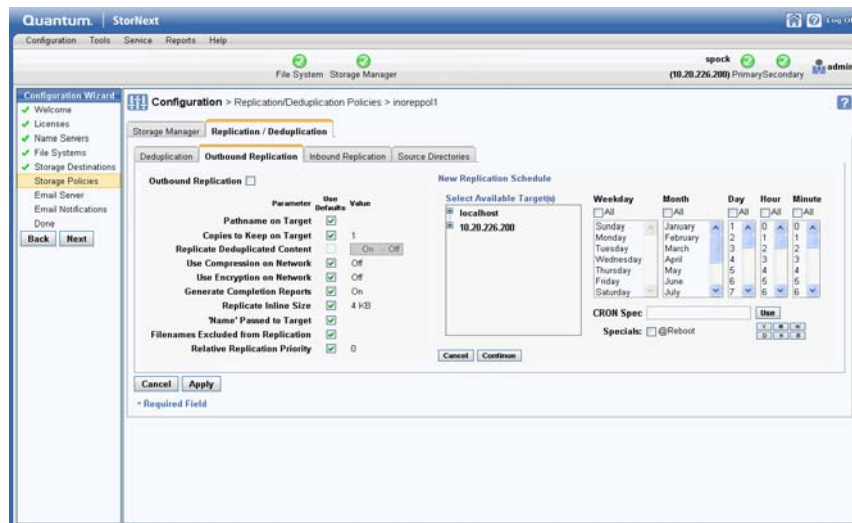
- 1 Click the **Outbound Replication** tab.

Figure 71 Storage Policies > New > Outbound Replication Tab



- 2 At the **Outbound Replication** field, enable outbound replication (going out from the source) by clicking the area to the right of the field so that **On** is displayed.
- 3 For each of the outbound replication parameters, either accept the default value by checking the box to the right of the parameter name, or uncheck the box to manually enter the value. (See the online help for more information about parameter definitions.)
- 4 (Optional) To create a new replication schedule, in the **Replication Schedules** box, click **New**. Additional fields appear where you can create a replication schedule.

Figure 72 Outbound
Replication Tab > Replication
Schedule

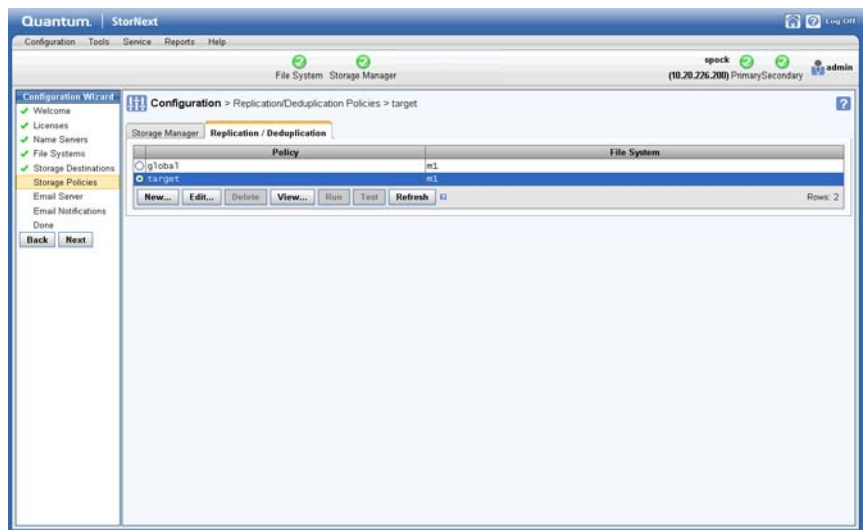


- 5 Under the heading **Select Available Targets**, select the target file system on the target server.
- 6 Create a schedule by making a selection in every column. If you select none of the schedule columns, this creates an unscheduled policy that must be run manually. The schedule shown in [Figure 72](#) will run at midnight every day. (See the online help for more information about how to enter the fields on this screen.)
- 7 Click **Continue** to complete the schedule and target selections.
- 8 Click **Apply** to finish creating the policy with the options from all of the tabs.
- 9 After a message informs you that the policy was created successfully, click **OK**.

Enter Inbound Replication Information

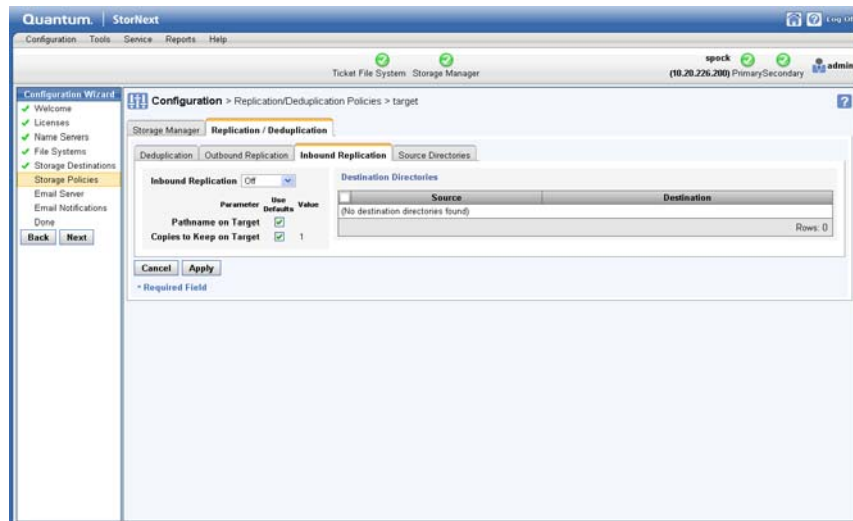
- 1 On the **Configuration > Storage Policies** screen, select the replication/deduplication policy named "target" for the replication target file system, and then click **Edit**.

Figure 73 Configuration > Storage Policies Screen (Select "target")



- 2 When the **Configuration > Storage Policies > Edit > target** screen appears, Click the **Inbound Replication** tab.

Figure 74 Storage Policies >
Edit > target > Inbound
Replication Tab



3 At the **Inbound Replication** field, select **On**.

Note: If you do not turn on replication, the process will fail and you will receive an error message saying, “Replication disabled on target.” It is VERY IMPORTANT that you enable replication by setting Inbound Replication to On.

4 Click **Apply** to finish editing the policy with your selected options.

Configuration Steps Summary

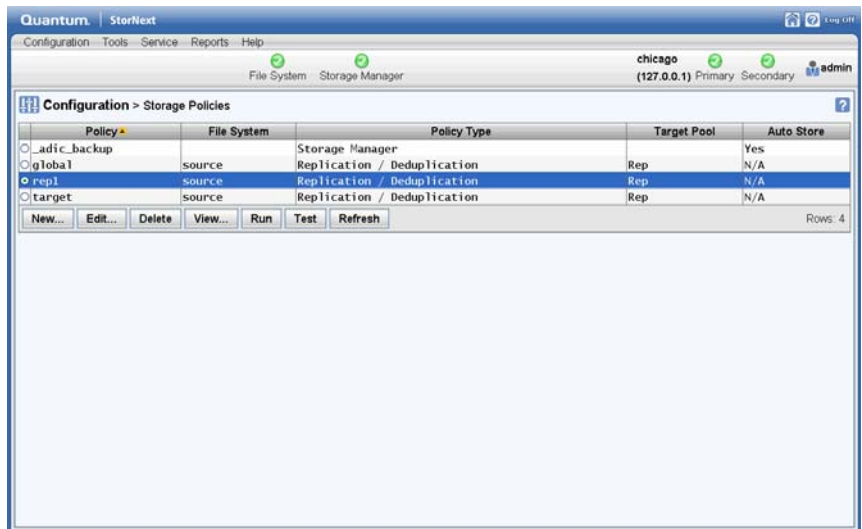
The preceding four configuration steps accomplished the following:

- Created a source replication policy and associated a source directory with it
- Selected a target file system on a target host machine and left the target directory unspecified, which uses the directory name of the source
- Set a replication schedule that runs every day at midnight
- Enabled inbound in the target policy
- Enabled outbound replication in the source policy

The contents of the source directory (additions and deletions) will now be replicated to the target directory every night. You can test this by

running the policy manually at any time on the **Configuration > Storage Policies** screen. (Select the policy you want to test and then click **Run**.)

Figure 75 Configuration > Storage Policies (Run Policy)



Scheduling Replication Blackouts (Optional)

The Replication Blackout feature provides bandwidth management by allowing you to select of a time period when you do not want replication to run. When a blackout is not in effect, replication data transfer occurs automatically in the background as data changes in the source directories, but the replicated files do not appear in the target directory until the replication policy is run.

You can set a blackout period on the source or target file system (or both) in the file systems' global policy. During the blackout period, both replication data transfer and the realization of file copies on the target are prevented from starting.

A blackout period for a source file system prevents automatic starting new data transfers or scheduled policies. However, manually started policies do run, and perform the necessary data transfers.

A blackout period for a target file system prevents all inbound data transfers from starting, which blocks both manually and automatically started source policies.

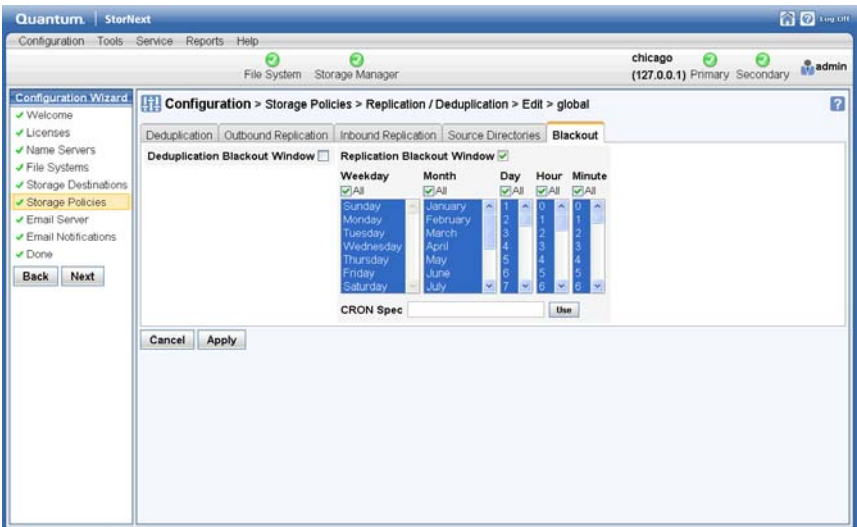
However, note the following caveats about blackouts:

- Any replication attempt (whether scheduled or initiated from the command line) which starts during the blackout on the *source* will not be started unless the force option is used. Replications started before the blackout should complete.
- Any replication request which arrives at the *target* during its blackout will be rejected by the target. The source will retry replication until the process succeeds. Replications started before the blackout will complete.

Follow these steps to set up a blackout:

- 1 Choose the StorNext Configuration Wizard's **Storage Policies** task.
- 2 On the **Storage Policies** screen, select the "global" policy for the desired source or target file system, and then click **Edit**. The **Configuration > Storage Policies > Edit** screen appears.
- 3 Click the **Blackout** tab.

Figure 76 Storage Policies > New > Blackout Tab



- 4 Click the box to the right of the **Replication Blackout Window** heading to display scheduling fields.

- 5 Specify the weekday(s), month(s), day(s), hour(s) and minute(s) when you would like to block replication from starting automatically.
- 6 Click **Apply** to save the changes in the replication/deduplication storage policy.

Optional HA and Multilink Configuration

When the High Availability (HA) feature is used with replication, a virtual IP (vIP) address must be configured to allow a replication source to use a single IP address to access whichever server is currently performing the Primary function in the target HA cluster.

The vIP is automatically moved to the correct server as part of the failover process for the HaShared file system. (See [Virtual IP \(vIP\)](#) on page 138 for an expanded definition.)

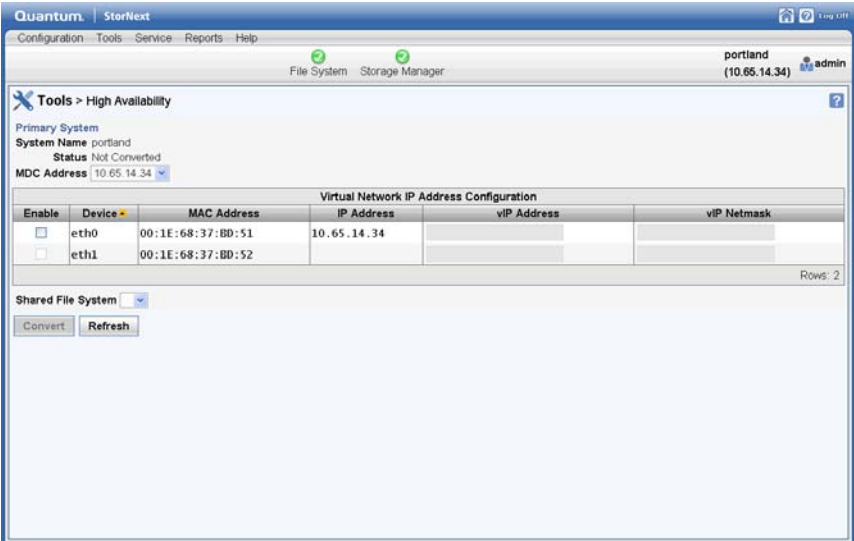
It is easiest to set up the vIP during the initial HA conversion. The vIP configuration items appear automatically at this time if a license exists for replication. It is not necessary to have a replication policy configured.

The IP address used for the vIP must be statically allocated and routable to the physical network interface of both servers in the HA cluster. Please request this IP address and netmask information from your network administrator before starting the HA conversion.

Note: This step describes only the tasks necessary for configuring replication on an HA system. For general instructions about configuring HA, see [Converting to HA](#) on page 217.

- 1 Choose **High Availability > Convert** from the **Tools** menu. The **HA (Convert)** screen appears.

Figure 77 Tools > HA
Convert Screen



- 2 At the **Shared File System** field, select from the dropdown list the file system that will be dedicated to StorNext HA internal functions.
- 3 At the **MDC Address** field, select from the dropdown list the primary system's IP address for use in communicating between HA MDCs.
- 4 Since this HA system runs a blockpool, you must configure a Virtual IP Address (viP). Under the heading **Virtual Network IP Address Configuration**, check **Enable** and then enter the viP (virtual IP) Address and viP Netmask provided by your network administrator.
- 5 Click **Convert** to convert the primary node to HA.
- 6 When the confirmation message appears, click **Yes** to proceed or **No** to exit without converting.
- 7 When a message informs you that the operation was completed successfully, click **OK**. The configuration items for the Secondary System will be added to the page.

- 8 At the **System Name** field, enter the IP address of the Secondary System to use for communications between HA MDCs, and then click **Scan Host**.
- 9 Select the IP address of the physical interface to associate with the vIP, and then click **Convert**.
- 10 When the confirmation message appears, click **Yes** to proceed or **No** to exit without converting.
- 11 When a message informs you that the conversion was completed successfully, click **OK** to continue.

Setting the IP Address of the Blockpool Server in HA Clusters

The default location of the blockpool server process is `localhost`. This is not sufficient for HA Clusters where the blockpool server moves with the Primary status to the redundant server in a failover of the HA Shared file system.

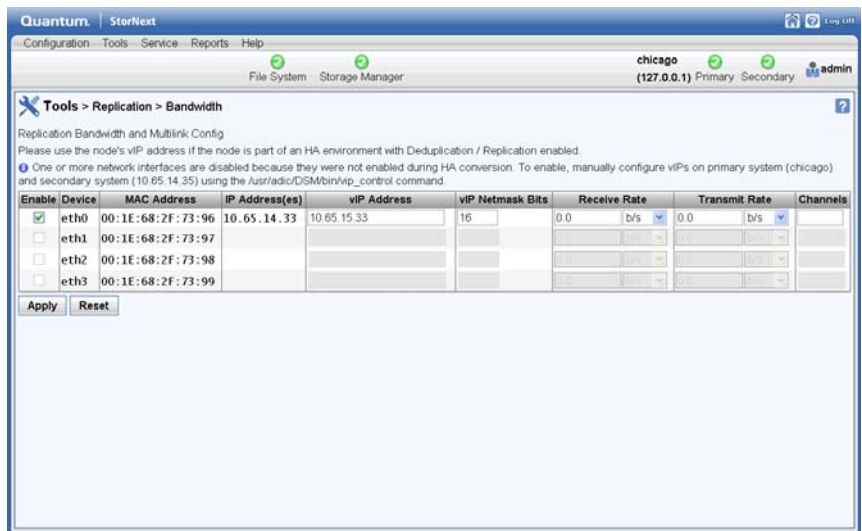
- 1 Return to the StorNext Configuration Wizard's **Storage Policies** task.
- 2 Locate the Deduplication/Replication file system, select its global policy, and then click **Edit**. (This step must be repeated for each Deduplication/Replication enabled file system.)
- 3 Click the **Deduplication** tab.
- 4 At the **Address for Replication and Deduplication** field, click the **Inherit** button.
- 5 Replace the `localhost` value with the vIP address in the **Override** box. (**Override** appears after you click **Inherit**.)
- 6 Click **Apply**.
- 7 When the confirmation message appears, click **Yes** to proceed or **No** to exit. (In this case you can safely ignore the warning about associated directories.)
- 8 When a message informs you that the operation was completed successfully, click **OK** to continue.
- 9 Repeat steps 2 thru 8 for each file system.

Configuring Multilink

Virtual IPs are also used in an HA environment if the multilink feature is configured. A virtual IP address is required for each NIC card you use for replication.

- 1 Choose **Replication/Deduplication > Replication Bandwidth** from the **Tools** menu. The **Tools > Replication > Bandwidth** screen appears.

Figure 78 Tools > Replication > Bandwidth Screen



- 2 The **Replication Bandwidth** screen displays a list of NIC cards available for replication. Select **Enable** for each NIC card you want to include in the replication process.
- 3 Enter the following fields:
 - **VIP:** Enter the virtual IP address for the NIC. (Ask your network administrator for this address as well as the virtual netmask.)
 - **VIP Netmask:** Enter the virtual netmask for the NIC
 - **Receive Rate:** Enter the maximum data reception rate (expressed in bits per second) for the replication target. When replication data is received on the target, it will not exceed this speed. (For more information, see [Bandwidth Throttling](#).)

- **Transmit Rate:** Enter the maximum data transmission rate (expressed in bits per second) for the replication source. When replication data is transmitted to the target it will not exceed this speed. (For more information, see [Bandwidth Throttling](#).)
- **Channels:** Enter the number of channels you want enabled on the NIC.

4 Click **Apply** to save your changes.

Running Replication Manually (Optional)

If you did not specify a schedule in the replication source policy, the source directory will be replicated only if you manually run the policy. If you *did* specify a schedule, you can also replicate the source directory at any time by running the policy manually.

Follow these steps to manually run replication for any replication/deduplication policy (whether it was scheduled or not):

- 1 Choose the StorNext Configuration Wizard's **Storage Policies** task. (Alternatively, choose **Storage Policies** from the **Configuration** menu.) The **Configuration > Storage Policies** screen appears. (See [Figure 75](#).)
- 2 Select the policy you want to run, and then click **Run**.
- 3 When a message informs you that the job was successfully initiated, click **OK** to continue.
- 4 To view job progress, select **Jobs** from the **Reports** menu.

Replication Statuses and Reporting

StorNext provides three ways to monitor replication status:

- [Replication Reports](#): View reports showing information pertaining to storage policies and replication targets.
- [Replication Administration](#): View the current replication status.
- [StorNext Jobs](#): View currently running StorNext jobs, including replication.

Replication Reports

There are two reports for replication: **Policy Activity** and **Policy Summary**.

- The **Policy Activity** report shows replication performance statistics.
- The **Policy Summary** report shows replication-related information for each policy.

Both of these reports also show information related to deduplication.

Access these replication reports by choosing **Replication/Deduplication** from the **Reports** menu.

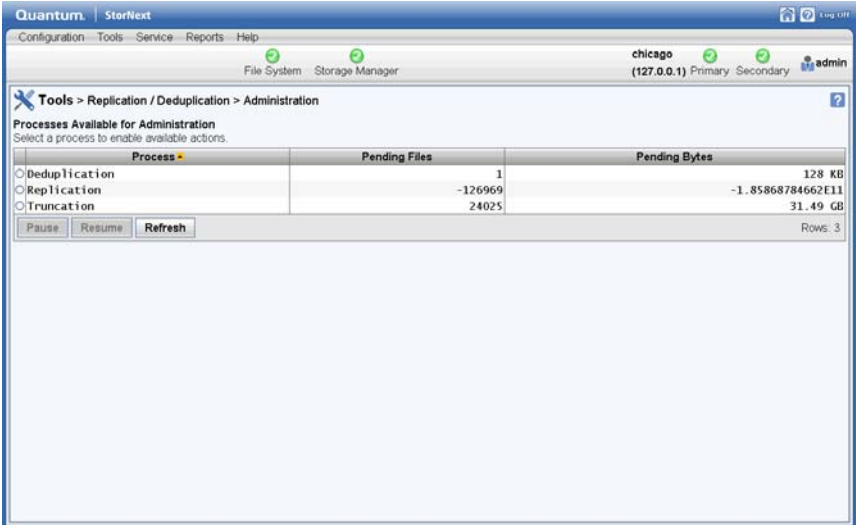
For more information about replication reports, see [Replication Deduplication Reports](#) on page 250.

Replication Administration

The **Administration** option available under the **Tools > Replication/Deduplication** menu allows you to view current replication process, or pause, resume, or stop replication.

After you choose **Administration** from the **Tools > Replication/Deduplication** menu, the **Tools > Replication/Deduplication > Administration** screen appears.

Figure 79 Tools > Replication/Deduplication > Administration Screen



The **Tools > Replication/Deduplication > Administration** screen shows the number of pending files and bytes remaining to replicate (or deduplicate or truncate).

Pausing and Resuming Replication

Near the bottom of the Administration screen are two buttons, **Pause** and **Resume**, which enable you to temporarily pause or resume replication (or deduplication or truncation) respectively. Before you pause or refresh, first select the process you want to pause or refresh: Replication, Deduplication or Truncation.

If you pause or resume replication, clicking the **Refresh** button updates the statuses shown on the Administration screen.

StorNext Jobs

At any time you can view currently running StorNext jobs, including replication. The **Reports > Jobs** screen shows the job ID and type of job, the start and end times, and the current status.

To view jobs, choose **Jobs** from the **Reports** menu. The **Reports > Jobs** report appears.

For more information about StorNext jobs, see [The Jobs Report](#) on page 230.

Troubleshooting Replication

The Troubleshooting appendix in this guide contains simple troubleshooting procedures related to replication. For more information, see [Troubleshooting Replication](#) on page 414.

For issues not covered in that section of the appendix, contact the Quantum Technical Support

Data Deduplication Overview

StorNext *data deduplication* refers to a specific approach to data reduction built on a methodology that systematically substitutes reference pointers for redundant variable-length blocks (or data segments) in a specific dataset. The purpose of data deduplication is to increase the amount of information that can be stored on disk arrays and to increase the effective amount of data that can be transmitted over networks.

For example, if the same 1 terabyte of file data appears in several different files, only one instance of that 1 terabyte needs to be retained. Each of those several files can use the same data bytes from a common storage source when the data is needed.

Quantum's deduplication not only recognizes duplicate data in the entire file, but also recognizes duplicate data ranges within files. For example, if two 1TByte files share the same data from byte 10,000,000 through byte 500,000,000, those duplicate byte ranges can be identified and stored only once. Several files may contain the same data or some of the same data, and these files can all benefit from deduplication.

How Deduplication Works

When a file is initially created in a directory managed by StorNext deduplication, all of the application data is created in that file. Later, the file may be ingested by StorNext. During the ingest process the file will be split (logically) into segments called *blobs*, which is short for “binary large objects.”


Each blob is stored in the machine's blockpool, and has a unique blob tag associated with it. From the list of a file's blob tags, StorNext can reconstitute the file with data from the blockpool.

If several files contain the same blob, only one copy is stored in the blockpool.

If StorNext file truncation is enabled for the deduplication policy, the original file can be “truncated.” (This means that the space for the original file is released and can be re-used.) When part or all of the original file data is needed by an application, the data is retrieved from the blockpool. This concept of file truncation is similar to the file truncation available with StorNext Storage Manager.

The following graphic illustrates how deduplication works.

Figure 80 Deduplication



Deduplication and Replication

If StorNext deduplication is enabled in a replication source directory, it is the blobs that get replicated from the source machine to the target machine. This happens continuously during the first stage of replication,

which is data movement. If a blob is shared by more than one file, less data is transferred than when replication occurs without deduplication.

Replicated data moves from the source machine's blockpool to the target machine's blockpool. If the source and target machine are the same, then no data needs to move for replication Stage 1.

When the replication namespace realization occurs in replication Stage 2, the replicated files appear in the target directory as truncated files. The blob tags needed to reconstitute the file are replicated along with other file metadata during Stage 2. When replication is complete, an application can access the replicated file and data will be retrieved from the blockpool as needed.

Setting Up Deduplication

This section describes the steps necessary to configure data deduplication. The easiest way to configure your system for deduplication is to use the StorNext Configuration Wizard, but you can also use the Configuration menu's options to accomplish the same tasks.

Complete these tasks to set up and enable deduplication:

- Step 1: Enable replication/deduplication when you create (or edit) a source file system.
- Step 2: Specify the file system to use for the blockpool (this is done only once per machine.)
- Step 3: Create (or edit) a replication/deduplication storage policy with deduplication enabled on the Deduplication tab.

Step 1: Creating a Deduplication-Enabled File System

Create a file system as you normally would, or edit an existing file system.

- 1 In the Configuration Wizard, choose the **File Systems** task. (Alternatively, choose **File Systems** from the **Configuration** menu.)
- 2 On the **Options** tab, enable replication by selecting **Replication/Deduplication**.

- 3 Continue creating the file system as you normally would. (If you are editing an existing file system, click **Apply** to save your changes.) For more information about creating a file system, see [Step 4: File Systems](#) on page 29.

Step 2: Specifying the Blockpool

To use deduplication you must specify the file system on which the blockpool resides. If you have already enabled replication and a blockpool already exists, you can skip this step.

The process for specifying a blockpool for deduplication is identical to specifying a blockpool for replication. For more information, see [Step 2: Setting up the Blockpool](#) on page 149 in the Configuring Replication section.

Step 3: Creating a Deduplication-Enabled Storage Policy

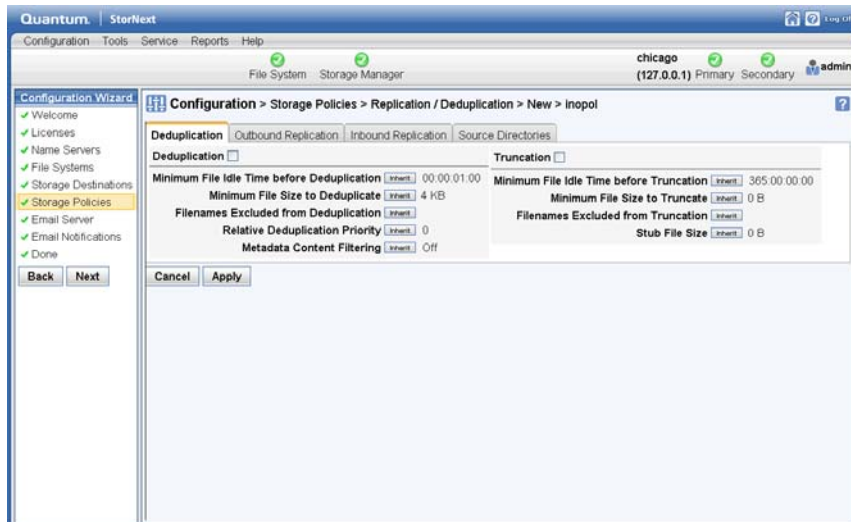
To enable deduplication you must either create a new replication/deduplication storage policy or edit an existing policy.

- 1 Choose the StorNext Configuration Wizard's **Storage Policies** task. (The **Configuration > Storage Policies** Screen appears.
- 2 If you are creating a new policy, click **New**. The **Storage Policies > New** Screen appears. (See [Figure 69](#).)

If you are editing an existing replication policy, select the policy you want to edit and then click **Edit**. Skip to Step 5.

- 3 Enter the following fields:
 - **Policy Class:** The name of the new policy you are creating
 - **Policy Type:** choose **Replication /Deduplication** to create a deduplication storage policy.
- 4 Click **Configure**. The **Replication /Deduplication Policy** screen appears.

Figure 81 Replication/
Deduplication Policy Screen



- 5 On the **Deduplication** tab, enable deduplication by clicking the field to the right of the **Deduplication** heading so that it says **On**.
- 6 Accept the displayed default values for other fields, or click the **Inherit** button beside the desired field to enter your own values. (For information about what to enter at each field, see the online help.)

Here are some especially important Deduplication parameters:

- **Minimum File Idle Time before Deduplication:** This parameter determines the interval of time for a file to remain idle before deduplication begins. The default value is 1 minute.
- **Minimum File Size to Deduplicate:** This parameter determines the minimum size a file must be in order to be eligible for deduplication. The default value is 4KB.

Data Deduplication Functions

This section describes the deduplication options on the Setup and Tools menus, which enable you to setup, administer, and manage data deduplication on your StorNext system.

Deduplication Administration

The **Tools > Replication/Deduplication > Administration** screen allows you to view the number of pending files and bytes remaining to be deduplicated (or replicated or truncate). (See [Figure 79](#).)

On this screen you can also pause or resume deduplication. The process for pausing or resuming deduplication is identical to pausing or resuming replications. (For more information, see [Pausing and Resuming Replication](#) on page 167.)

Deduplication Reports


There are two reports for deduplication: **Policy Activity** and **Policy Summary**.

- The **Policy Activity** report shows deduplication performance statistics.
- The **Policy Summary** report shows deduplication-related information for each policy.

Both of these reports also show information related to replication.

Access these deduplication reports by choosing **Replication/Deduplication** from the **Reports** menu.

For more information about deduplication reports, see [Replication Deduplication Reports](#) on page 250.



Chapter 7

Tools Menu Functions

The Tools Menu contains the following options:

- [User Accounts](#): Control user access to StorNext tasks
- [Client Download](#): Download SNFS client software
- [System Control](#): Stop or start the file system or StorNext Storage Monitor, and specify whether to automatically start StorNext at system startup
- [File and Directory Actions](#): Perform file-related and directory-related tasks on managed file systems such as storing and moving files, recovering and retrieving files and directories, and modifying file attributes.
- [File Systems](#)
 - **Label Disks**: Label disk drives
 - **Check File System**: Run a check on your file system (cvfsck) before expanding the file system or migrating a stripe group.
 - **Affinities**: Configure affinities for your file system.
 - **Migrate Data**: Migrate the file system's stripe group(s)
 - **Truncation Parameters**: Manage the file system's truncation parameters

- [Storage Manager](#)
 - **Storage Components:** View current status for libraries, storage disks, and tape drives; place one or more of these components online or offline
 - **Drive Pool:** Add, modify, or delete drive pools
 - **Media Actions:** Remove media from a library or move media from one library to another
 - **Storage Exclusions:** Specify file names to exclude from StorNext Storage Manager
 - **Truncation Exclusions:** Specify file paths to exclude from the truncation process
 - **Tape Consolidation:** Enter parameters for automatically consolidating space on tape media
 - **Library Operator:** Enter or eject media from the Library Operator Interface
 - **Software Requests:** View or cancel pending software requests
 - **Scheduler:** Schedule file system events including Clean Info, Clean Versions, Full Backup, Partial Backup, and Rebuild Policy
 - **Alternate Retrieval Location:** Specify a remote retrieval location to use in situations where files stored on tape or a storage disk.
 - **Distributed Data Mover (DDM):** Spread the distribution of data across several machines rather than the primary server.
- [Replication and Deduplication](#)
 - **Administration:** View current replication process, or pause, resume, or stop replication
 - **Replication Targets:** Add a host or directory for data replication, or edit existing replication targets
 - **Replication Bandwidth:** Configure replication bandwidth limits and multilink
- [HA](#)
 - **Convert:** Convert to a high availability configuration
 - **Manage:** Manage HA system parameters

User Accounts

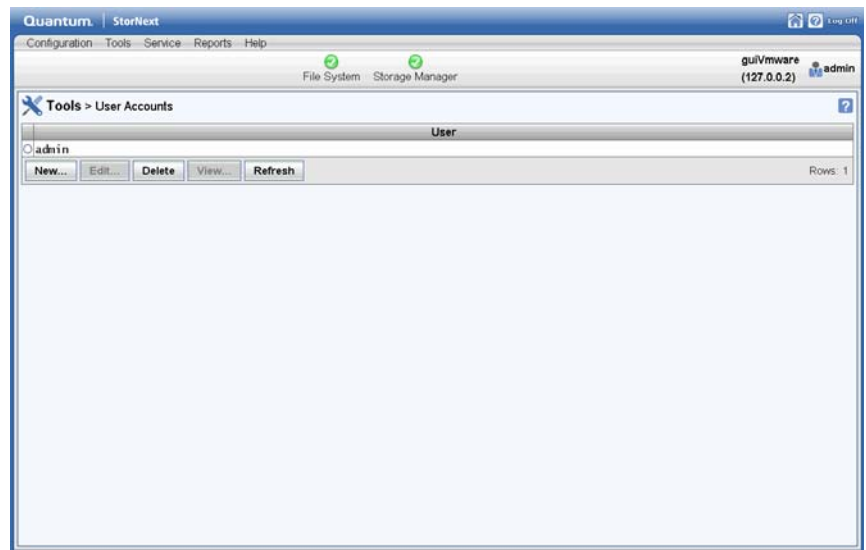
The Tools Menu's User Accounts option allows you to add new StorNext users and modify permissions for existing users. User Accounts is also where you change the admin's password.

Adding a New User

Follow this procedure to add a new StorNext user.

- 1 Choose **User Accounts** from the **Tools** menu. The **User Accounts** screen appears. All existing users and the admin are shown.

Figure 82 User Accounts
Screen



- 2 Click **New**. The **User Accounts > New** screen appears.

Figure 83 New User Screen

Quantum StorNext

Configuration Tools Service Reports Help

acadiann2.1.mdh.quantum.com (127.0.0.1) admin

File System Storage Manager Tickets (2 open)

Tools > User Accounts > New

User Name:

Password:

Confirm Password:

Access Control: ☐ Admin Defaults ☐ Operator Defaults ☒ General User Defaults

Admin Functions:

- ☐ Download Client Software
- ☐ Manage Email Notifications
- ☐ Manage Email Server
- ☐ Manage File Systems
- ☐ Manage HA
- ☐ Manage Licenses
- ☐ Manage Logging
- ☐ Manage Replication/Deduplication
- ☐ Manage Users
- ☐ Stop/Start System Components

Operator Functions:

- ☐ Cancel Software Requests
- ☐ Manage Admin Alerts
- ☐ Manage Backups
- ☐ Manage Drive Pools
- ☐ Manage Libraries
- ☐ Manage Storage Disks
- ☐ Manage Storage Manager
- ☐ Manage Tickets
- ☐ Perform File Actions
- ☐ Perform Library Operator Actions
- ☐ Perform Media Actions
- ☐ Perform Storage Component Actions
- ☐ Run Capture State
- ☐ Run Health Checks

General User Functions: ☒ Run Reports

Cancel Apply

- 3 In the **User Name** field, type the name the new user will enter at the User ID field when he or she logs on to StorNext.
- 4 In the **Password** field, type the password the new user will enter when logging on to StorNext.
- 5 Roles are grouped according to **Admin Functions**, **Operator Functions** and **General User Functions**. You can automatically pre-select all the functions for one of these three roles by clicking at the **Access Control** field **Admin Defaults**, **Operator Defaults** or **General User Defaults**. Selecting one of these roles for the new user makes it easy for you to add or remove functions by selecting or deselecting.
- 6 Select all the different roles you want the new user to have:

Admin Functions

Download Client Software	Manage Licenses
Manage Email Notifications	Manage Logging
Manage Email Server	Manage Replication/Deduplication
Manage File Systems	Manage Users
Manage HA	Stop/Start System Components

Operator Functions

Cancel Software Requests	Manage Tickets
Manage Admin Alerts	Perform File Actions
Manage Backups	Perform Library Operator Actions
Manage Drive Pools	Perform Media Actions
Manage Libraries	Perform Storage Component Actions
Manage Storage Disks	Run Capture State
Manage Storage Manager	Run Health Checks

General User Functions

Run Reports

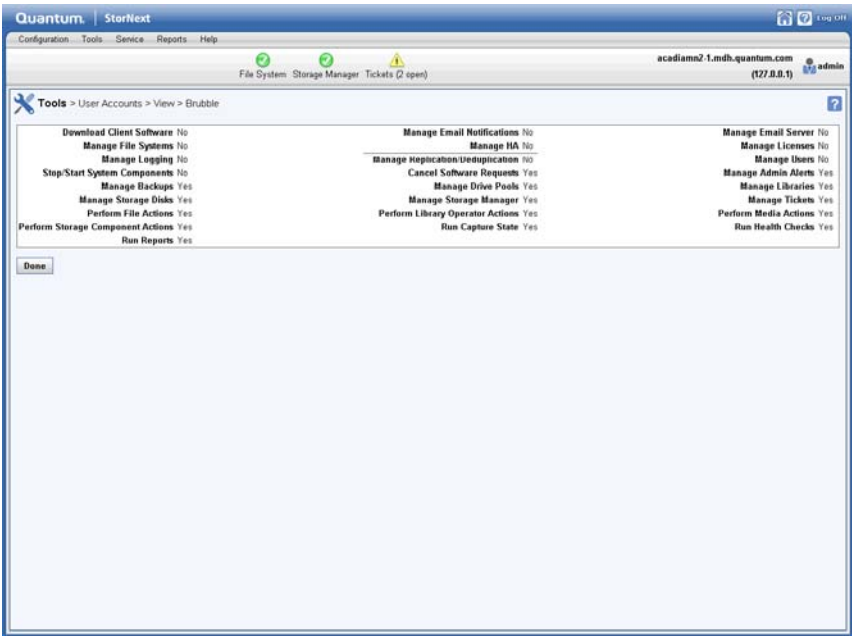
- 7 When you are satisfied with the permissions you have assigned, click **Apply** to save your changes. (To exit without saving, click **Cancel**.)
- 8 When a message informs you that the new user was successfully added, click **OK**.

Viewing an Existing User Profile

Follow this procedure to view an existing user’s profile.

- 1 From the **User Accounts** screen, select the user whose information you want to view, and then click **View**. A screen shows the parameters for the selected user.

Figure 84 View User Screen



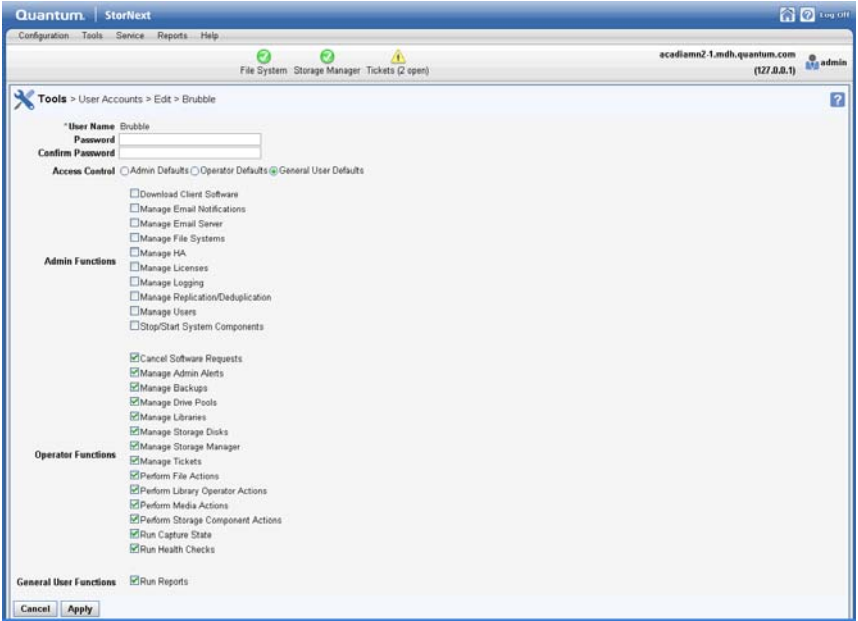
- 2 When you are finished viewing user profile information, click **Back** to return to the User Accounts screen.

Modifying an Existing User

Follow this procedure to modify an existing user's permission.

- 1 From the **User Accounts** screen, select the user whose information you want to modify, and then click **Edit**. A screen similar to the one where you added the user appears.

Figure 85 Edit User Screen



- 2 If you are editing the admin's information, you will be asked to enter and confirm the current password, and confirm that you want to modify the information for the admin. When prompted, click **Yes** to proceed.
- 3 As necessary, change the user's password and then modify permissions by selecting or deselecting roles.

Note: Only an admin user can change the admin password.

- 4 When you are satisfied with the changes you have made, click **Apply** to save your changes. (To exit without saving, click **Cancel**.)
- 5 When a message informs you that the new user was successfully modified, click **OK**.

Deleting an Existing User

Follow this procedure to delete an existing StorNext user.

- 1 From the **User Accounts** screen, select the user you want to delete, and then click **Delete**. (See [Figure 82](#).)

- 2 When the confirmation message appears, click **Yes** to proceed, or **No** to return to the **User Accounts > [admin name]** screen without saving.

Note: You cannot delete the admin.

- 3 When a message informs you that the new user was successfully deleted, click **OK**.

Client Download

The StorNext client software lets you mount and work with StorNext file systems. Note that Distributed LAN Client (DLC) client and DLC gateways are included with the standard client software packages. For more information about DCL, see [About Distributed LAN Clients](#) on page 2 and [Distributed LAN Server and Client Network Tuning](#) on page 261.

In addition to StorNext client software, Distributed Data Mover is also downloadable here. For more information about installing and using Distributed Data Mover, see [Distributed Data Mover \(DDM\)](#) on page 121.

To ensure successful operation, before you install the client software verify that the client system meets all operating system and hardware requirements listed below.

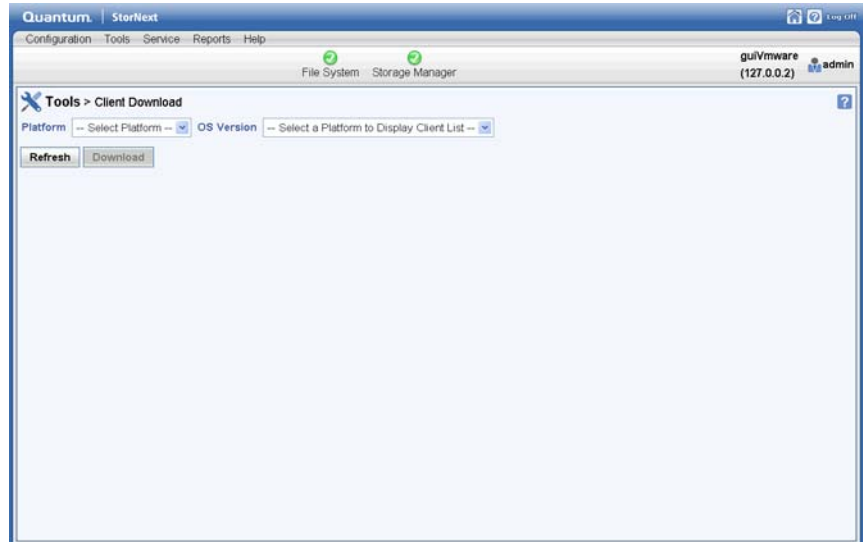
To install the StorNext client software, first download the client software from the metadata controller (MDC) as described in [Downloading Client Software](#).

After downloading the client software, install and configure it using the appropriate method for your operating system. For more information about installing and configuring client software, see the instructions in the StorNext Installation Guide.

To download client software:


- 1 Choose **Client Download** from the **Tools** menu. The **Tools > Client Download** screen appears.

Figure 86 Client Download
Screen



- 2 Select from the **Platform** list the desired operating system.
- 3 Select from the **OS Version** list the desired operating system version corresponding to the platform you selected.
- 4 When a window appears containing a link to the client software download location, click the link to begin downloading.

Figure 87 Client Download
Link



- 5 Click **Download** to begin the process.
- 6 When prompted, choose the **Save to Disk** option, and then click **OK**.
- 7 Browse to the location where you want to save the file, and then click **Save**.
- 8 After the client package has been saved, click **Done**.

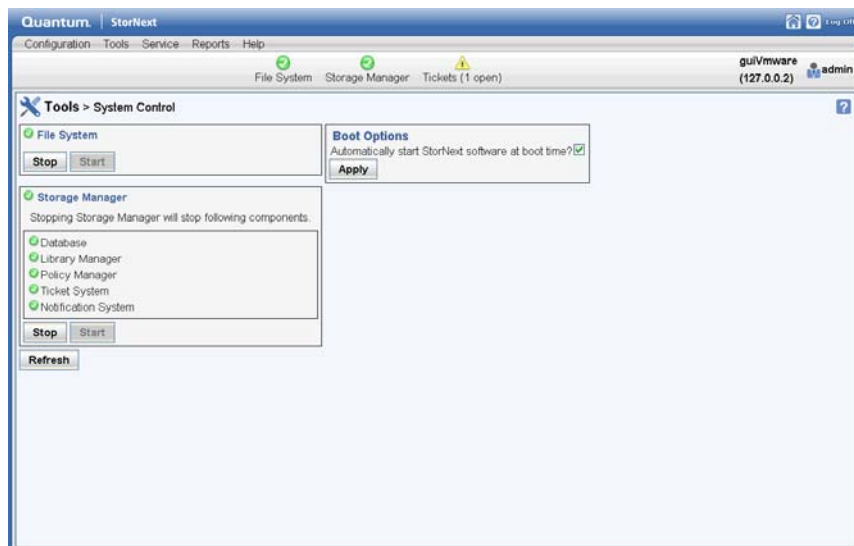
- 9 Continue with the installation procedure for your operating system as described in the *StorNext Installation Guide* or the online help.

System Control

The System Control screen enables you to tell at a glance whether StorNext File System and StorNext Storage Manager are currently started. In the case of Storage Manager, you can also see which individual components are currently started or stopped. From this screen you can start or stop File System and Storage Manager, and also specify whether you want StorNext to start automatically whenever your system is rebooted.

To access the **System Control** screen, choose **System Control** from the **Tools** menu. The **Tools > System Control** screen appears.

Figure 88 System Control Screen



Starting or Stopping StorNext File System

Most StorNext operations require that the StorNext File System be started, although there may be times when you need to stop the File System.

Click **Start** to start the File System, or **Stop** to stop the File System.

Starting or Stopping StorNext Storage Manager

StorNext Storage Manager includes the following components:

- Database
- Library Manager
- Policy Manager
- Ticket System
- Notification System

There are conditions which could cause one or more component to stop. If this happens, starting the Storage Manager restarts these stopped components.

Click **Start** to start the Storage Manager, or **Stop** to stop the Storage Manager.

Refreshing System Status

When there is a change in system status, sometimes there is a delay in updating the status. Click **Refresh** to immediately update the GUI system status.

Specifying Boot Options

If you would like StorNext to automatically start File System and Storage Manager whenever your system starts, select the option **Automatically start StorNext software at boot time?** and then click **Apply**.

File and Directory Actions

The Tools menu's **File and Directory Actions** option enables you to perform various actions on the files and directories in your library.

(Storage Manager is required to perform all of the file and directory actions described in this section.)

To access the **Tools > File and Directory Actions** screen, choose **File and Directory Actions** from the **Tools** menu. The following information is displayed for the available files:

- **Name:** The name of the file
- **Owner:** The file owner
- **Group:** The group to which the file belongs
- **Size:** The size (in bytes) of the file
- **Last Modified:** The date and time when the file was last modified

At the top of the screen is a dropdown list of Available Actions you can perform for files. Select the file for which you want to perform the action, and then choose one of these options from the Available Actions list:

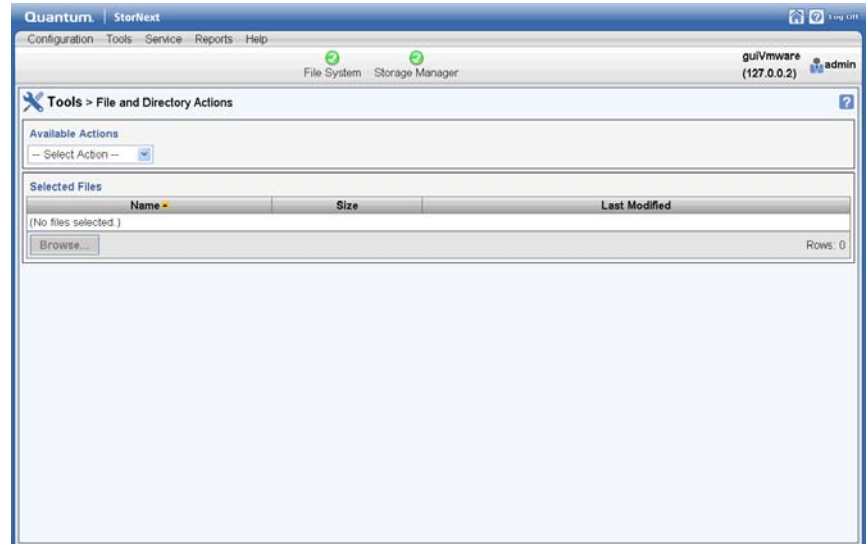
- [Store Files](#)
- [Change File Version](#)
- [Recover Files](#)
- [Recover Directories](#)
- [Retrieve Files](#)
- [Retrieve Directory](#)
- [Truncate Files](#)
- [Move Files](#)
- [Modify File Attributes](#)
- [View File Information](#)

Store Files

Choose this option to store files by policy or custom parameters.

- 1 Choose **File and Directory Actions** from the **Tools** menu. The **Tools > File and Directory Actions** screen appears.

Figure 89 File and Directory
Action Screen



- 2 Select the file you want to store. If necessary, click **Browse** and then click **All Managed Directories** to view a list of the managed directories. Select the directory containing the files to be stored. Mark the files of interest and then click **Continue** to select them.
- 3 To store the selected file according to policy, at the **Store Parameters** field, select **By Policy**.
 - c Click **Apply**.
 - d When the confirmation message appears, click **Yes** to store the file, or **No** to abort.
- 4 To store the selected file according to custom parameters, at the **Store Parameters** field, select **Custom**.
 - a Enter the following fields:
 - **Number of Copies**: Indicate the number of copies to store
 - **Truncate Files Immediately**: Select this option to truncate files immediately after storing
 - **Tape Drive Pool**: Select the tape drive pool for the selected file
 - **Minimum File Size**: Specify the minimum file size
 - **Media Type**: Specify the tape drive media type

b Click Apply.

When the confirmation message appears, click **Yes** to mount the store the file, or **No** to abort.

Change File Version

Choose this option to change the file version to a new version.

- 1 If you have not already done so, choose **File and Directory Actions** from the **Tools** menu. The **Tools > File and Directory Actions** screen appears. (See [Figure 89.](#))
- 2 Choose **Change File Version** from the **Available Actions** dropdown menu.

Figure 90 Change File Version Screen

Quantum | StorNext

Configuration Tools Service Reports Help

File System Storage Manager Tickets (1 open)

192.168.255.128 Primary Secondary admin

Tools > File and Directory Actions > Change File Version

Available Actions

Change File Version

Selected File

Name	Owner	Group	Size	Last Modified
(No file selected)				

Browse

Rows: 0

New File Version

Version	Inactive Since	Last Modified (When Stored)
(No file selected)		

Cancel Reset Apply

- 3 Select the file whose version want to change. If necessary, click **Browse** to navigate to the file location and then select the file.
- 4 At the **New File Version** field, select the new version to which you want to change for the selected file.
- 5 Click **Apply**.
- 6 When the confirmation message appears, click **Yes** to dismount the media, or **No** to abort.
- 7 Repeat steps 2 - 6 to change versions for additional files.

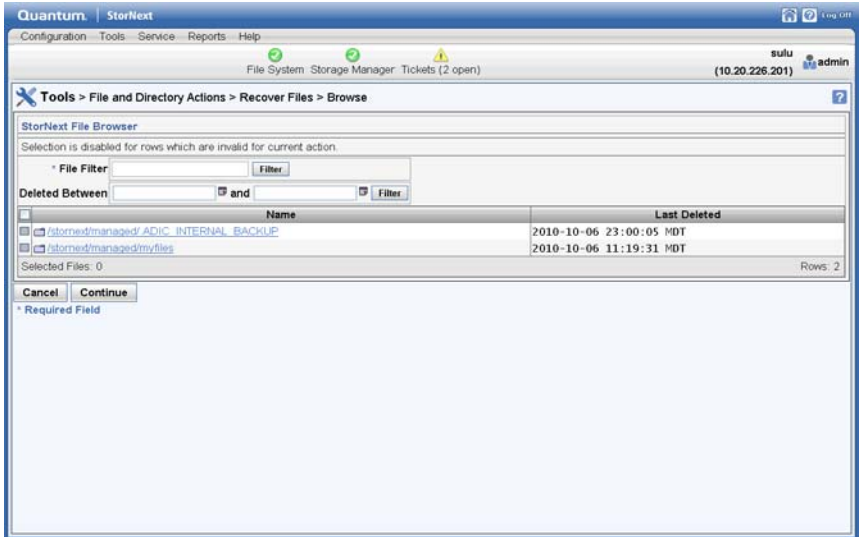
Recover Files

Choose this option to recover previously deleted files.

- 1 If you have not already done so, choose **File and Directory Actions** from the **Tools** menu. The **Tools > File and Directory Actions** screen appears. (See [Figure 89](#).)
- 2 Choose **Recover Files** from the **Available Actions** dropdown menu.
- 3 Click **Browse**. The **Tools > File and Directory Actions > Recover Files > Browse** screen appears.

The **Browse** screen shows a list of managed directories which contain files eligible for recovering.

Figure 91 Recover Files Browse Screen



- 4 To select files to recover, click the name of the desired directory.
- 5 When that directory's files are displayed, select the file(s) you want to recover by clicking the square checkbox to the left of the file's name. You can also select all files by clicking the checkbox next to the **Name** heading.
- 6 If desired, you can use one or both filters to restrict the number of files displayed:
 - **File Filter:** This filter enables you to restrict the search to files whose names contain any part of the string you enter at this

field. For example, if you enter **p**, only files which include the letter “p” in their file names are shown.

To use the File Filter, enter a letter or string in the **File Filter** field, and then click **Filter**.

- **Deleted Between:** When you enter a beginning and ending date range, only files which were deleted between those dates (inclusive) are included in search results.

To use the Deleted Between filter:

- a Position the cursor in the first field, and then click the calendar icon.
 - b Select the desired date and time, and then click the blue X icon to close the calendar.
 - c Repeat steps a and b at the second date field.
 - d Click the **Filter** button to the right of the **Deleted Between** field to apply the filtering criteria you entered.
- 7 When you are ready to recover the selected files, click **Continue**. You are returned to the **Tools > File and Directory Actions** screen, and the selected files are shown.
 - 8 Click **Apply** to start a job to recover the selected files, or click **Cancel** to abort the recovery process.
 - 9 If you clicked **Apply**, a message informs you that the recovery job has started. Click **OK** to continue.
 - 10 Choose **Reports > Jobs** to view the results of the file recovery operation.

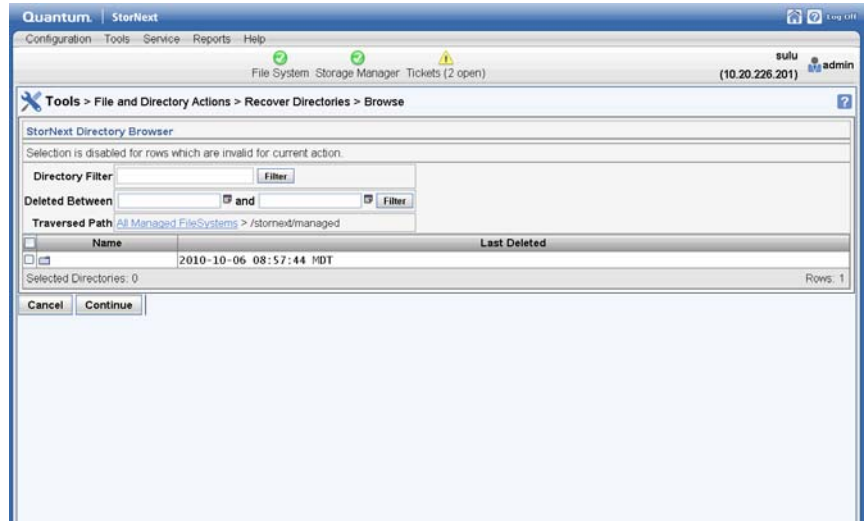
Recover Directories

Choose this option to recover previously deleted directories.

- 1 If you have not already done so, choose **File and Directory Actions** from the **Tools** menu. The **Tools > File and Directory Actions** screen appears. (See [Figure 89](#).)
- 2 Choose **Recover Directories** from the **Available Actions** dropdown menu.
- 3 Click **Browse**. The **Tools > File and Directory Actions > Recover Directories > Browse** screen appears.

The **Browse** screen shows a list of managed directories which are eligible for recovering.

Figure 92 Recover Directories
Browse Screen



- 4 To select directories to recover, click the square checkbox to the left of the directory's name. You can also select all directories by clicking the checkbox next to the **Name** heading.
- 5 If desired, you can use one or both filters to restrict the number of files displayed:
 - **Directory Filter:** This filter enables you to restrict the search to directories whose names contain any part of the string you enter at this field. For example, if you enter **p**, only directories which include the letter "p" in their directory names are shown.
To use the Directory Filter, enter a letter or string in the **Directory Filter** field, and then click **Filter**.
 - **Deleted Between:** When you enter a beginning and ending date range, only directories which were deleted between those dates (inclusive) are included in search results.
To use the Deleted Between filter:
 - a Position the cursor in the first field, and then click the calendar icon.

- b Select the desired date and time, and then click the blue X icon to close the calendar.
 - c Repeat steps a and b at the second date field.
 - d Click the **Filter** button to the right of the **Deleted Between** field to apply the filtering criteria you entered.
- 6 When you are ready to recover the selected directories, click **Continue**. You are returned to the **Tools > File and Directory Actions** screen, and the selected directories are shown.
- 7 If desired, enter one or both of the **Advanced Options**:
 - **Destination Directory**: Specify an alternative path name for the directory to which you want to save the recovered directory.
 - **Files Active At**: Click the calendar icon and specify a date and time to include only the files inside the recovered directory which were active (not deleted) on or before the specified date. Click the blue X icon to close the calendar.

The 'Destination Directory' and 'Files Active At' options are valid only when recovering a past instance of a directory. (The normal functionality of a recover operation is to recover files or directories that have been removed.) When recovering an instance of a directory, the directory may still exist (along with its contents,) but what is desired is a 'snapshot' of which files were active in the directory at a point in time.

Consequently, when recovering a directory instance you must use the 'Destination Directory' option to specify a new destination because the source directory may already exist and you do not want to recover over that directory.

Note: After recovering an instance you end up with an entirely new managed directory with no relation to the source.

Also when recovering an instance of a directory, the time of the snapshot is needed. The 'Files Active At' option allows you to specify a specific time for the snapshot.

- 8 Click **Apply** to start a job to recover the selected directories, or click **Cancel** to abort the recovery process.
- 9 If you clicked **Apply**, a message informs you that the recovery job has started. Click **OK** to continue.

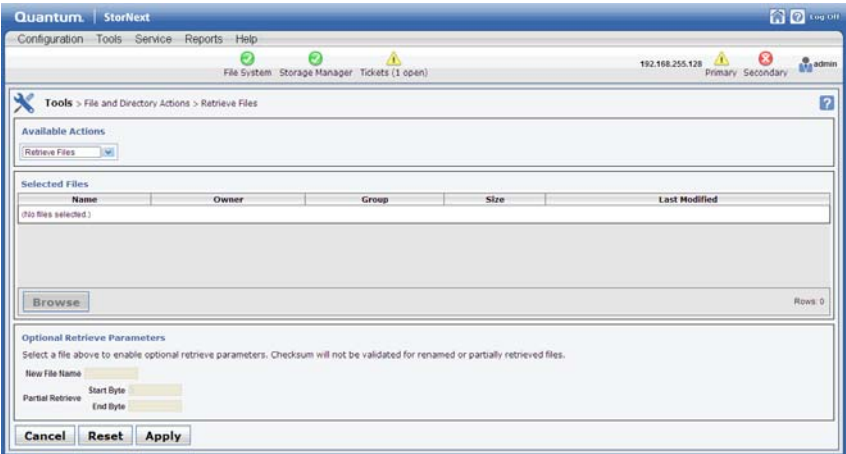
- 10 Choose **Reports > Jobs** to view the results of the directory recovery operation.

Retrieve Files

Choose this option to retrieve files which have been truncated but not deleted.

- 1 If you have not already done so, choose **File and Directory Actions** from the **Tools** menu. The **Tools > File and Directory Actions** screen appears. (See [Figure 89](#).)
- 2 Choose **Retrieve Files** from the **Available Actions** dropdown menu.

Figure 93 Retrieve Files Screen



- 3 Select the file you want to retrieve. If necessary, click **Browse** to navigate to the file location and then select the file.
- 4 If desired, enter the following **Optional Retrieve Parameters** for the selected file.
 - **New File Name:** Enter a new name to assign to the selected file upon retrieval
 - **Partial Retrieve Start Byte and End Byte:** To do a partial file retrieval, enter the file's starting and ending bytes.

When you enter these optional retrieve parameters, checksum is not validated for the selected file.

- 5 Click **Apply**.


- 6 When the confirmation message appears, click **Yes** to proceed, or **No** to abort.
- 7 Repeat steps 2 - 6 to retrieve additional files.

Retrieve Directory

Choose this option to retrieve directories.

- 1 If you have not already done so, choose **File and Directory Actions** from the **Tools** menu. The **Tools > File and Directory Actions** screen appears. (See [Figure 89.](#))
- 2 Choose **Retrieve Directory** from the **Available Actions** dropdown menu.

Figure 94 Retrieve Directory Screen



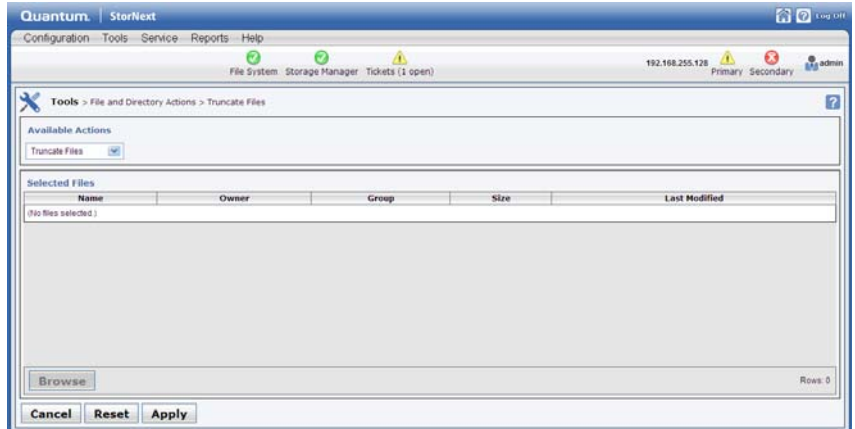
- 3 Select the directory you want to retrieve. If necessary, click **Browse** to navigate to the directory location and then select the directory.
- 4 Click **Apply**.
- 5 When the confirmation message appears, click **Yes** to proceed, or **No** to abort.
- 6 Repeat steps 2 - 5 to retrieve additional directories.

Truncate Files

Choose this option to truncate files.

- 1 If you have not already done so, choose **File and Directory Actions** from the **Tools** menu. The **Tools > File and Directory Actions** screen appears. (See [Figure 89.](#))
- 2 Choose **Truncate Files** from the **Available Actions** dropdown menu.

Figure 95 Truncate Files Screen



- 3 Select the file you want to truncate. If necessary, click **Browse** to navigate to the file location and then select the file.
- 4 Click **Apply**.
- 5 When the confirmation message appears, click **Yes** to proceed, or **No** to abort.
- 6 Repeat steps 2 - 5 to truncate additional files.

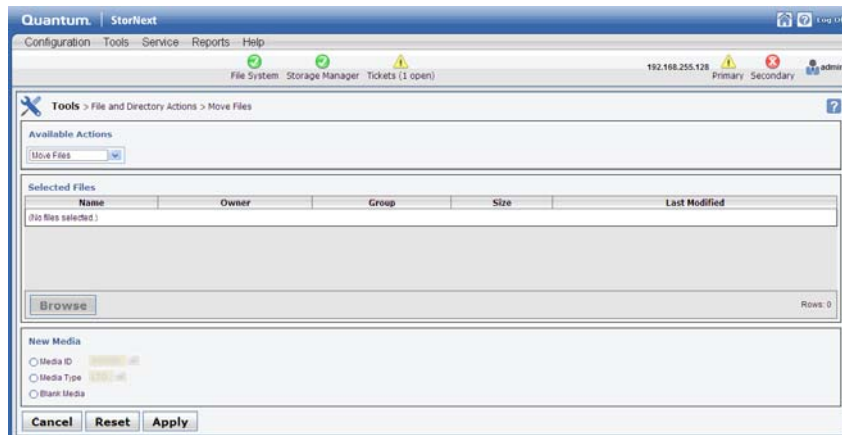
Move Files

Choose this option to move files.

- 1 If you have not already done so, choose **File and Directory Actions** from the **Tools** menu. The **Tools > File and Directory Actions** screen appears. (See [Figure 89.](#))

- 2 Choose **Move Files** from the **Available Actions** dropdown menu.

Figure 96 Move Files Screen



- 3 Select the file you want to move. If necessary, click **Browse** to navigate to the file location and then select the file.
- 4 At the **New Media** field, select one of these options for the file:
 - **Media ID:** Specify the unique identifier for the media to which you are moving the selected file
 - **Media Type:** Specify the media type for the media to which you are moving the selected file
 - **Blank Media:** Select this option if you are moving the selected file to blank media
- 5 Click **Apply**.
- 6 When the confirmation message appears, click **Yes** to proceed, or **No** to abort.
- 7 Repeat steps 2 - 6 to move additional files.

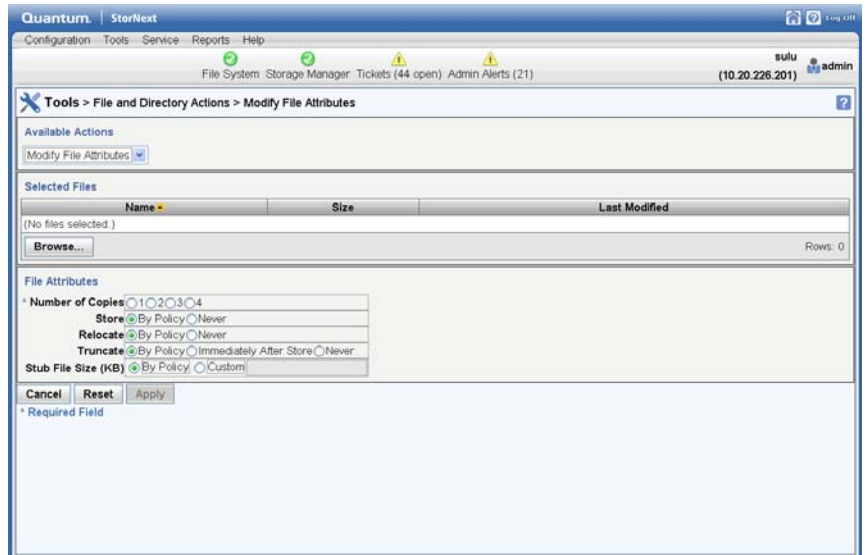
Modify File Attributes

Choose this option to modify attributes for the selected file.

- 1 If you have not already done so, choose **File and Directory Actions** from the **Tools** menu. The **Tools > File and Directory Actions** screen appears. (See [Figure 89](#).)

- 2 Choose **Modify File Attributes** from the **Available Actions** dropdown menu.

Figure 97 Modify File Attributes Screen




- 3 Select the file whose attributes you want to change. If necessary, click **Browse** to navigate to the file location and then select the file.
- 4 At the **File Attributes** field, enter these options. Required fields are marked with an asterisk. (For information about what to enter at each field, see the online help.)
 - **Number of Copies**
 - **Store**
 - **Relocate**
 - **Truncate**
 - **Stub File Size**
- 5 Click **Apply**.
- 6 When the confirmation message appears, click **Yes** to proceed, or **No** to abort.
- 7 Repeat steps 2 - 6 to apply attributes to additional files.

View File Information

Choose this option to view detailed information about the selected file.

- 1 If you have not already done so, choose **File and Directory Actions** from the **Tools** menu. The **Tools > File and Directory Actions** screen appears. (See [Figure 89](#).)
- 2 Choose **View File Info** from the **Available Actions** dropdown menu.

Figure 98 View File Info Screen



- 3 Select the files whose attributes you want to view. If necessary, click **Browse** to navigate to the file location and then select the file.
- 4 Click **File Info** to view information.
- 5 Click **Done** when you are finished viewing file information.

File Systems

The **Tools > File Systems** menu contains options that enable you to perform the following file system-related tasks:

- **Label Disks:** Apply EFI or VTOC label names for disk devices in your StorNext libraries
- **Check File System:** Run a check on StorNext file systems prior to expanding or migrating the file system
- **Affinities:** Allocate additional storage to a file system by creating a new stripe group in the file system configuration file, and assigning new disks to the stripe group
- **Migrate File System:** Move data files from a source file system to a destination stripe group, freeing stripe groups so they can be removed from an existing StorNext file system
- **Truncation Parameters:** Enter truncation parameters for your file systems in order to free up file storage that isn't being actively used

These tasks are described in [Chapter 4, File System Tasks](#)

Storage Manager

The **Tools > Storage Manager** menu contains options that enable you to perform the following Storage Manager-related tasks:

- **Storage Components:** View your system's libraries, storage disks, and tape drives, and place those devices online or offline
- **Drive Pool:** View, add, edit, or delete drive pools (groups of tape drives allocated for various administrator-defined storage tasks)
- **Media Actions:** Perform various actions on the storage media in your library
- **Storage Exclusions:** Specify file names to exclude from StorNext Storage Manager

- **Truncation Exclusions:** Specify file paths to exclude from the truncation process
- **Tape Consolidation:** Enter parameters for automatically consolidating space on tape media
- **Library Operator Interface:** The StorNext Library Operator Interface allows you to perform media-related actions remotely from the library
- **Software Requests:** View current software requests in progress or cancel a request. (A *software request* is a StorNext GUI request to the command line interface to perform an action.)
- **Scheduler:** Schedule tasks to run automatically based on a specified schedule
- **Alternate Retrieval Location:** Specify a remote retrieval location to use in situations where files stored on tape or a storage disk.
- **Distributed Data Mover:** Spread the distribution of data across several machines rather than the primary server.

These tasks are described in [Chapter 5, Storage Manager Tasks](#)

Replication and Deduplication

The **Tools > Replication/Deduplication** menu options enable you to perform the following tasks related to replication:

- **Administration:** View current progress for replication, data deduplication, and truncation operations. Also pause or resume replication, deduplication and truncation.
- **Replication Targets:** Add replication hosts and mount points to your replication targets, and edit properties for existing hosts and mount points. Also delete unwanted replication targets.
- **Replication Bandwidth:** Configure multilinks and bandwidth limits for replication.

Replication and deduplication tasks are described in [Chapter 6, Replication and Deduplication](#)

HA

The **Tools > HA** menu options enable you to perform the following HA-related tasks:

- **Convert:** Convert a shared file system to high availability configuration
- **Manage:** View the current status of the file systems on your HA system and perform various HA-related functions such as starting or stopping nodes on the HA cluster

These tasks are described in [Chapter 9, Converting to HA](#)



Chapter 8

Service Menu Functions

The StorNext **Service Menu** contains the following options:

- **Health Check:** Perform one or more health checks on StorNext and view recent health check results
- **Capture State:** Obtain and preserve detailed information about the current StorNext system state
- **System Backup:** Run a backup of StorNext software
- **Admin Alerts:** View informational messages about system activities
- **Tickets:** View, edit, or close service tickets generated for the system
- **Logging:** Enables robust debugging mode for advanced tracing

The Health Check Function

The Health Check feature enables you to run various diagnostic checks on your StorNext system. This screen shows the available tests, as well as the start time, finish time, and status for the last time each test ran.

Here are the diagnostic tests available:

- **Archive:** Verify that all configured archives are online

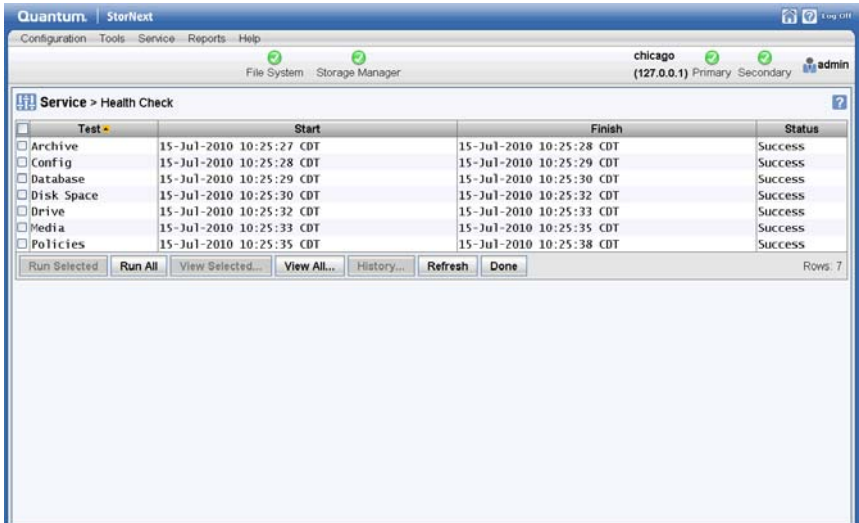
- **Config:** Verify that affinities are configured correctly in SNSM for managed file systems, and that SNSM-managed file systems are identified and configured correctly
- **Database:** Verify database integrity.
- **Disk Space:** Verify that enough disk space exists for the SNSM database tables, logging, 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 this procedure to run a health check.

- 1 Choose **Health Check** from the **Service** menu. The **Service > Health Check** screen appears.

Figure 99 Health Check Screen



- 2 Select one or more tests to run by clicking the desired check. (There is no need to hold down the **Control** key while clicking.) To deselect a test, click it again.
- 3 Click **Run Selected** to run the tests you selected. Or, to run all tests, click **Run All**.

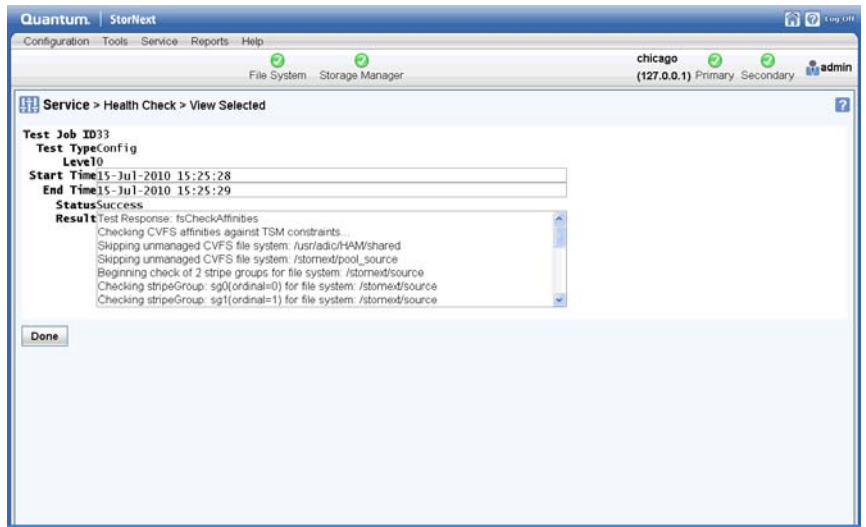
Viewing the Health Check Results

After a test has been run successfully (as indicated by “Success” in the Status column), you can view test results.

- 1 To view results for one or more tests, select the desired tests and then click **View Selected**.
- 2 To view results for all successfully completed tests, click **View All**.
- 3 When you are finished viewing, click **Done**.

Regardless of which View option you choose, test results are shown for the last successful tests completed regardless of the date on which they ran. For example, if you select two tests and then click **View Selected**, there might be an interval as long as a week or more between the finish dates for the two tests.

Figure 100 Health Check > View Selected Screen

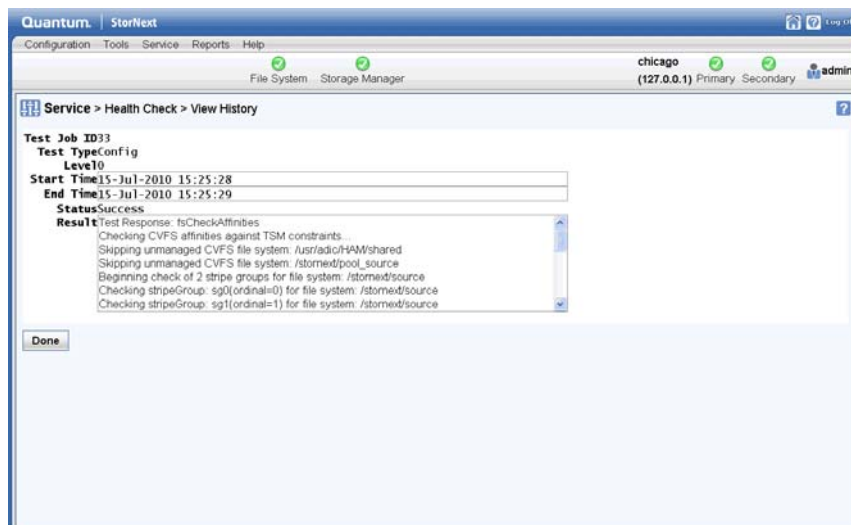


Viewing Health Check Histories

You can also view a history (up to five runs) of each health check test that has been run. This can be useful for comparing results over a time span.

- 1 Select the tests whose history you want to view.
- 2 Click **History**. The **Service > Health Check > View History** screen appears.

Figure 101 Health Check > View History Screen



- 3 When you are finished viewing history information, click **Done** to return to the **Service > Health Check** screen.

The Capture State Function

The StorNext Capture State feature 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.

Running Capture State 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.

Creating a Capture State Log

Follow this procedure to create a Capture State log:

- 1 Choose **Capture State** from the **Service** menu. The **Service > Capture State** screen appears. Any previously captured snapshots are shown.

Figure 102 Capture State Screen



- 2 Click **Capture**. The **Capture State Status** window is shown. The capture file appears after the process completes.
- 3 Click **Download** to save the generated file.
- 4 To view the file, choose the **Open with** option and then click **Browse** to navigate to an application such as WinZip capable of reading tar.gz files.
- 5 To save the file, choose the **Save to Disk** option and then navigate to the location where you want to save the file.

Deleting a Previous System State Capture

Follow this procedure to delete an unwanted Capture State file.

- 1 If you have not already done so, choose **Capture State** from the Service menu. The **Service > Capture State** screen appears. (See [Figure 102](#) on page 207.) All previously captured snapshots are shown.
- 2 Select the file you want to delete, and then click **Delete**.
- 3 When a confirmation screen prompts you to confirm that you want to delete the file, click **Yes** to continue or **No** to abort.
- 4 After the status screen informs you that the file was successfully deleted, click **OK**.

Creating a Capture State for an HA Secondary Node

When you use the Capture State feature on an HA system, system information for the primary node is captured by default.

To create a Capture State file for the secondary node, click **Capture Secondary**. Information about your secondary node is captured and saved to a file on the primary node. After the capture process completes, this file appears in the list of Capture State files.

As with Capture State files for the primary node, you can download or delete Capture State files for the secondary node. The processes for downloading or deleting Capture State files for the secondary node is identical to downloading or deleting a Capture State file for the primary node.


Note: The Capture Secondary function applies only to HA systems.

The System Backup Function

The Service menu's Backup option allows you to perform a full or partial backup.

- 1 Choose **Backups** from the **Service** menu. The **Service > Backup** screen appears.

Figure 103 Backup Screen



- 2 Click **Full Backup** to perform a full backup, or click **Partial Backup** to perform a partial backup.
- 3 After a message informs you that the backup was initiated successfully, click **OK**.

The Admin Alerts Function


Admin alerts are informational messages about system activities you might want to be aware of, but are not necessarily an error condition. For example, issues related to the Distributed Data Mover feature generate admin alerts. Admin alerts typically 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
- 1 To view admin alerts, select **Admin Alerts** from the **Service** menu. The **Service > Admin Alerts** screen appears.

Figure 104 Admin Alerts
Screen



- 2 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 an individual alert by selecting the desired alert and then clicking the **Delete** button
 - Delete all alerts by clicking the **Delete All** button

The Tickets Function


The Service menu's Tickets option allows you to view 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. On this screen you can view the ticket number, current status, priority, date and time the ticket was last updated, and a brief summary of the error condition.

By default, tickets are listed with the most recently opened tickets displayed first. If desired, you can click the column headers to change the sorting. For example, click the Ticket heading to display tickets in ascending or descending numerical order.

Viewing Ticket Information

- 1 From the StorNext home page, choose **Tickets** from the **Service** menu. The **Service > Tickets** screen appears.

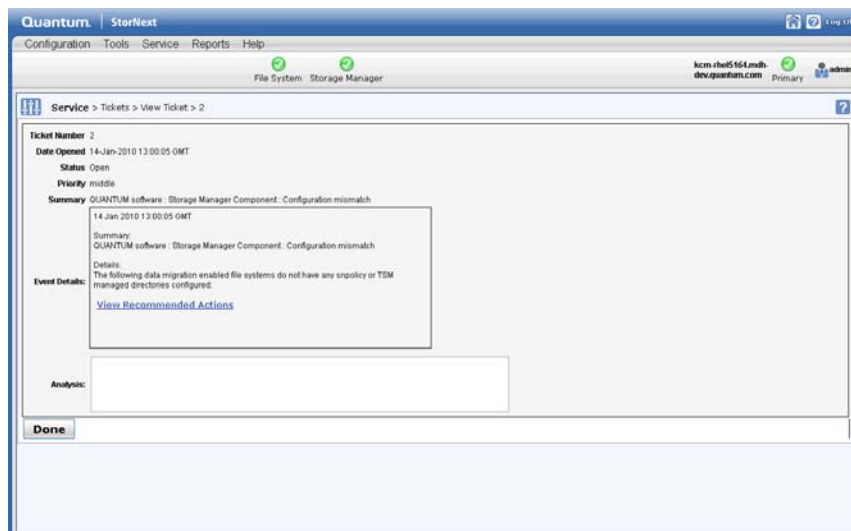
Figure 105 Tickets Screen



The **Service > Tickets** screen provides the following information:

- **Ticket:** The RAS ticket number, displayed in the order in which it was created
 - **Status:** 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
- 2 If desired, change the display by choosing **Show All Tickets**, **Show Closed Tickets**, or **Show Open Tickets** in the dropdown list at the bottom of the screen.
 - 3 Highlight the ticket you wish to view, and then click **View**. The **Service > Tickets > View Ticket > [number]** screen appears.

Figure 106 Tickets > View Ticket Screen



This screen provides the following information:

- **Ticket Number:** The number of the ticket in the displayed ticket list
- **Date 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:** A brief description of the ticket.
 - **Event Details:** Detailed information about event that triggered the ticket, including a link that allows you to View Recommended Actions which will help you correct the fault or condition
 - **Analysis:** Any user-entered comments pertaining to the fault or condition, such as a recommended action
- 4 To see recommended actions for the ticket, click the **View Recommended Actions** link. The **Recommended Actions** screen appears provides information and steps to correct the condition or fault that generated the RAS ticket. Follow the instructions on the screen to correct the condition or fault. When you are finished viewing the recommended actions, close the window.
 - 5 When you are finished viewing ticket information, click **Done** to return to the **Service > Tickets** screen.

Editing Ticket Information

Follow this procedure to add comments or notes to the ticket in the **Analysis** field:

- 1 Select the desired ticket and then click **Edit**. The **Service > Tickets> Edit Ticket > [number]** screen appears.

Figure 107 Tickets > Edit
Ticket Screen



- 2 Make your comments or notes in the **Analysis** field.
- 3 Click **Apply** to save your changes. When you are ready to return to the previous screen, click **Close**. (To return to the previous screen without saving your changes, click **Cancel**.)

Closing Tickets

When you no longer need to retain ticket information, you can close (delete) selected tickets or all tickets by following this procedure:

- 1 To close a specific ticket, select the desired ticket and then click **Close**.
- 2 To delete all tickets, click **Close All**.

Logging

The Service menu's Logging option is a robust debugging tool which enables you to turn on tracing for various system components. The feature is useful if you have been asked by Quantum Service personnel

to enable debugging for one or more components in order to help them identify and diagnose a particular error.

When logging (debugging) is enabled, information is copied to the same location as regular log files. (For more information about logs, see [StorNext Logs](#) on page 229.)


Note: The default value for the four system components for which you can enable logging is “disabled.” Enabling logging can have a minor impact on overall system performance, so you should not enable logging for a component unless you have been instructed to do so by a Quantum Support representative.

Enabling Logging

Follow these steps to enable logging:

- 1 From the StorNext home page, choose **Logging** from the **Service** menu. The **Service > Logging** screen appears.

Figure 108 Logging Screen



- 2 As necessary, enable logging for any of the following system components:
 - **Application Debug Logging:** This option enables debugging for StorNext.

- **Replication/Deduplication Debug Logging:** This option enables debugging for StorNext policies such as SNPolicy.
 - **Web Services Debug Header Logging:** This option enables debugging for Web-related components specific to Web headers.
 - **Web Services Debug Content Logging:** This option enables debugging for Web-related components specific to Web content.
- 3 Click **Apply** to enable debugging for the selected components. (Or, click **Cancel** to abort.)



Chapter 9

Converting to HA

The StorNext High Availability (HA) feature allows you to operate a redundant server that can quickly assume control of the primary server's operations in the event of software, hardware and network failures.

This chapter describes how to configure HA for StorNext. For a much more detailed discussion about how HA works, see [Appendix C, High Availability Systems](#).

HA Overview

The StorNext HA feature is a special StorNext configuration with improved availability and reliability. The configuration consists of two similar servers, shared disks and possibly tape libraries. StorNext is installed on both servers. One of the servers is dedicated as the initial primary server and the other the initial standby server.

StorNext File System and Storage Manager run on the primary server. The standby server runs StorNext File System and special HA supporting software.

The StorNext failover mechanism allows the StorNext services to be automatically transferred from the current active primary server to the standby server in the event of the primary server failure. The roles of the servers are reversed after a failover event. Only one of the two servers is

allowed to control and update StorNext metadata and databases at any given time. The HA feature enforces this rule by monitoring for conditions that might allow conflicts of control that could lead to data corruption.

StorNext provides two main HA functions: **Convert (to) HA** and **Manage HA**.

HA Terms and Concepts

This section defines key terms and concepts you should become familiar with before converting to an HA system.

Failover

Failover is the process of passing control of a file system from an FSM on one MDC to a standby FSM on a redundant MDC. When that FSM is for the HaShared file system, Primary status transfers to the redundant MDC along with all the processing that occurs only on the Primary MDC. This includes all the HaManaged FSMs, the Storage Manager processes, and the blockpool server. When an FSM does not relinquish control cleanly, an HA Reset can occur to protect against possible corruption of file system metadata and Storage Manager databases. (See Primary Node and Secondary Node.)

Primary Node

The *primary node* is the main server in your configuration. Processing occurs on this server until system failure makes it necessary to switch to another server. Also known as the *local node*. The primary status is transient and dynamic, not fixed to a specific machine.

Secondary Node

The *secondary node* is the redundant or secondary server in your configuration. If your primary server fails or shuts down, processing automatically moves to this secondary server so there is no interruption in processing. Like primary status, the secondary status is transient and dynamic, not fixed to a specific machine. Also known as the *peer node*.

Virtual IP (vIP)

Virtual IP or *vIP* is a fixed IP address that is automatically associated with the Primary MDC to provide a static IP address for replication and deduplication access to the target server in an HA cluster, and for access to the blockpool.

Following are some general requirements for vIP addresses as they apply to HA:

- 1 The vIP should be static (currently StorNext supports only static IP for HA).
- 2 The NIC should have a *physical* IP address assigned.
- 3 The vIP should be a real and unique IP address.
- 4 The vIP should be reachable by other nodes, and you should also be able to reach other nodes from the vIP address. For this reason, Quantum recommends that the vIP address be on the same subnet of the physical IP address of the same NIC.

When the NIC is also involved in multilink communication, the following additional requirement applies:

- The grouping address (taking the first configured maskbits of the IP address) of the physical and vip IPs on the same NIC should be the same, and unique on the node.

Your local Network Administrator should provide a reserved IP address and netmask for this purpose.

For more information about multiple-NIC configurations, see [Configuring Multiple NICs](#) on page 342.

Virtual Netmask

This is a 32-bit mask used to divide a virtual IP address into subnets and specify the network's available hosts.

HA Reset

This is the HA mechanism introduced in StorNext 4.0 to replace the previous *STONITH* mechanism. Like STONITH, HA Reset also has two nodes with one operating as primary or active node, and the other operating as the secondary or standby node. The primary node can reset itself on the hardware level. This new HA feature does not require a power brick to reset a node.

Preparing for HA Conversion

Before you convert to an HA system, you should assess your needs and current configuration. At a minimum, both the primary and secondary node should meet the minimum configuration requirements outlined in the *StorNext Installation Guide*.

You will also need to reserve an IP address in your local domain for use as the virtual IP address for using the HA cluster as a replication/deduplication target, so obtain an IP address and netmask from your network administrator.

Caution: Before you attempt this or any other major system configuration change, you should make a complete backup before proceeding.

Pre-Conversion Steps

Before converting to HA, you should perform the following steps:

- 1 Identify two servers, which must have similar hardware running the same version of Linux, and have identical LAN and SAN connectivity. For example, on multiple Ethernet port connections, both systems must be connected to the same Ethernet ports (eth0 on System A and System B going LAN1, eth1 on System A and System B going to LAN2, etc.)
- 2 Synchronize the clocks on both systems.
- 3 Install StorNext on both servers.
- 4 Enter StorNext license information on both server nodes.

Note: Although licenses must be entered on both HA MDCs, StorNext must be run only on the secondary for this purpose.

Also, if you manually edit the `license.dat` file, you must restart StorNext after making changes.

- 5 Launch StorNext on one server.

- 6 Configure an unmanaged file system for use as the HA shared file system. This unmanaged file system must be a file system that is not used for replication. (For more information about creating a file system, see [Step 4: File Systems.](#))

HA and Distributed LAN Clients

On a StorNext HA system using the StorNext Distributed LAN Client/Server (DLC) feature:

When configuring DLC Server on the MDCs of an HA cluster, it must be configured by-hand on each MDC. Service will be lost when an HA Reset occurs, so DLC clients should be configured to access the DLC file systems through both MDCs.

This practice allows for the best and highest availability of the DLC capability. Ideally, each node in the HA pair should have the same number of NICs and be on the same networks.

For more information about multiple-NIC configurations, see [Configuring Multiple NICs](#) on page 342.

Converting to HA

This section describes the configuration steps necessary to convert a node to a primary HA server. Converting to HA consists of selecting your dedicated unmanaged StorNext file system for use as the controlling shared file system, and then instructing StorNext to convert the node to HA.

Note: The **Convert** menu option will be unavailable (grayed out) on the **Tools** menu if you do not have a secondary system. If you have not already done so, specify a secondary system by using the Name Servers function. For more information, see [.Step 3: Name Servers](#) on page 22.

Following are some other things you should be aware of concerning the HA conversion process:

- The HA shared file system **MUST** be configured as an unmanaged file system.

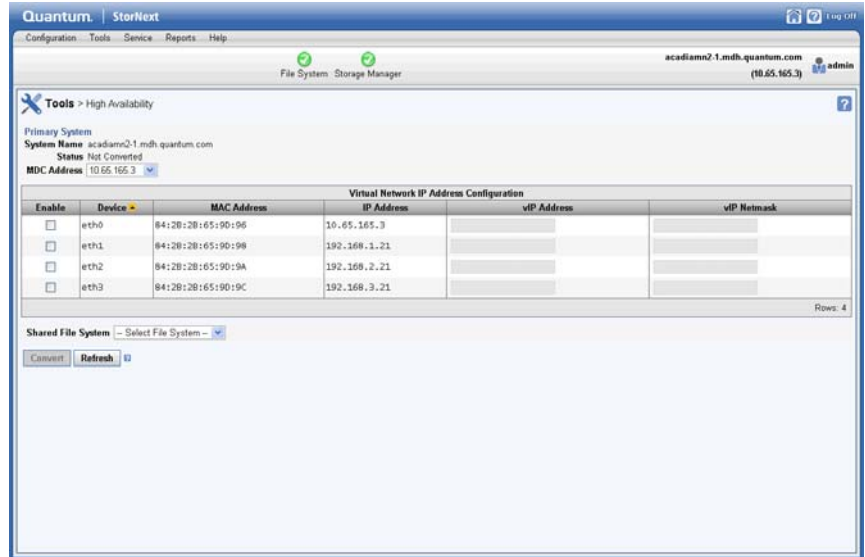
- The conversion process converts one node at a time. The second node should be converted as soon as possible after the first node.
- StorNext operating files will be moved to the HaShared file system, and this move cannot easily be reversed.
- Following conversion, the Primary node is identified by the vIP for Replication/Deduplication.
- Replication/Deduplication policies must be changed to use the vIP:
 - The global policy for each file system must use it as the “Address for Replication and Deduplication”
 - Replication policies must use it as the Target for Replication Schedules
 - If multilink is configured, the vIP address should be used.
- **Prior to the conversion you MUST verify that users and groups created by StorNext on both MDCs have the same (numeric) uid/gid. If they do not have the same uid/gid, permission problems will occur after converting to HA.**

HA Conversion Procedure

Follow these steps to configure HA:

- 1 Choose **High Availability > Convert** from the **Tools** menu. The **Tools >High Availability** screen appears.

Figure 109 Tools > HA Screen



- At the **Shared File System** field, select the shared file system you want to convert to HA.

Caution: Once you convert a file system to HA you cannot undo the operation, so choose carefully when selecting the file system.

- At the **MDC Address** field, select one IP address to be placed in the `ha_peer` file for use in administrative operations between the MDCs in the HA Cluster.
- If your HA cluster also runs the blockpool, select **Enable** and then enter the virtual IP address and virtual netmask. (Ask your network administrator for the vIP address and netmask.)
- Click **Convert** to convert the primary node to HA.
- Enter the IP address of the secondary system on the same LAN, and then click **Scan**. If you do not already have licenses for the secondary system in the license file, you will be required to switch to the license page to import them before continuing. (The information comes from the individual `license.dat` files on both MDCs. StorNext merges the two into one file before converting the secondary.)

Note: Until you have performed the scan, you cannot import the license file for the secondary system using StorNext's import function. After you have performed the scan you can import licenses for the secondary. Following the conversion to HA, the license file will contain both primary and secondary licenses and will be present on both servers.

- 7 Click **Convert** to convert the secondary system.

Managing HA


The StorNext Manage HA screen is used to monitor the current statuses of your primary and secondary servers.

The screen includes **Enter Config Mode** and **Exit Config Mode** buttons to place the HA Cluster in a state that allows the Primary MDC to restart CVFS and individual FSMs without incurring an HA Reset, failover of any file systems, or transfer of Primary status to the peer MDC. This is required for making configuration changes to the HaShared file system through the GUI.

Follow these steps to lock the HA cluster and enter Config mode, and subsequently to exit Config mode:

- 1 Choose **High Availability > Manage** from the **Tools** menu. The **Manage High Availability** screen appears.

Figure 110 Manage HA Screen



- 2 Click **Enter Config Mode**.
- 3 When the confirmation message appears, click **Yes** to proceed or **No** to abort.
- 4 Click **OK** when a message informs you that the HA cluster was successfully locked.
- 5 When you are ready to unlock the cluster and exit Config mode, click **Exit Config Mode**. All file systems will be stopped on the primary MDC and then restarted on both MDCs.
- 6 When the confirmation message appears, click **Yes** to proceed or **No** to abort.
- 7 Click **OK** when a message informs you that the HA cluster was successfully unlocked.

HA Statuses and Reporting

StorNext does not currently have an HA report, but you can use the HA log to track activity related to the nodes in your HA system, such as whether a node is currently active, whether synchronization has completed, and so on.

Troubleshooting HA

The Troubleshooting appendix in this guide contains simple troubleshooting procedures pertaining to HA. For more information, see [Troubleshooting](#).



Chapter 10

StorNext Reports


The Reports Menu contains the following options:


- [StorNext Logs](#): Access logs of StorNext operations
- [The Jobs Report](#): View a list of pending and completed jobs on the system
- [The Files Report](#): View information about specific files, such as the owner, group, policy class, permissions, and copy information
- [The Drives Reports](#): View information about the drives in your libraries, including the serial number and current state and status
- [The Media Report](#): View information pertaining to the media in selected libraries or all libraries, including the media ID, type and class, as well as the number of files and the last access date and time
- [The Relations Report](#): View the names of the policy classes which correspond to the managed directories in your system
- [The File Systems Report](#): View file system statistics including active clients, space, size, disks, and stripe groups
- [The SAN Devices Report](#): View information about devices currently attached to your SAN, including disks/LUNs, libraries, and tape drives
- [The Tape Consolidation Report](#): View statistics on the tape consolidation (defragmenting) process


- [The SAN and LAN Clients Report](#): View statistics for StorNext clients, including the number of connected clients and distributed LAN clients, and client performance
- [The LAN Client Performance Report](#): View information about distributed LAN clients and servers, including read and write speed
- [Replication Deduplication Reports](#)
 - [Policy Activity Report](#): View replication and deduplication performance statistics
 - [Policy Summary Report](#): View replication and deduplication information for each policy
- [The Distributed Data Mover Report](#): View activity related to the Distributed Data Mover feature.


Report Navigation Controls


If a log or report spans more than one screen, navigation controls at the bottom of the screen allow you to select a page by number, or to view one of these pages.





Click  to go to the first page

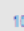





Click  to skip backwards ten pages

Click  to go to previous page

Click  to go to the next page

Click  to skip ahead ten pages

Click  to go to the last page

Click a specific page number          to go to that page

StorNext Logs

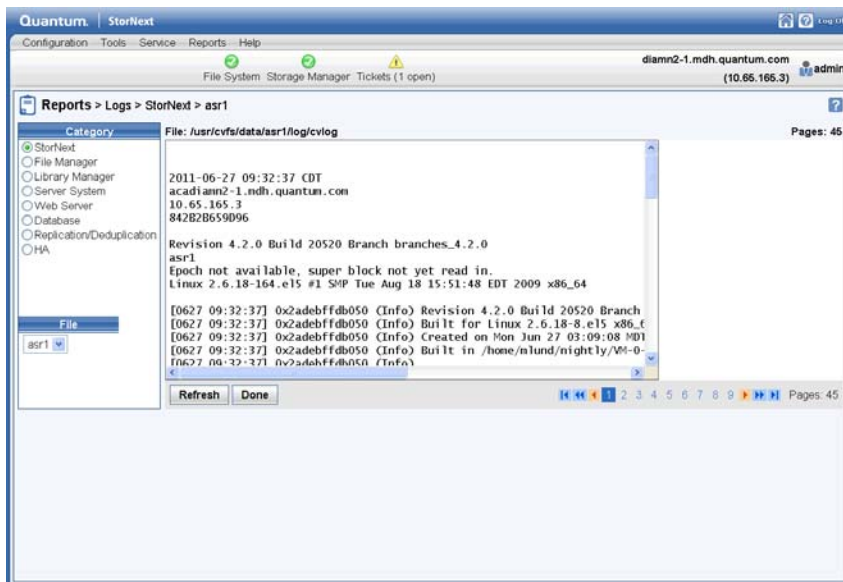
Report menu options enable you to access and view any of the following types of logs:

- **StorNext Logs:** Logs about each configured file system.
- **File Manager Logs:** Logs that track storage errors, etc. of the Storage Manager.
- **Library Manager Logs:** Logs that track library events and status
- **Server Logs:** Logs that record system messages.
- **StorNext Web Server Logs:** Various logs related to the web server.
- **StorNext Database Logs:** Logs that track changes to the internal database.
- **Replication/Deduplication Logs:** Logs that track the progress and status of data replication or data deduplication operations.
- **HA Logs:** Logs that track activity related to the nodes in your HA system, such as whether a node is currently active, whether synchronization has completed, and so on.

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 Select **Logs** from the **Reports** menu. The **Reports > Logs** screen appears.

Figure 111 Reports > Logs
Screen



- 2 On the left side of the screen, select the type of log you wish to view.
- 3 If desired, select a file system different from the default one shown beneath the log categories.

The log you selected automatically appears. (If the log does not appear, click **Refresh**.) If the log spans more than one screen, use the navigation controls at the bottom of the screen as described in [Report Navigation Controls](#) on page 228.

The Jobs Report

The Jobs Report provides information about previously run jobs on your file systems. Jobs include all actions performed for file systems, such as make, stop, start, check, and so on. Use the navigation controls at the bottom of the screen if there are multiple screens of jobs.

Use the following procedure to run the Jobs Report.

- 1 Choose **Jobs** from the **Reports** menu. The **Reports > Jobs** report appears.

Figure 112 Jobs Report

ID	Job	Attributes	User	Start Time	End Time	Status
51	RAS Tickets Close All		admin	2010-07-16 15:59:26 CDT	2010-07-16 15:59:26 CDT	Success
50	Media Dismount		admin	2010-07-16 15:51:20 CDT	2010-07-16 15:51:20 CDT	Failure
49	Tape Drive Modify	2 ON		2010-07-16 15:44:03 CDT	2010-07-16 15:44:04 CDT	Success
48	Tape Drive Modify	1 ON		2010-07-16 15:44:03 CDT	2010-07-16 15:44:03 CDT	Success
47	Media Dismount		admin	2010-07-16 15:43:05 CDT	2010-07-16 15:43:07 CDT	Failure
46	Tape Drive Modify	2 ON		2010-07-16 15:42:32 CDT	2010-07-16 15:42:32 CDT	Success
45	Tape Drive Modify	1 ON		2010-07-16 15:42:11 CDT	2010-07-16 15:42:11 CDT	Success
44	Media Dismount		admin	2010-07-16 15:40:41 CDT	2010-07-16 15:41:18 CDT	Success
43	Media Dismount		admin	2010-07-16 15:39:38 CDT	2010-07-16 15:40:15 CDT	Success
42	Store Files	1000	admin	2010-07-16 15:12:51 CDT	2010-07-16 15:21:35 CDT	Failure

The Jobs Report includes the following information:

- **ID:** The job ID number.
- **Job:** The job name assigned by StorNext for the type of action performed (for example, "FileSystem Make").
- **Attributes:** The name of the related file system, mount point, policy, etc. on which the job was performed. For example, if the job was to start the file system, the name of that file system appears in the Attributes column.
- **User:** The logged in user who initiated the job.
- **Start and End Time:** The times the job was started and ended.
- **Status:** The job's final or current status, such as Success or Failure.

Viewing Detailed Job Information

To view detailed information about a specific job, select the desired job and then click **View** to see the information on a new screen. When you are finished viewing that job's information, click **Done**.

Filter Options

The **Status Filter** allows you to refine the displayed list of jobs according to Success, Failure, Warning, Working, Unknown, or All. Choose one of these criteria to restrict the displayed list of jobs to your selection. After you select a Status Filter option, click **Refresh** to resort and view the jobs list with your selected criteria.

The **Type Filter** works either together or separately from the Status Filter. The Type Filter allows you to refine the displayed list of jobs according to a specific job action:

All	Tape Drive Modify	File System Start	File System Metadump
Unknown	Media Add	File System Stop	Cancel Request
Policy Add	Media Delete	File System Check	System Service Start
Policy Delete	Media Modify	File System Rename	System Service Stop
Policy Modify	Media Move	File System Expand	Convert to HA
Run Store Policy	Media Remove	File System Scan For New Storage	Store Files
Schedule Add	Media Eject	RAS Ticket Close	Change File Version
Schedule Delete	Media Purge	RAS Ticket Close All	Recover Files
Schedule Modify	Media Mount	RAS Ticket Analysis Update	Recover Directory
Schedule Reset	Media Dismount	Email Server Add	Retrieve Files
Storage Disk Add	Media Reclassify	Email Server Delete	Retrieve Directory
Storage Disk Delete	Media Assign to Policy Class	Email Server Modify	Truncate Files
Storage Disk Modify	Media Transcribe	Email Notification Add	Move Files

File System Add	Media State Change	Email Notification Delete	Modify File Attributes
File System Delete	Media Clean by Media ID	Email Notification Modify	Health Check
File System Modify	Media Clean by File System	NAS Share Add	Capture State
File System Move Stripe Groups	Media Clean by Policy Class	NAS Share Delete	Add Drive Pool
Library Add	Media Import Mailbox	CIFS Shares Delete All	Delete Drive Pool
Library Delete	Media Import Bulk Load	CIFS Windows Domain Join	Add Email Contact
Library Modify	File System Make	CIFS Windows Domain Leave	Modify Email Contact
Tape Drive Add	File System Mount	CIFS Workgroup User Add	Delete Email Contact
Tape Drive Delete	File System Unmount	System Date/Time Modify	System Backup

Exiting the Jobs Report screen.

When you finished viewing the Jobs Report, click **Done**.


The Files Report

The Files Report provides general information about selected files, as well as specific details if you require more granular information.

Use the following procedure to run the Files Report.

- 1 Choose **Files** from the **Reports** menu. The **Reports > Files** report appears.


Figure 113 Files Report



The **Report > Files** screen shows the following information about the files in your system:

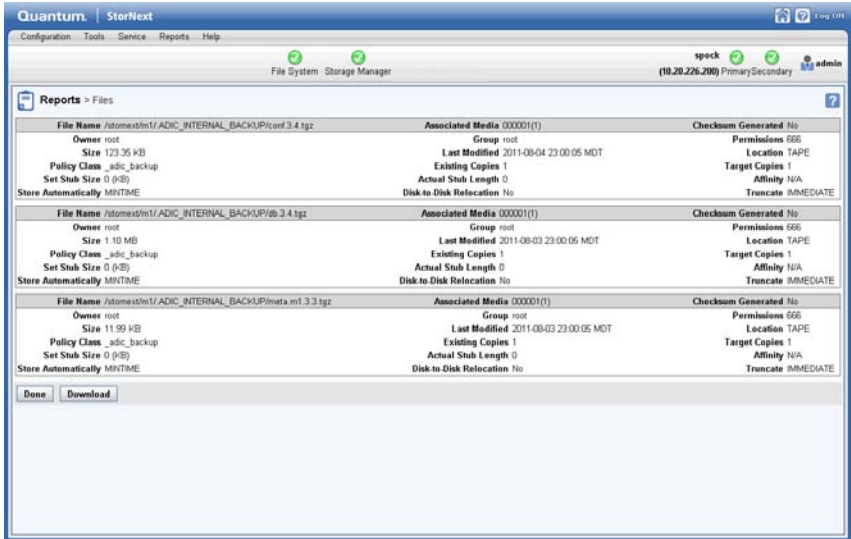
- **Name:** The name of the file
 - **Size:** The current size of the file
 - **Last Modified:** The date when the file's contents were last modified
- 2 To locate another file, click **Browse** to display the **StorNext File Browser**.

Figure 114 StorNext File
Browser



- 3 Do one of the following:
 - Check the box to the left of the desired folder (directory) to select all files in the folder
 - Click the folder name and then select files individually. (Hold down the **Shift** key and click to select contiguous files, or hold down the **Control** key and click to select multiple non-contiguous files.)
- 4 Click **Continue** to proceed and return to the **Reports > Files** screen.
- 5 Click **File Info** to view detailed information for the files you selected. The **File Info** screen appears.

Figure 115 File Info Screen



- 6 To download the report, click **Download**.
- 7 When you finished viewing report information, click **Done**.

The Drives Reports

The Drives Report provides a list of drives in your system and enables you to view details about selected drives.

Use the following procedure to run the Drives Report.

- 1 Choose **Drives** from the **Reports** menu. The **Reports > Drives** report appears.

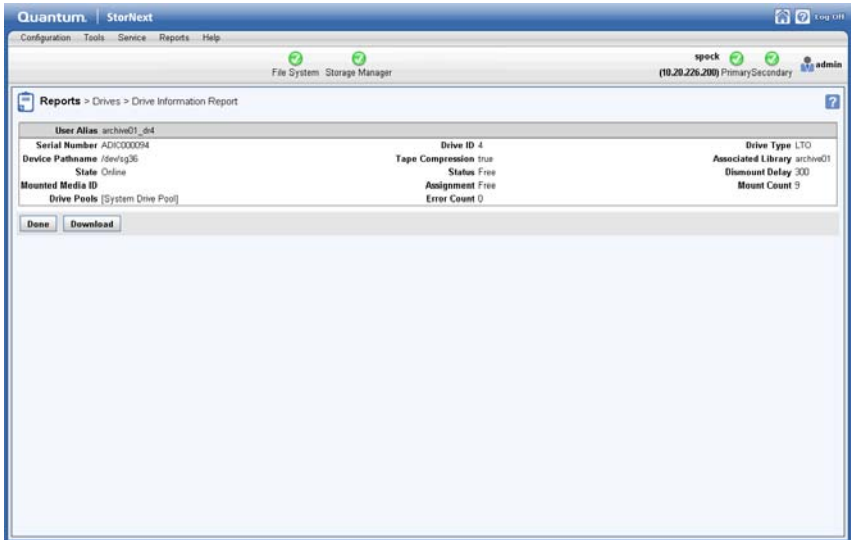
Figure 116 Drives Report

Serial Number	State	Status	User Alias	Mounted Media	Dismount Delay	Compression
ADIC000094	Online	Free	archive01_dr4		300	true
ADIC000095	Online	Mounted	archive01_dr2	000004	300	true
ADIC000092	Online	Free	archive01_dr5		300	true
ADIC000093	Online	Free	archive01_dr3		300	true

The **Report > Drives** screen shows the following information about your drives:

- **Serial Number:** The drive’s serial number
 - **State:** The current state of the drive, such as Online or Offline
 - **Status:** The drive’s current status, such as Free or Mounted
 - **User Alias:** The user identifier for the drive
 - **Mounted Media:** The mounted media number
 - **Dismount Delay:** The interval of time before the drive is dismounted
 - **Compression:** Specifies whether compression is enabled (True) or disabled (False)
- 2 To view information for one or more specific drives, select the desired drive and then click **View Drive Information Report**. The **Drives > Drive Information Report** appears.

Figure 117 Drive Information Report



- 3 To download the report, click **Download**.
- 4 When you finished viewing report information, click **Done**.

The File Systems Report

The File Systems Report provides a list of parameters and statistics about configured StorNext file systems.

Use the following procedure to run the File System Report.

- 1 Choose **File Systems** from the **Reports** menu. The **Reports > File Systems** report appears.

Figure 118 File Systems Report

File System: pool_target		Mounted on: /stornext/pool_target		Status:
Stripe Group 1: sg0	Content: MJ	Status: Up	Read: Enabled	Write: Enabled
Total Space 99.98 GB	Depth 1	Realtime IO Limit 0/sec	Non-realtime IO Reserve 0/sec	LUNs HP_EVA_0011
Reserved Space 0 B	Breadth 4 MB	Realtime Bandwidth Limit 0/sec	Non-realtime Clients 0	
Free Space 99.92 GB	Multipath Method Rotate	Realtime IO Commit 0/sec	Non-realtime IO Hint 0/sec	
Stripe Group 2: sg1	Content: U	Status: Up	Read: Enabled	Write: Enabled
Total Space 99.99 GB	Depth 1	Realtime IO Limit 0/sec	Non-realtime IO Reserve 0/sec	LUNs HP_EVA_0012
Reserved Space 4.13 GB	Breadth 2 MB	Realtime Bandwidth Limit 0/sec	Non-realtime Clients 0	
Free Space 96.21 GB	Multipath Method Rotate	Realtime IO Commit 0/sec	Non-realtime IO Hint 0/sec	
File System: target		Mounted on: /stornext/target		Status:
Stripe Group 1: sg0	Content: MJ	Status: Up	Read: Enabled	Write: Enabled
Total Space 99.98 GB	Depth 1	Realtime IO Limit 0/sec	Non-realtime IO Reserve 0/sec	LUNs HP_EVA_0008
Reserved Space 0 B	Breadth 4 MB	Realtime Bandwidth Limit 0/sec	Non-realtime Clients 0	
Free Space 99.42 GB	Multipath Method Rotate	Realtime IO Commit 0/sec	Non-realtime IO Hint 0/sec	
Stripe Group 2: sg1	Content: U	Status: Up	Read: Enabled	Write: Enabled
Total Space 199.99 GB	Depth 2	Realtime IO Limit 0/sec	Non-realtime IO Reserve 0/sec	LUNs HP_EVA_0009
Reserved Space 4.13 GB	Breadth 2 MB	Realtime Bandwidth Limit 0/sec	Non-realtime Clients 0	HP_EVA_0010
Free Space 196.81 GB	Multipath Method Rotate	Realtime IO Commit 0/sec	Non-realtime IO Hint 0/sec	

Refresh Done Rows: 2

The File Systems Report provides the following information about your file systems:

- **File System Name:** The name of the file system
- **Mount Point ("Mounted on"):** The file system's mount point location
- **Status:** The file system's current status, indicated a green check mark icon (Active), a yellow exclamation mark icon (Warning), or a red X icon (Stopped).

2 When you finished viewing report information, click **Done**.

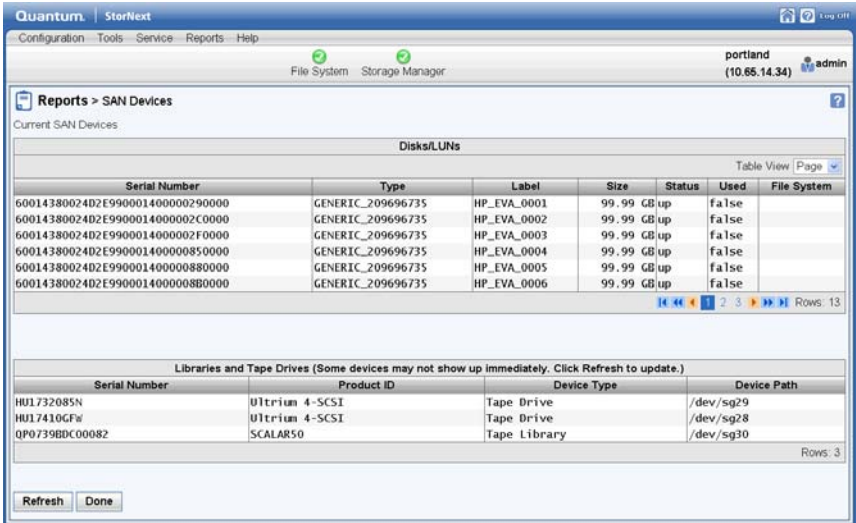
The SAN Devices Report

The SAN Devices Report shows a list of details for all currently configured devices attached to your SAN.

Use the following procedure to run the SAN Devices Report.

- 1 Choose **SAN Devices** from the **Reports** menu. The **Reports > SAN Devices** report appears.

Figure 119 SAN Devices Report



The SAN Devices Report provides the following information:

- **Disks and LUNs**
 - **Serial Number:** The disk's or LUN's serial number or path name.
 - **Type:** The device type.
 - **Label:** The label, if any, assigned to the device.
 - **Size:** The total capacity for the device.
 - **Status:** The device's current status. Statuses include:
 - **Used:** Indicates whether the device is currently in use (true or false).
 - **File System:** The name of the file system with which the device is associated.
- **Libraries and Tapes Drives**
 - **Serial Number:** The serial number of the library or tape drive.
 - **Product ID:** The model number or product name of the library or tape drive.
 - **Device Type:** The type of device: Tape Library or Tape Drive.

- **Device Path:** The path name for the device.
- 2 If desired, click **Refresh** to manually update (refresh) the report data.

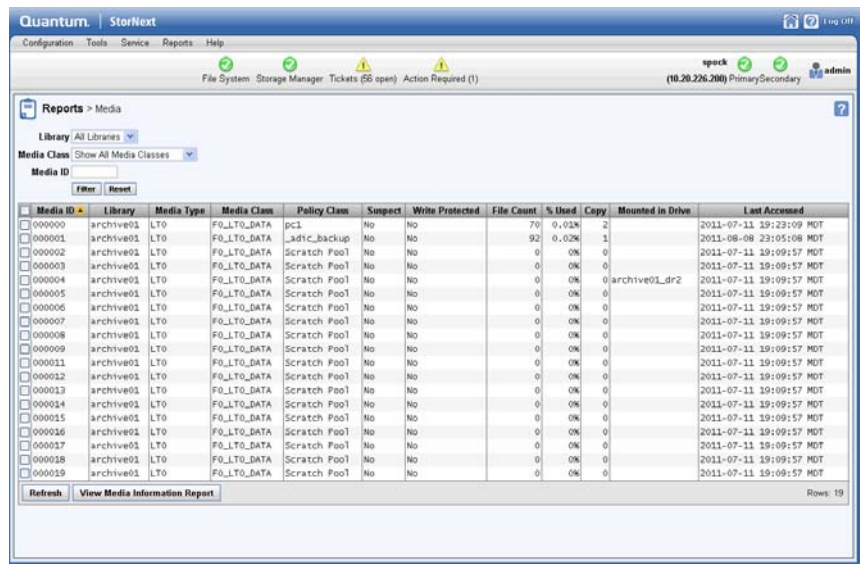
The Media Report

The Media Report shows a list of details for all media in a selected library or all libraries.

Use the following procedure to run the Media Report.

- 1 Choose **Media** from the **Reports** menu. The **Reports > Media** report appears.

Figure 120 Media Report



- 2 Select from the **Library** pulldown list one of these options:
- **All Libraries:** Select this to view information for all media in all libraries

- **A selected library:** Select a specific library whose media information you want to view
- 3 Select from the **Media Class** pulldown list one of these options:
- **Show All Media Classes:** Display information for all media in all classes
 - **Show Data Media Class:** Display only information for media available to be used for data migration that are not considered full
 - **Show Migrate Media Class:** Display only information for media used for data migration which are considered full
 - **Show Blank Media Class:** Display only information for blank media
 - **Show Clean Media Class:** Display only information for cleaning media
 - **Show Backup Media Class:** Display only information for media used for backup purposes
- 4 If desired, enter one or more characters at the **Media ID** field to restrict the list of media to those whose IDs contain the character(s).
- 5 Click **Filter** to apply the filtering options.
- 6 To update the information displayed, you can click **Refresh** at any time.

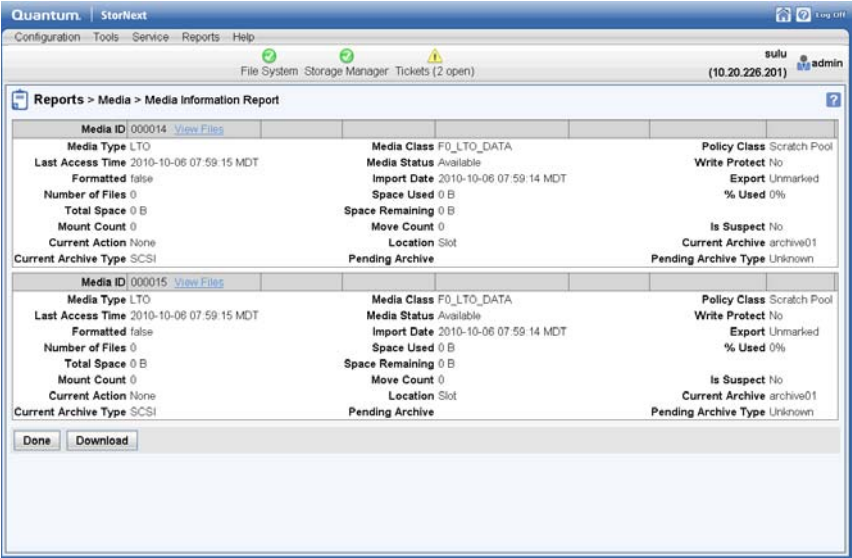
The Media Report provides the following information for each piece of media that meets the criteria you selected from the **Library**, **Media Class** and **Media ID** fields:

- **Media ID:** The unique identifier for the media
- **Library:** The name of the library in which the media currently resides
- **Media Type:** The type of media
- **Media Class:** The media class to which the media belongs
- **Policy Class:** The policy class to which the media belongs
- **Suspect:** Indicates whether the media is considered suspect (possibly unreliable or defective)
- **Write Protected:** Indicates whether the media is write protected

- **File Count:** The number of files saved on the media
 - **% Used:** Indicates the percentage of the media which is currently used
 - **Copy:** Indicates the policy class copy number on the media
 - **Mounted in Drive:** Indicates whether the media is currently mounted in a drive
 - **Last Accessed:** Indicates the date and time when the media was last accessed
- 7 To view a report for a particular piece of media, select the desired media from the list. To select multiple media, hold down the Control key while you click additional media. To select all media, click the checkbox to the left of the **Name** heading.

After you have selected media, click **View Media Information Report**. This report allows you to see all files on the selected media. The report looks similar to the one shown in [Figure 121](#).

Figure 121 Media Information Report



- 8 If desired, you can save the report output as a CSV file (Microsoft Excel format) by clicking **Download**.
- 9 When you are finished viewing the information report, click **Done**.


The Relations Report

The Relations Report shows the pathname of the managed file system’s directory and the corresponding policy class name.

Use the following procedure to run the Relations Report.

- 1 Choose **Relations** from the **Reports** menu. The **Reports > Relations** report appears.

Figure 122 Relations Report



- 2 When you are finished reviewing the report output, click **Done**.

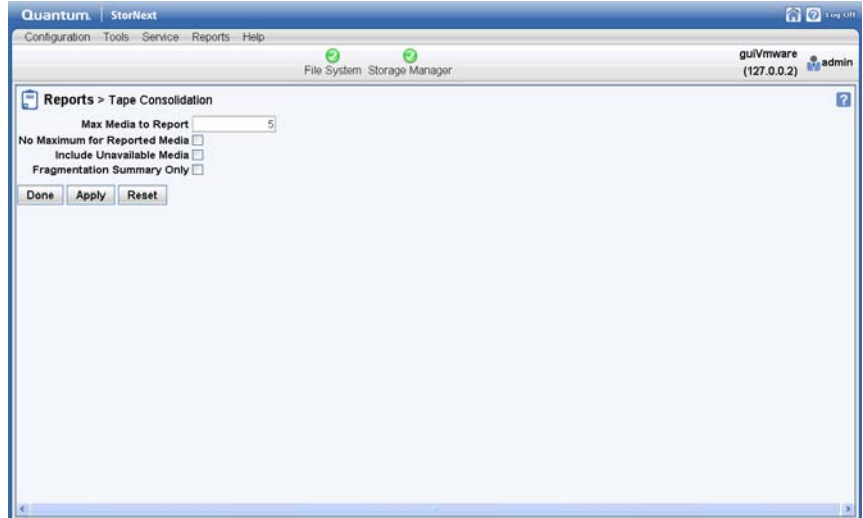
The Tape Consolidation Report

The Tape Consolidation Report shows information about the tape consolidation process, also known as defragmentation.

Use the following procedure to run the Tape Consolidation Report.

- 1 Choose **Tape Consolidation** from the **Reports** menu. The **Reports > Tape Consolidation** report appears.

Figure 123 Tape Consolidation Report



- 2 Enter the following fields which determine report parameters:
 - **Max Media to Report:** The maximum number of media included in the consolidation process
 - **No Maximum for Reported Media:** Indicate that there is no limit for the number of media included in the report
 - **Include Unavailable Media:** Specify whether to include currently unavailable media in the report
 - **Fragmentation Summary Only:** Specify whether to provide only a high-level summary of consolidation results
- 3 Click **Apply** to save and apply the report parameters you just entered.
- 4 When you are finished, click **Done** to generate the report.
- 5 The report is generated as a job. To view the report output, choose **Jobs** from the **Reports** menu and then select a Tape Consolidation Report job.
- 6 When you are finished reviewing the report output, click **Done**.

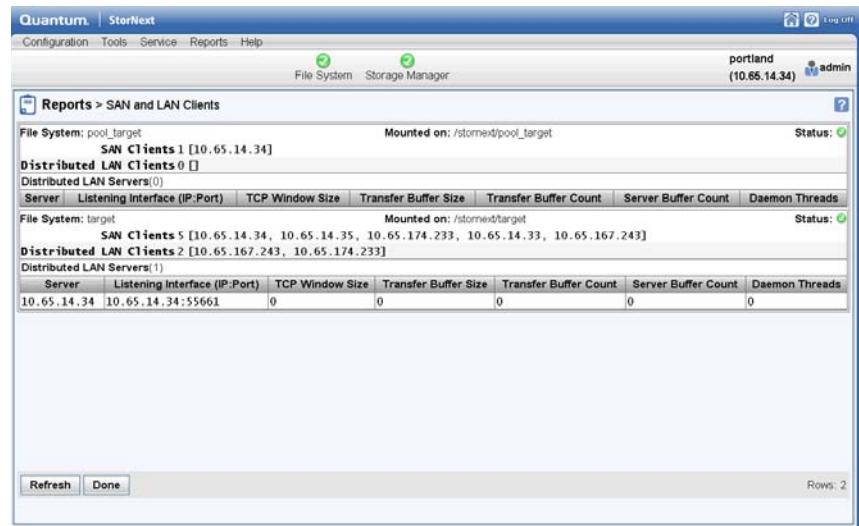
The SAN and LAN Clients Report

The SAN and LAN Clients Report provides statistics for StorNext clients, including the number of StorNext SAN clients and distributed LAN clients, and client performance.

Use the following procedure to run the SAN and LAN Clients Report.

- 1 Choose **SAN and LAN Clients** from the **Reports** menu. The **Reports > SAN and LAN Clients** report appears.

Figure 124 SAN and LAN Clients Report



The SAN and LAN Client 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.
- **Status:** The file system's current status (Normal, Error, or Warning)
- **SAN Clients:** The total number of physically connected StorNext SAN clients, and the IP address of the current client.
- **Distributed LAN Clients:** The total number of StorNext distributed LAN clients.

- **Distributed LAN Servers:** The total number of distributed LAN servers for the file system.
 - **Server:** The names of the distributed LAN servers.
 - **Distributed LAN Clients:** The names of distributed LAN clients.
 - **Listening Interface (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)
 - **Server Buffer Count:**
 - **Daemon Threads:** The maximum number of daemon threads used by the distributed LAN server. (Default: 8)
- 2 If desired, click **Refresh** to manually update (refresh) the report data.
 - 3 Click **Done** when you are finished viewing the report.

The LAN Client Performance Report

The LAN Client Performance Report provides information about distributed LAN clients, including read and write speed.

Use the following procedure to run the LAN Client Performance Report.

- 1 Choose **LAN Client Performance** from the **Reports** menu. The **Reports > LAN Client Performance** report appears.

Figure 125 LAN Client Performance Report

File System	Server	Client	Client Interface	Read Bytes/Sec	Write Bytes/Sec
FourLUN	10.65.187.19	10.65.176.62	10.65.176.62:53325	0 B	0 B
FourLUN	10.65.187.19	10.65.176.62	192.168.200.1:33383	0 B	0 B
FourLUN	10.65.187.19	10.65.186.252	10.65.186.252:43074	0 B	0 B
FourLUN	10.65.187.19	10.65.186.252	192.168.200.2:33753	0 B	0 B
FourLUN	10.65.180.117	10.65.176.62	10.65.176.62:55111	0 B	0 B
FourLUN	10.65.180.117	10.65.176.62	192.168.200.1:50943	0 B	0 B
FourLUN	10.65.180.117	10.65.186.252	10.65.186.252:44108	0 B	0 B
FourLUN	10.65.180.117	10.65.186.252	192.168.200.2:55737	0 B	0 B
SingTelLUN	10.65.187.19	10.65.176.62	10.65.176.62:47608	0 B	0 B
SingTelLUN	10.65.187.19	10.65.176.62	192.168.200.1:54677	0 B	0 B

The LAN Client Performance Report provides the following information:

- **File System:** The name of the file system supporting the clients.
 - **Server:** The name of the distributed LAN server on the indicated file system.
 - **Client:** The IP address for the corresponding client interface listed in the **Client Interface** column.
 - **Client Interface:** The name of the distributed LAN client for the indicated file system and distributed LAN server.
 - **Read Bytes:** The number of bytes read by the distributed LAN client.
 - **Write Bytes:** The number of bytes written by the distributed LAN client.
- 2 If desired, click **Refresh** to manually update (refresh) the report data.
- 3 Click **Done** when you are finished viewing the report.

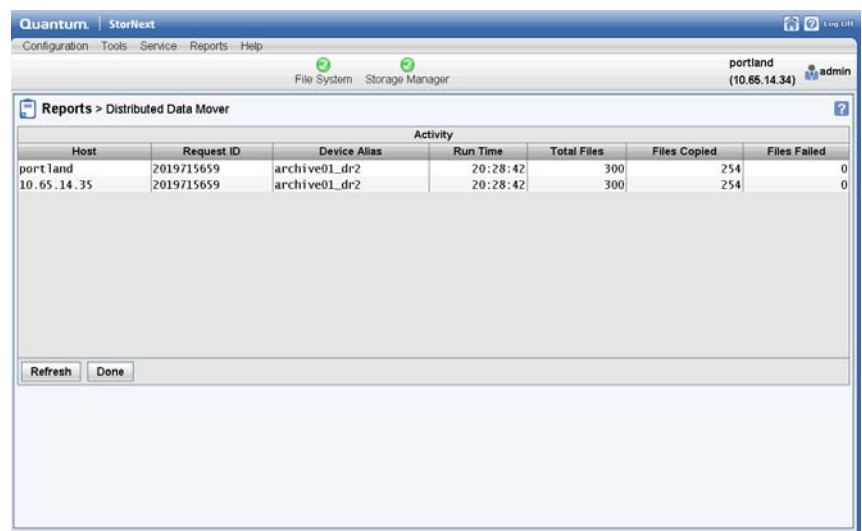
The Distributed Data Mover Report

The Distributed Data Mover Report shows a list of details pertaining to the Distributed Data Mover feature.

Use the following procedure to run the Distributed Data Mover Report.

- 1 Choose **Distributed Data Mover** from the **Reports** menu. The **Distributed Data Mover Report** appears.

Figure 126 Distributed Data Mover Report



The Distributed Data Mover Report provides the following information:

- **Host:** The name of the machine on which the source data resides.
- **Request ID:** The identification number for the move-data request.
- **Device Alias:** The alias of the destination device to which data is moved.
- **Run Time:** The time the data movement occurred.
- **Total Files:** The total number of files moved.

- **Files Copied:** The total number of files copied in the move-data process.
 - **Files Failed:** The number of files that were not moved during the move-data process
- 2 If desired, click **Refresh** to manually update (refresh) the report data.
 - 3 When you are finished viewing report information, click **Done**.

Replication Deduplication Reports

StorNext provides these reports that contain information pertaining to replication and deduplication:

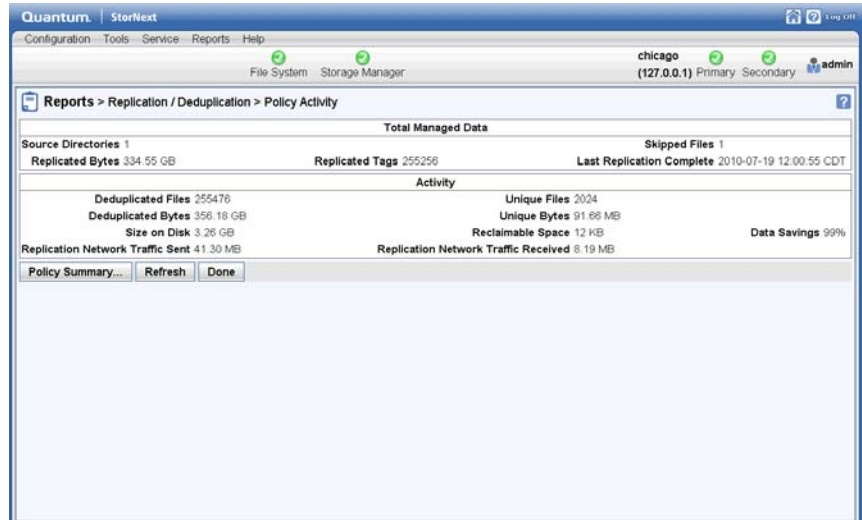
- **Policy Activity:** This report shows statistics related to replication and deduplication. This report also provides statistics such as savings realized by replication, and the current progress of ongoing replications.
- **Policy Summary:** This report shows information about replication storage policies created to support the data replication and deduplication processes. This report also provides statistics such as savings realized by replication, and the current progress of ongoing replications.

Policy Activity Report

Use the following procedure to run the Replication/ Deduplication Policy Report.

- 1 Choose **Replication/ Deduplication > Policy Activity** from the **Reports** menu. The **Reports > Replication/ Deduplication > Policy Activity** report appears.

Figure 127 Replication/
Deduplication Policy Activity
Report



The Replication/ Deduplication Policy Activity Report provides the following information:

- **Total Managed Data**
 - **Source Directories:** The number of source directories from which StorNext checked for data to replicate.
 - **Replicated Bytes:** The number of bytes of data replicated.
 - **Replicated Tags:** The number of replication tags applied to files.
 - **Skipped Files:** The number of files not included in the replication process.
 - **Last Replication Complete:** The date and time the last replication was finished.
- **Deduplication Activity**
 - **Deduplicated Files:** The number of files deduplicated.
 - **Deduplicated Bytes:** The number of bytes of data deduplicated.
 - **Size on Disk:** The size (in bytes) of deduplicated data on disk.
 - **Network Sent:** The number of bytes of deduplicated data sent over the network.

- **Unique Files:** The number of unique files included in the deduplication process.
 - **Unique Bytes:** The number of unique bytes included in the deduplication process.
 - **Reclaimable Space:** The amount of reclaimable space realized by data deduplication.
 - **Network Received:** The number of bytes of deduplicated data received over the network.
 - **Data Savings:** The percentage of data savings realized by data deduplication.
- 2 If desired, click **Refresh** to manually update (refresh) the report data.
 - 3 To view a report showing replication policy summary information, click **Policy Summary**.

Policy Summary Report

Use the following procedure to run the Replication/ Deduplication Policy Summary Report.

- 1 Choose **Replication/ Deduplication > Policy Summary** from the **Reports** menu. The **Reports > Replication/ Deduplication > Policy Summary** report appears.

Figure 128 Replication/
Deduplication Policy Summary
Report

Quantum | StorNext |

Configuration Tools Service Reports Help

File System Storage Manager

chicago (127.0.0.1) Primary Secondary admin

Reports > Replication / Deduplication > Policy > Summary

Running Policy	File System	Source Directory	Replicated				Replication			Replication Network Traffic	
			Directories	Files	Skipped	Bytes	Average Rate	Estimated Completion	Last Completed	Sent	Received
repl	source	/stornext /source/rep1	256	255256		1 334.55 GB	N/A	N/A	2010-07-19 12:00:55 CDT	41.30 MB	8.19 MB

Details...

Completion Report...

Policy Activity...

Refresh

Done

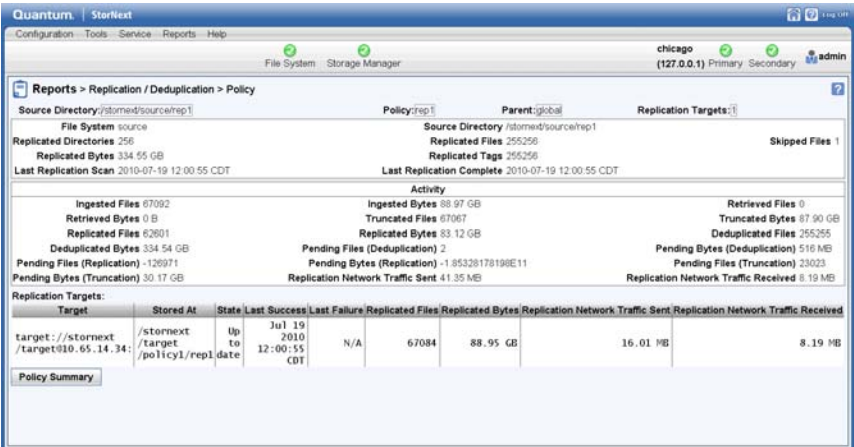
Rows: 1

The Replication/ Deduplication Policy Summary Report provides the following information:

- **Policy:** The name of the replication storage policy.
- **File System:** The name of the file system for which replication is enabled.
- **Source Directory:** The name of the source directory from which information is replicated.
- **Replicated**
 - **Directories:** The number of replicated directories to date.
 - **Files:** The number of replicated files to date.
 - **Skipped:** The number of files skipped by the replication process.
 - **Bytes:** The total number of data bytes replicated to date.
- **Replication**
 - **Average Rate:** The approximate rate at which data was replicated from the source to the target.
 - **Estimated Completion:** The estimated time replication is currently scheduled to complete.
 - **Last Completed:** The date and time the last replication was finished.

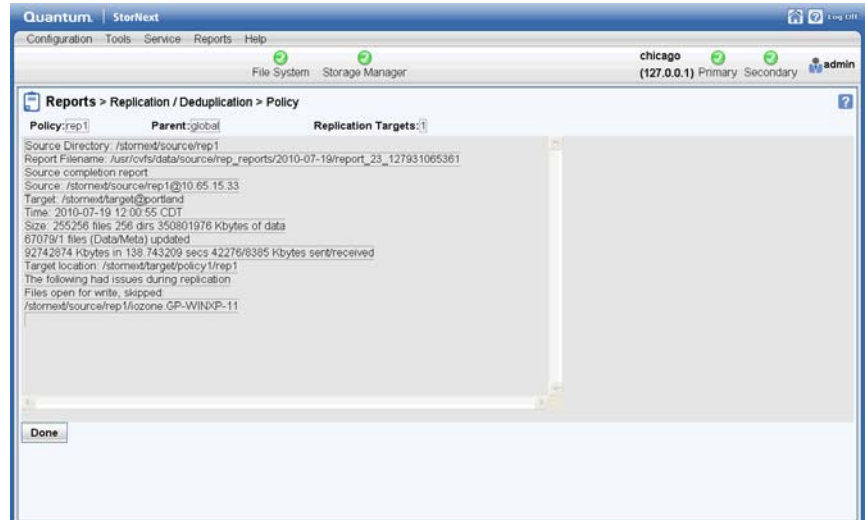
- **Network**
 - **Sent:** The amount of replicated data sent from the source.
 - **Received:** The amount of replicated data received on the target.
- 2 To update report information, click **Refresh**.
- 3 To view details for a particular policy, select the desired policy and then click **Details**.

Figure 129 Replication/
Deduplication Policy Details
Report



- 4 To view a report showing completed replication, click **Completion Report**.
- 5 When you are finished viewing this report, click **Done**.

Figure 130 Replication/
Deduplication Policy
Completion Report



6 To view the Policy Activity report, click **Policy Activity**.



Chapter 11

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.

Linux Configuration File Format

Beginning with StorNext 4.0, the configuration file for Linux is now in XML format. The Linux configuration file is now identified by a `.cfgx` extension rather than `.cfg` for Windows systems.

There are some differences between the XML and `.cfg` formats. Specifically, the Reserved Space parameter is called `reservedSpace` in the XML format, and its value must be either true or false.

For additional information and examples of both configuration file formats, see the *StorNext Installation Guide*.

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.** On Linux systems, multiple Ethernet interfaces may be configured as a single bond interface using the Linux bonding driver. The bond interface may be then be configured for use by the StorNext Distributed LAN server. In this case a Distributed LAN client may have only a single Ethernet interface. LAN clients running Linux may also be configured to use a bond interface. 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.

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.

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`.

Note: In the Windows XML format configuration file, the FS block size parameter is called `fsBlockSize`. Regardless of the difference in parameter names (`fsBlockSize` and `FsBlockSize`) used in the Windows and UNIX configuration files, the requirements are identical for Windows and UNIX systems.

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. Existing file systems managed by StorNext Storage Manager **MUST** have a journal size of at least 64M. The TSM portion of SNSM may not start if the journal size for a managed file system is less than 64M.

Note: In the Windows XML format configuration file, the journal size parameter is called `journalSize`. Regardless of the difference in parameter names (`journalSize` and `JournalSize`) used in the Windows and UNIX configuration files, the requirements are identical for Windows and UNIX systems.

The following table shows recommended settings. Choose the setting that corresponds to your configuration.

<code>FsBlockSize</code>	<code>JournalSize</code>
16KB	16MB
64KB	64MB

This setting is adjustable using the `cvupdatefs` utility. For more information, see the `cvupdatefs` man page.

Example: `JournalSize 16M`

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.

Changing StorNext's Default Session Timeout Interval

By default, StorNext automatically logs out the current user after thirty minutes of inactivity. Follow the procedure below to change the timeout interval. (All steps must be performed from the command line.)

- 1 Stop StorNext by entering the following:

```
service stornext_web stop
```

- 2 Edit the config file /usr/adic/tomcat/webapps/ROOT/WEB-INF/web.xml by entering the following:

```
<!-- Set session timeout to 30 minutes -->

<session-config>
    <session-timeout>30</session-timeout>
</session-config>
```

- 3 Restart the StorNext GUI by entering the following:

```
service stornext_web start
```

For HA, these steps must be performed on both servers.

Caution: These changes alter the StorNext configuration file, so you should exercise caution. If you have any doubts about your ability to manually change the configuration file as described, do not attempt this procedure unless you have assistance from Quantum Technical Support.

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	Feature or Category	Description
AIX	Clients	<p>Clients on AIX systems may not unmount a file system after running the <code>fsstress</code> command against that file system. The client must then be rebooted to release the mount.</p> <p>This issue occurs on AIX systems when the <code>fsstress</code> command is run using <code>mknod</code> on the same command line. To prevent encountering this behavior, do not include <code>mknod</code> when running <code>fsstress</code>. Otherwise, you will be required to reboot the client without a successful unmount.</p> <p>Note: A StorNext 4.2.1 AIX client supports only AIX version 6.1</p>

Operating System	Feature or Category	Description
Solaris	StorNext labels	<p>Solaris hosts may need to rescan disk devices after StorNext labels have been applied.</p> <p>In particular, when a StorNext label is put on a LUN less than 1TB in size, Solaris hosts will not be able to use that LUN until they have done a device rescan. A device rescan is accomplished with a boot flag:</p> <pre>reboot -- -r</pre> <p>This issue will be addressed in a future StorNext release.</p> <p>In the meantime, work around this issue by rescanning devices using the boot flag <code>reboot -- -r</code></p> <p>If the labeling operation was performed on a Solaris host, that host does not need to do the rescan. However, some intermediate versions of the Solaris 10 Kernel Jumbo Patch break the necessary functionality to support this; please be sure you have applied the latest Solaris 10 Kernel Jumbo Patch before labeling any StorNext LUNs.</p>
Linux	Linux Multipath Support (the <code>rr_min_io</code> setting in the Linux DM Multipath Driver)	<p>Current versions of the Linux DM Multipath driver assign a default value of 1000 for <code>rr_min_io</code> which is too high for most configurations having multiple active paths to storage. Using a smaller value such as 32 will typically result in significantly improved performance. Refer to the RedHat or SuSE documentation provided with your version of Linux for details on how to apply this setting.</p> <p>Note: Experimentation may be required to determine the optimal value.</p>

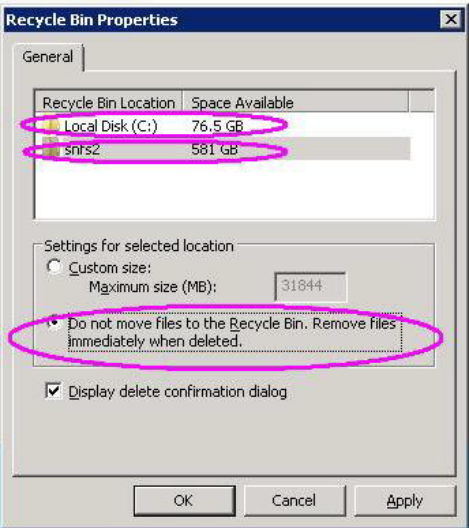
Operating System	Feature or Category	Description
Linux	StorNext File System	<p>StorNext File System does not support the Linux <code>sendfile()</code> system call.</p> <p>This issue causes Apache web servers to deliver blank pages when content resides on StorNext file systems.</p> <p>This issue also affects Samba servers running on Linux.</p> <p>The workaround is to disable <code>sendfile</code> usage by adding the following entry into the Apache configuration file <code>httpd.conf</code>:</p> <pre>EnableSendfile off</pre> <p>The workaround for Samba servers is to add the following line into the configuration file:</p> <pre>sendfile=no</pre>
	HA	<p>Changing the <code>haFsType</code> parameter in a file system configuration file to one of the HA types, and then (re)starting its FSM enables HA-specific features that change the functionality of StorNext.</p> <p>When the <code>HaShared</code> or <code>HaManaged</code> types are configured, other changes must be completed by successfully running the <code>cnvt2ha.sh</code> script, which is indicated by the creation of the <code>/usr/adic/install/.snsm_ha_configured</code> touch file (<code>\$SNSM_HA_CONFIGURED</code> environment variable). No conversion is done or necessary for SNFS only (<code>HaUnmanaged</code>) configurations.</p> <p>If the conversion is not successfully completed, the <code>HaManaged</code> FSMs will not start, and the <code>HaShared</code> FSM will cause an HA Reset when it is stopped.</p> <p>To remedy this situation, edit every FSM configuration file to set its <code>haFsType</code> parameter to <code>HaUnmonitored</code>, then run the following commands to avoid the HA Reset in this special case only:</p> <pre>touch /usr/cvfs/install/.vip_down_hint service cvfs stop</pre>

Operating System	Feature or Category	Description
Linux	System logs	<p>Due to the way Linux handles errors, the appearance of SCSI “No Sense” messages in system logs can indicate possible data corruption on disk devices.</p> <p>This affects StorNext users on Red Hat 4, Red Hat 5, SuSe 9, and SuSe 10.</p> <p>This issue is not caused by StorNext, and is described in detail in StorNext Product Alert 20.</p> <p>For additional information, see Red Hat 5 CR 468088 and SuSE 10 CR 10440734121.</p>
	Migrating metadata controllers	<p>StorNext users migrating their metadata controllers from Apple Xsan to Linux should be aware of the following upgrade considerations:</p> <ul style="list-style-type: none"> • If the file system is running Xsan 2.1.1 or earlier, it should be a simple upgrade: just replace the MDC. • If the file system is running Xsan 2.2 or later with “NamedStreams No” (which is the default for Xsan 2.2,) it should also be a simple upgrade: just replace the MDC. • If the file system is running Xsan 2.2 or later with “NamedStreams Yes,” you must completely remake (reformat) the file system. For obvious reasons, you should do a complete backup before migrating.
	Subtree Check option	<p>Subtree Check Option in NFS No Longer Supported</p> <p>Although supported in previous StorNext releases, the subtree_check option (which controls NFS checks on a file handle being within an exported subdirectory of a file system) is no longer supported as of StorNext 4.0.</p>

Operating System	Feature or Category	Description
Linux	FQDN	<p>SuSe Linux distributions automatically associate the FQDN of the local machine with the address 127.0.0.2 in the /etc/hosts file. There is no benefit from doing this when the machine is connected to a network that can resolve its name to an IP address.</p> <p>However, the existence of this entry can sometimes cause a failure of configuration synchronization within and between the server computers in an HA configuration. For this reason, the 127.0.0.2 entry should be deleted from the /etc/hosts file.</p>
	Software Firewalls	<p>Software firewalls such as “iptables” on Linux and Windows Firewall can interfere with the proper functioning of StorNext and result in unexpected errors unless specifically configured for use with StorNext.</p> <p>Quantum strongly recommends that all software firewalls be disabled on systems used as StorNext clients and servers. If required, StorNext can be configured for use with hardware firewalls.</p> <p>For more information, refer to the fsports man-page or help file and the “Ports Used By StorNext” section in the <i>StorNext Tuning Guide</i>.</p>
Windows and Linux	Symbolic links to StorNext directories	<p>If you create a symbolic (soft) link in Linux to a directory on a StorNext file system, the link cannot be used by Windows. Windows also cannot process a symbolic link which contains a path to a file in another file system.</p>
Windows	Window backup utility	<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	Feature or Category	Description
Windows	Upgrades on Windows Vista	<p>StorNext upgrades on Vista machines can fail in the middle of installation. This problem is caused by the way Windows Vista handles software upgrades. A related error is described in Microsoft article 263253. Microsoft has a utility called the Windows Installer Cleanup Utility that removes files left behind by incomplete installations. Access the Microsoft website and search for article ID 290301.</p> <p>To work around this issue, follow these steps:</p> <ol style="list-style-type: none"> 1 Click Start, and then click Run. 2 In the Open box, type Regedit and then click OK. 3 On the Edit menu, click Find. 4 In the Find what box, type Snfs_XXX.dat and then click Find Next. 5 If the search result selects a string value called PackageName, continue with these steps. Otherwise, repeat steps 3-4. 6 Double-click the PackageName string value. 7 In the Value data box, change the installation directory path to the new pathname. For example if the old installation directory path contained OCT10, change that to the current path (e.g, NOV12.) 8 On the Registry menu, click Exit.
	Offline Notification Feature	<p>The StorNext Offline Notification feature is intended for single user systems only. Do not install this feature on systems where multiple users might be logged on at the same time.</p>

Operating System	Feature or Category	Description
Windows	Recycle bin	<p>If you are using the StorNext client software with Windows Server 2003, Windows Server 2008, Windows XP, Windows Vista or Windows 7, 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.)</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	Feature or Category	Description
Windows	Recycle bin (cont.)	<p>(Disabling the Recycle Bin, Continued)</p> <p>For Windows Server 2008, Windows Vista and Windows 7 systems, you must disable the Recycle Bin on C: and the File system name:</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. 

Operating System	Feature or Category	Description
All	File systems and stripe groups	<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
	Managed file systems	For managed file systems only, the maximum recommended directory capacity is 50,000 files per single directory. (This recommendation does not apply to unmanaged file systems.)
	Internet Explorer 8	<p>When using StorNext with Internet Explorer 8, warnings about insecure and secure items may be generated. These warnings can be ignored and require no response. As a workaround, follow these steps to prevent the warning messages from appearing:</p> <ol style="list-style-type: none"> 1 Launch Internet Explorer 8. 2 Choose Internet Options from the Tools menu. 3 Click the Advanced tab. 4 Under the Security heading, make sure the option Display mixed content is checked. 5 Click OK.
	StorNext Licensing	After you purchase or update a feature license and then enter license information through the StorNext GUI, you should restart StorNext services to ensure that your new license is recognized. Certain StorNext features such as replication may not recognize your license until services are restarted.

Operating System	Feature or Category	Description
All	Upgrade	Before attempting to upgrade from a previous StorNext release, make sure you have free space on the file system. If the file system is nearly full when you begin the upgrade, serious errors may occur or the upgrade could fail. Best practice is to maintain an area on the file system which is not used for data or system files, but is reserved as an empty buffer to ensure that upgrades and other operations complete successfully.
	StorNext home page	The capacity indicators on the StorNext home page provide <i>approximations</i> and may not accurately summarize the actual current capacity. If you require accurate, up-to-the-minute capacity information, click the Capacity areas of the home page to view current capacity.
	Backups	Quantum recommends making two or more backup copies to minimize vulnerability to data loss in the event of hardware failure.
	Tape drives	StorNext does not support hot-swapping tape drives. When replacing or adding new tape drives you must first stop StorNext before installing the new drive.
	Cluster-Wide Central Control	The StorNext Cluster-Wide Central Control file (nss_cctl.xml) is used to enforce the cluster-wide security control on StorNext nodes (client nodes, fsm nodes, and nodes running cvadmin). This file is placed on an nss coordinator server. Currently the nss coordinator server capable of parsing this xml file must be on the Linux platform.
	Xsan	It is not possible to delete data within a StorNext policy relation point from an Xsan client via the Finder. Rather, data must be deleted using the shell.
	Converting file systems	StorNext does not currently support converting from a managed file system to an unmanaged file system.

Operating System	Feature or Category	Description
All	Labels	<p>Disks with existing non-StorNext labels may not show up in the StorNext GUI in order to protect non-StorNext disks from being accidentally overwritten. If you need to label a disk that is not visible in the StorNext GUI, use the <code>cvlabel</code> command to label the disk or use <code>cvlabel -U</code> to remove the existing label from the disks. (Refer to the <code>cvlabel</code> man pages for instructions on labeling and unlabeled drives.)</p> <p>Caution: Modifying the label on an active non-StorNext disk can make the disk unusable. Proceed with caution.</p>
	Replication	<p>When creating or editing a replication storage policy, there is a field on the Outbound Replication tab called "Filenames Excluded from Replication." This field allows you to exclude specific files from the replication process.</p> <p>This field works the same way as a UNIX shell which lets you pattern match names. For example, entering <code>*.0</code> core would exclude all .o files and also files named "core." You could also skip all core files by entering <code>rep_skip=core*</code>.</p>
	Replication/ Deduplication	<p>If you are using the Deduplication or Replication feature, part of the installation process is to update the on-disk index. The time required to complete this part of the installation process times may vary depending on the size of your licensed blockpool, drive performance, and other factors. As a general guideline, allow approximately five minutes for a 10TB blockpool.</p>

Operating System	Feature or Category	Description
All	Replication/ Deduplication	<p>When either the StorNext GUI or the <code>snpolicy</code> command are used to create or modify a replication/deduplication policy, a policy text file is written to the file system.</p> <p>Example: Suppose that <code>/stornext/photos/</code> is the mount point for file system named photos. If a policy named <code>pol_replicate_1</code> is created in that file system, a text copy of the policy information called <code>/stornext/photos/.rep_private/config/pol_replicate_1</code> is created.</p> <p>If the file system is damaged and has to be recreated, the policy must also be recreated. This is simpler to do beginning with the StorNext 4.1 release because a backup copy of the policy text file is saved whenever a policy is created or updated. (The backup copy is saved as a file named <code>/usr/cvfs/data/fsname/policy_history/policyname.date_time</code>.)</p> <p>In the previous example, the file system name (fsname) is photos and the policy name is <code>pol_replicate_1</code>. So, the backup copy would have a name like this:</p> <p><code>/usr/cvfs/data/photos/policy_history/pol_replicate_1.2010-10-29_14-07-13</code></p> <p>Note: The backup copy directory is not in the same file system is photos.</p> <p>If Storage Manager is used on the machine, all the policy backup files will be backed up along with the regular Storage Manager backups.</p> <p>Quantum suggests that after upgrading to StorNext 4.2.1 you run the command <code>snpolicy_gather -b > "some_file"</code></p> <p>This will save a copy of your current configuration. The <code>-b</code> option will also create a copy of policy information in the <code>usr/cvfs/data/fsname/policy_history</code> directory.</p>

Operating System	Feature or Category	Description
All	Deduplication	<p>If a deduplication candidate is removed before blockpool processing is completed, errors such as the following may be sent to the syslog:</p> <pre>Oct 2 15:22:00 orleans Blockpool[16403]: E: [5] (Store Local) Error storing file "/stornext/ source/ __CVFS_Handle.000474F892EBB65E000E000000000000000 000000292BF4". Error opening file "/stornext/source/ __CVFS_Handle.000474F892EBB65E000E000000000000000 000000292BF4". No such file or directory. Errors such as these may appear serious, but there is no reason for concern. If you receive these errors, no action is required.</pre>
	HA	<p>On HA systems only:</p> <p>The /usr/cvfs/config/ha_peer file supports some essential HA features by providing an address for HA administrative communications between the MDCs in an HA Cluster. If CVFS is started without this file having correct information, the probability of an HA Reset increases. To correct this condition, restore the ha_peer file to the IP address of the peer MDC, and restart StorNext by running the following command:</p> <pre>service cvfs restart</pre> <p>Note: The peer will be Primary after running this command.</p> <p>If the ha_peer file is removed for any length of time while StorNext is running, the snhamgr(1) HA Manager subsystem could stop functioning, which impacts the GUI HA Manage status page and the starting and stopping of CVFS, as well as any command line use of snhamgr itself. If this occurs, restore the ha_peer file to the IP address of the peer MDC, and then restart the HA Manager service by running the following command:</p> <pre>service snhamgr restart</pre>

Operating System	Feature or Category	Description
All	HA	<p>On HA systems only:</p> <p>You may receive the following incorrect error message when scanning for a secondary MDC from the StorNext Convert to HA page:</p> <pre>WARN com.quantum.qutosgui.jsf.ha.HaMBean - doScanHost: Secondary system cannot be same as the primary system.</pre> <p>This message is generated if <code>/usr/adic/util/cnvt2ha.sh</code> fails for any reason (for example, if the file system exists on the secondary, if a shared file system can't mount, etc). Upon secondary conversion failures, StorNext resets the <code>ha_peer</code> file to 255.255.255.255 on the secondary. Since the conversion fails, the primary <code>ha_peer</code> file is not updated and faulty comparison logic causes the erroneous error message (255.255.255.255 == 255.255.255.255).</p> <p>The workaround consists of two steps:</p> <ol style="list-style-type: none"> 1 Remove the <code>/usr/cvfs/config/ha_peer</code> file from the secondary system. 2 Reset the StorNext processes on the secondary system by running <code>/etc/init.d/stornext_web restart</code>.
	HA	<p>When exiting HA Config mode, StorNext will be stopped, which will also 'fuser' any processes which have files open on the file system from either node.</p> <p>Prepare systems and users for this eventuality before entering HA Config mode.</p>

Operating System	Feature or Category	Description
All	HA	<p>On HA systems only:</p> <p>When a non-managed file system is converted to a managed file system in an HA pair, it is possible for the FSMPM on the secondary MDC to continue to operate this FSM as non-managed, which incorrectly allows the FSM to start on the secondary MDC.</p> <p>Restarting the CVFS service corrects the problem. Quantum recommends taking the following steps as a temporary workaround after converting any non-managed file systems to managed file systems:</p> <ol style="list-style-type: none"> 1 Complete the configuration changes 2 Make sure that CVFS is running on the secondary MDC, and wait 120 seconds to be sure that the configuration-file changes have been synchronized to the secondary MDC 3 Restart CVFS on the secondary by issuing "service cvfs restart" 4 Issue the command "cvsadmin -e fsmlist" on the secondary MDC, and make sure that the output shows the FSM as "State: Blocked (waiting for MDC to become HA primary)"
	HA	<p>Use caution when configuring the netmask for the HA Virtual Interface (VIP).</p> <p>The VIP is an alias IP address that is associated with a real interface. For example, if the VIP is based on eth0, eth0:ha will be created as the VIP.</p> <p>The netmask you associate with the VIP should generally be the same as that of the base interface, but in no case should it be more specific. For example, if the netmask on eth0 is 255.255.224.0 (a /19), then configuring the VIP netmask as anything more than a /19, such as a /24 (255.255.255.0) would be incorrect.</p> <p>Using the same /19 mask on both eth0 and eth0:ha is the correct approach.</p> <p>Note: The above applies only when the IP address of the VIP falls into the subnet defined by the base interface's IP address and mask.</p>

Operating System	Feature or Category	Description
All	HA	<p>Understanding the performance of FSM failover in StorNext High Availability installations:</p> <p>When a failover of any file system occurs, the new FSM notices if any clients had a file exclusively opened for writes, and waits up to 35 seconds for those clients to reconnect. In the case of an HA Reset of the Primary MDC, that MDC is not going to reconnect, so the failover to FSMs on the Secondary MDC and the promotion of that MDC to Primary status can be delayed by 35 seconds.</p> <p>The StorNext system exclusively opens files on the HaShared file system, but assumes that only the Primary MDC does this and waives the delay for that one file system. Quantum advises against running user processes other than StorNext processes on HA MDCs for performance, reliability and availability reasons. In the event that processes running on the Primary MDC have files exclusively open for writes on other file systems, the availability of those file systems to all clients will be delayed by 35 seconds following an HA Reset event.</p>
	Quotas	When you enable or disable quotas using the CLI <code>cvadmin</code> command, the change does not persist after rebooting. In order to permanently enable or disable quotas, you must modify the Quotas parameter of the file system config file.
	fsretrieve	<p>If you run multiple <code>fsretrieve</code> commands simultaneously to find files (for example, <code>find -type -f xargs fsretrieve</code>), you might receive error messages because doing this taxes system resources.</p> <p>Instead, use the recursive retrieve command. When you use this command the files under a directory are retrieved in batches, and more sorting is done to put files in tape order for increased performance. Run recursive retrieve by entering</p> <pre>% fsretrieve -R .</pre>

Operating System	Feature or Category	Description
All	Distributed LAN	<p>Distributed LAN Clients in HA Environments:</p> <p>Each HA node must have its own dpserver files detailing the NICs on that node. The dpserver files are not synchronized between HA pairs. If the Distributed LAN Server is configured after converting to HA, the file system(s) running as Distributed LAN servers must be unmounted and mounted again to service DLC requests.</p> <p>When deduplication/replication is enabled, one or more Virtual IP Addresses (VIPs) provides access to the Primary MDC (where the blockpool server is running). In StorNext startup and failover situations, the VIP is dynamically associated with a physical address on the Primary server. Do not use VIP interfaces when setting up the dpserver configuration file, or it will not be available when the node is running as Secondary. The physical interface and IP address should be used in this situation.</p>
	Stripe group expansion	<p>StorNext does not support expansion on stripe groups containing mixed-sized LUNs. For example, if you create a file system that has two different-sized disks in a userdata only stripe group and then attempt to add a new disk to that stripe group and expand it, the expansion will fail.</p>
	dpserver	<p>In some cases the physical IP address must be included in the dpserver file in addition to the interface name. Note these conditions:</p> <ul style="list-style-type: none"> • When there is one IP address associated with a NIC interface, the interface name alone is a sufficient identifier • If there are multiple IP addresses associated with a NIC interface, one IP address is required in addition to the interface name • On HA systems, the physical IP address is required if virtual IP is configured for the NIC interface. (See also the following entry, "Distributed LAN Clients in HA Environment.")

Operating System	Feature or Category	Description
All	Truncation	By design, replication or deduplication must be completed before data files can be truncated if these files are associated with both a replication/dedup policy and a Storage Manager policy. Even if the Storage Manager policy is configured with the "Truncate Immediately" option, the truncation may not occur at store time unless the file has been replicated or deduplicated.
	DXi Virtual Tape Library Compatibility	Note the following recommendations and limitations for using DXi as a virtual tape library for StorNext: <ul style="list-style-type: none"> • Recommended library emulation: "ADIC Scalar i2000" • Recommended tape drive emulation: "IBM LTO-x" or "HP LTO-x" • DDM (Distributed Data Mover): This feature is currently not supported due to lack of full SCSI3 support in DXi.
	Affinities	When a file system with two affinities is to be managed by the Storage Manager, the GUI forces those affinities to be named tier1 and tier2. This will cause an issue if a site has an existing unmanaged file system with two affinities with different names and wants to change that file system to be managed. There is a process for converting a file system so it can be managed but it is non-trivial and time consuming. Please contact Quantum Support if this is desired. Note: The restriction is in the StorNext GUI because of a current system limitation where affinity names must match between one managed file system and another. If a site was upgraded from a pre-4.0 version to post-4.0, the affinity names get passed along during the upgrade. For example, if prior to StorNext 4.0 the affinity names were <i>aff1</i> and <i>aff2</i> , the GUI would restrict any new file systems to have those affinity names as opposed to <i>tier1</i> and <i>tier2</i> .



Appendix B

Additional Replication and Deduplication Information

This appendix contains detailed information about how replication and deduplication work, and about the underlying processes.

Replication Configuration File

StorNext includes a configuration file called `snpolicyd.conf` located at `/usr/cvfs/config/snpolicyd.conf`.

The `snpolicyd.conf` file provides a way to configure the `snpolicyd` process, which handles most aspects of StorNext replication and deduplication.

The man page for `snpolicyd.conf` contains detailed syntax, examples and instructions for modifying this file.

The remaining sections in this appendix also make reference to this file.

Replication Terminology and Conventions

StorNext has two kinds of policy:

- Storage Manager (SM) policies
- Replication/Deduplication policies.

For the sake of simplicity, in this appendix Replication/Deduplication policies will be called "snpolicyd" policies. Snpolicyd is the name of the Linux daemon that interprets and acts upon the policies.

StorNext users often talk about Storage Manager storing files to tape or retrieving files from tape, but Storage Manager can also use storage disk (called SDISK in SM) for storing files. In this appendix when we mention writing to or reading from tape, it includes using SDISK.

Copies and Versions

It's important to understand what "copies of a file or directory" means. There are several meanings, and this section attempts to clarify where StorNext supports additional copies and versions of a file or directory.

Context 1: "Number of copies to keep on target" is one property of an snpolicyd policy. This parameter specifies the number of replicated directories to keep on the target file system for a source directory. Remember, the replication process involves replicating a source directory and all of files and subdirectories that it contains. You can create from 1 to 16 target directories, depending on the "Number of copies to keep on target". Number of copies in this context means the number of target directory instances. By default, the different directories are differentiated by names like `dir`, `dir.1`, `dir.2`, and so on.

Context 2: Number of target file systems for an snpolicyd replication source policy. When configuring replication, you can specify up to three target file systems. For example, you could specify file system `/stornext/bk` on machine `hist1`, and file systems `/snfs/backup` and `/snfs/dr` on machine `host2`. Each of these directories can be a target of a replication source directory. The replication process is not complete

until each of the three target file system targets have been completely made.

If a replication source policy specified 10 for the "Number of copies to keep on target" and specified 3 target file systems, you would eventually have 30 replication directories: 3 after the first replication, 6 after the second replication, etc.

Context 3: Storage Manager number of copies. Storage Manager stores 1 through 4 copies of a file. The number of files is configured under the Steering tab when editing or creating a Storage Manager storage policy. (Actually, 4 is the default maximum number of copies. The number can be larger than 4.) Each of these copies is a copy of the same file contents.

Context 4: Storage Manager version. Versions refer to changed contents of the same file. By default, Storage Manager keeps ten versions of a file. Unlike Storage Manager copies, Storage Manager versions refers to different file contents. If there is a file called "log" that changes every day, and Storage Manager stores "log" every day, after ten days there would be ten versions of "log". The `fsversion` command is used to examine and change the Storage Manager versions of a file.

Context 5: Storage Manager File Recovery. When a file is removed from a Storage Manager relation point, the previous copies stored by Storage Manager are still on media and in the SM database. These previous versions may be recovered using the `fsrecover` command. There is no a limit to the number of SM instances which can be recovered in this manner. Eventually the administrator may use the `fsclean` command to clean up older versions of SM media. After running `fsclean`, files that used to reside on the media can no longer be recovered with the `fsrecover` command.

Replication Target Directories

Replication results in a directory on the target that represents the files that were in the source directory at the time of the replication. The source and target directories could be on the same machine (node) or different machines. Also, StorNext can replicate either deduplicated data or non-deduplicated data.

Number of Replication Copies

When a source directory is replicated to a target there can be from 1 through 16 replicated target directories that reflect replications of the source at different times. The number of copies is specified by the "Copies to Keep on Target" parameter on the Inbound Replication tab or Outbound Replication tab. You enter parameters on these tabs when configuring a snpolicyd storage policy.

The "Copies to Keep on Target" selection allows values of 1 through 16, and also a special case called in-place. We will not discuss the in-place selection in this section.

First, let's consider the case where "Copies to Keep on Target" is 2. Each time a replication occurs a new target directory is created. This target directory might have the same name as the previous target directory, but it is actually a new directory. The new directory reflects files added, deleted, and changed since the previous replication.

It is important to understand that in this example the target is a *new* directory. This has implications that might not be immediately obvious. For one thing, it means we cannot use the target directory in exactly the same way as we might use the source directory. Following is an explanation and examples.

Example: Copies on Target = 2

In this example, we replicate source directory `/stornext/snfs1/photos`, a directory in file system `/stornext/snfs1`, to a target directory `/stornext/backup/photos` in file system `/stornext/backup`. (For this example it doesn't matter whether the two file systems are on the same node machine or on different machines.) Since we are keeping two copies on the target, we will usually have two directories on the target:

- `/stornext/backup/photos` - *most recent replication*
- `/stornext/backup/photos.1` - *previous replication*

When the next replication occurs, the following directory changes Take place:

- The previous replication `/stornext/backup/photos.1` is removed
- The most recent replication `/stornext/backup/photos` is renamed `/stornext/backup/photos.1`
- The new replication appears in `/stornext/backup/photos`

Now consider a Linux shell process that is executing inside directory `/stornext/backup/photos`. When the next replication occurs, the directory still exists but is named `/stornext/backup/photos.1`. If the Linux shell executes the command `ls -l`, the `ls` command lists the *previous* contents of `photos` - the directory now named `photos.1`.

When the replication after that occurs, the original directory is removed. When the shell executes `ls -l`, the command shows no files since the original directory and its contents have been removed.

Thus, a process executing inside a replication directory may see files in the directory at one time and see no files a while later. This is different behavior than we would expect to see when a process is executing inside the original source directory.

Similar surprising behavior occurs if the replicated directory is NFS exported or Samba/CIFS shared. Suppose directory `/stornext/backup/photos` is NFS exported on the target machine. The directory can be NFS mounted on another Linux or Unix machine. The mounted NFS file system can generate errors (input/output error, stale NFS file handle) on the client when the original directory changes due to replication.

The bottom line is that you must be aware that changes occur with the replicated directory. The replicated directory should not be used as a substitute for the original source directory unless you take precautions to isolate the application from unexpected changes.

Isolating a Replication Target Directory

To isolate a replication target directory, use the `snpolicy` command's `-exportrepcopy` option. This operation is available only from the command line, not from the StorNext GUI.

First, use the `-listrepcopies` option on the target node to determine the association between the target copy number and the target directory to use. The `-listrepcopies` output provides the "key" value for the policy used to implement this replication. For example, if the target file system is `/snfs/rep`, use the command:

```
/usr/cvfs/bin/snpolicy -listrepcopies=/snfs/rep
```

Here is the relevant part of the command output:

```
source://snfs/sn1@10.65.170.108:/project?key=402 ->
target://snfs/rep@node2:?key=406
```

```
0 -> /snfs/rep/project
1 -> /snfs/rep/project.1
2 -> /snfs/rep/project.2
3 -> /snfs/rep/project.3
```

The copy number appears in the left column, and the realization directory for that copy number is shown after the "->".

There are two "keys" shown in the sample output. The first key (402) is the key value for the source policy controlling the replication. The second key value (406) controls the replication realization of the target.

Let's say you want to copy files back from `/snfs/rep/project.2`. To isolate `/snfs/rep/project.2` you would use this command:

```
/usr/cvfs/bin/snpolicy -exportrepcopy=/snfs/rep/ --
key=406 -copy=2 --path /snfs/rep/project_temp
```

This command renames the directory `/snfs/rep/project.2` to `/snfs/rep/project_temp` and prevents the policy daemon from affecting this directory, in case replications for this target policy become activated again during the recovery process.

The `-path` argument is optional: you can do only the `exportrepcopy` operation and use the directory name `/snfs/rep/project.2` when recovering replicated files.

The point of this is that using the `-exportrepcopy` option allows you to use the directory without having to worry about it changing, or files disappearing as you do your work.

Once a directory has been isolated in this manner, it can then be transformed into a replication source directory for rereplication to another file system and/or machine.

Final Recommendation For Target Directories

You should not change the contents of a replication target directory. It should be treated as a "read-only" replica, even though StorNext does not enforce a read-only restriction.

If you change a file in a replication target directory you may be changing the file contents in other target directories due to the "hard-link" usage in replication. Furthermore, if you change or add files in a directory, that directory may disappear due to subsequent replications. (Using `exportrepcopy` avoids this second issue.)

What if you want to change an existing source directory into a target directory? This can be done, but without special configuration care the original source policy assignment will be lost. A directory can have only one snpolicyd policy assigned to it (and all of the files and subdirectories it contains.) If you change the policy assignment, the characteristics specified in the previous policy are forgotten.

StorNext snpolicyd Policies

You can create and edit StorNext snpolicyd policies from the StorNext GUI or with the snpolicy command. These snpolicyd policies differ from StorNext Storage Manager (SM) policies in several respects. Following is a summary of some of the similarities and differences between these two kinds of policies.

Storage Policy Option	Storage Manager Policy	snpolicyd Policy
<i>Configurable via the StorNext GUI?</i>	Yes. Select the Storage Policies menu's Storage Manager option.	Yes. Select the Storage Policies menu's Replication / Deduplication option.
<i>Configurable via the command line?</i>	Yes. Use fs commands such as fsaddclass and fsmodclass	Yes. Use the snpolicy command.
<i>Where are policy internals stored?</i>	In Storage Manager Database. One database per machine.	In the managed file system, in a private directory.
<i>Is the policy used across file systems?</i>	Yes. One policy can be used in multiple directories and multiple file systems.	No. Policies apply to one file system, but can be applied to multiple directories.
<i>Functions?</i>	Store (to tape or SDISK), retrieve, truncate files.	Deduplicate, replicate, truncate files.
<i>How are truncated files retrieved?</i>	The entire file must be retrieved.	Only portions of the file containing needed regions may be retrieved.
<i>Schedules?</i>	fspolicy / schedules stored in Database.	Linux crontab scheduling.

Storage Policy Option	Storage Manager Policy	snpolicyd Policy
<i>Management daemon?</i>	multiple fs_... processes	snpolicyd
<i>Previous file versions recoverable?</i>	Yes. Recover previous tape version with the fsrecover command. Up to 10 tape versions.	Yes. Previous replicated copies can be kept in previous replication directories. Up to 16.

Example

You create an snpolicyd policy with the StorNext GUI or with the snpolicy command. The snpolicy command is in directory /usr/cvfs/bin. Command line configuration must be done by the Linux root user.

Suppose you create directory /stornext/snfs1/photos in file system /stornext/snfs1 on machine host1. You then use the StorNext GUI to create a replication policy named photo_rep to replicate this directory to file system /stornext/backup on machine host2. (As in the previous example, the policy was configured to keep two copies on the target.)

Now use the snpolicy command to see more internal details about the policy called photo_rep.

Use this command:

```
/usr/cvfs/config/snpolicy -dumppol/stornext/snfs1/photos
```

The command's output looks like this:

```
inherit=photo_rep
key=1720399
root=/stornext/snfs1/photos
dedup=off
dedup_filter=off
max_seg_size=1G
max_seg_age=5m
dedup_age=1m
dedup_min_size=4K
dedup_seg_size=1G
dedup_min_round=8M
dedup_max_round=256M
```



```
dedup_bfst="localhost"  
fencepost_gap=16M  
trunc=off  
trunc_age=365d  
trunc_low_water=0  
trunc_high_water=0  
rep_output=true  
rep_report=true  
rep_target="target://stornext/backup@host2:"  
rep_copies=2
```

There is a lot of output, most of which we don't have to consider now. Some of the important values are:

- `inherit=photo_rep`: This means the policy controlling this directory receives its parameters from the policy named `photo_rep`. Remember, when you create a policy you give it a name, and the policy name belongs to the file system. There could be a different policy named `photo_rep` in a different file system, and there would be no connection between the two `photo_rep` policies.
- `rep_output=true`: This means the policy is a source of replication.
- `rep_copies=2`: This means you want to keep two copies (instances) of the replicated directory on the target file system.
- `rep_target="target://stornext/backup@host2:"`: This tells you the replication target directory is a directory in file system `/stornext/backup` on machine `host2`. But which directory name will be used in that file system? Since you did not specify anything else, the source directory name will be used. In this case the source directory name in the source file system is `photos`, so the target directory names will be `/stornext/backup/photos` and `/stornext/backup/photos.1`.
- `dedup=off`: This means the files in this directory are not deduplicated before being replicated. Deduplication and replication are discussed in another section.

One comment about a field *not* in the command output. Since there is no line for `rep_input=true`, this means this directory is not a replication target directory. This is not surprising. While it is true that a replication target can also be a replication source, that is an advanced case not covered here.

Replication Copies = 2 (Detail)

In this section we'll examine in more detail what two copies on the target (`rep_copies=2`) means.

Assume that we begin with files `file1`, `file2`, and `file3` in the source directory. After the first replication, we expect to see three files in the target directory `/stornext/ backup/photos`.

After running the command: `ls -l /stornext/backup/photos`, the output looks like this:

```
total 4144
-rwxr-xr-x 2 testuser root 1388936 Jan 26 10:11 file1
-rw-r--r-- 2 testuser root 1430896 Jan 26 10:11 file2
-rw-r--r-- 2 testuser root 1397888 Jan 26 10:12 file3
```

Notice the "link count" of 2 in front of the user name for each file. This means that each of these files has two links - two names. One name is the entry in directory `/stornext/backup/photos`. The other name is a name in a subdirectory of `/stornext/backup/.rep_private`. As its name suggests, directory `.rep_private` in the root of a managed file system contains internal information used to manage replication and deduplication.

Digression

Following is some additional detail which you may want to skip the first time you read this section.

Below is output from the command `ls -l /stornext/backup/.rep_private`:

```
total 144
drwx----- 19 root root 2057 Jan 26 10:12
00047DA110919C87
drwx----- 3 root root 2054 Jan 26 10:12 config
drwx----- 3 root root 2056 Jan 25 14:11 oldest
drwx----- 3 root root 2116 Jan 26 10:13 pending
drwx----- 3 root root 2132 Jan 26 10:13 queued
drwx----- 2 root root 2048 Jan 21 16:56 source_state
drwx----- 3 root root 2048 Jan 20 17:13 target
```

```
drwx----- 2 root root 2116 Jan 26 10:13 target_state
drwx----- 2 root root 2255 Jan 26 10:13 tmp
```

This output shows a list of directories underneath `.rep_private`. The directory we are interested in now is `00047DA110919C87`. Where did the directory name `00047DA110919C87` come from? It is the file system ID of the source file system, a unique string which can be used to identify that file system.

If you execute the command `ls -l /stornext/backup/.rep private/00047DA110919C87`

you would see one or more directories similar to this:

```
drwx----- 3 root root 2052 Jan 26 09:30 1720408
drwx----- 3 root root 2063 Jan 26 10:13 1720418
drwx----- 3 root root 2048 Jan 21 12:12 475
```

Here the directory names are 1720408, 1720418, and 475. Those names actually reflect the inode number of the directories on the source file system. In this case the directory we want is 1720418.

If you execute the command `ls -l /stornext/backup/.rep private/00047DA110919C87/1720418`

you would see the following:

[illegible]

The important files are the three in the middle named 0x1a4065.0, 0x1a4066.0, and 0x1a4067.0. As you can see from the file owner (testuser) and file lengths, these three files are the other three names (links) of file1, file2, and file3.

(End of Digression)

Second Replication

The "standard" case - when we have replicated files once - is that the link count for the target file will be two.

Now let's say that we add `file4` and `file5` to the source directory and replicate again. After the second replication, target directory `/stornext/backup/photos` contains the following:

```
total 6864
-rwxr-xr-x 3 testuser root 1388936 Jan 26 10:11 file1
-rw-r--r-- 3 testuser root 1430896 Jan 26 10:11 file2
-rw-r--r-- 3 testuser root 1397888 Jan 26 10:12 file3
-rwxr-xr-x 2 testuser root 1388994 Jan 26 11:02 file4
-rwxr-xr-x 2 testuser root 1388965 Jan 26 11:03 file5
```

Target directory `/stornext/backup/photos.1` contains the previous replication:

```
total 4144
-rwxr-xr-x 3 testuser root 1388936 Jan 26 10:11 file1
-rw-r--r-- 3 testuser root 1430896 Jan 26 10:11 file2
-rw-r--r-- 3 testuser root 1397888 Jan 26 10:12 file3
```

Notice that `file1`, `file2`, and `file3` each have a link count of 3. One link (name) is in directory `photos`, another link is in directory `photos.1`, and the third is the `snpolicyd` "internal" link in the `.rep_private` directory. The two new files, `file4` and `file5`, appear only in the new directory and in the `.rep_private` directory. They have a link count of 2.

Since `file1`, `file2`, and `file3` are really the same file in directories `photos` and `photos.1`, no extra disk storage is needed for these files when replicating again. In general, when you use replication with more than one copy retained on the target, no additional storage is needed for unchanged files. If a file is changed, both the old and the new version are retained, so additional storage is needed in this case. (Unless deduplication is also used, which is discussed later.)

Now let's make two changes. Say we remove `file4` in the source directory and modify `file2`. After the next replication, target directory `photos` contains the following:

```
total 5200
-rwxr-xr-x 3 testuser root 1388936 Jan 26 10:11 file1
```

```
-rw-r--r--    2 testuser root 1123155 Jan 26 11:20 file2
-rw-r--r--    3 testuser root 1397888 Jan 26 10:12 file3
-rwxr-xr-x    3 testuser root 1388965 Jan 26 11:03 file5
```

Target directory photos.1 contains:

total 6864

```
-rwxr-xr-x    3 testuser root 1388936 Jan 26 10:11 file1
-rw-r--r--    1 testuser root 1430896 Jan 26 10:11 file2
-rw-r--r--    3 testuser root 1397888 Jan 26 10:12 file3
-rwxr-xr-x    2 testuser root 1388994 Jan 26 11:02 file4
-rwxr-xr-x    3 testuser root 1388965 Jan 26 11:03 file5
```

The three files, file1, file3 and file5, were unchanged, so they have the expected link count of 3. One name occurs in photos, one in photos.1, and the third in a subdirectory of .rep_private. Since file2 was changed in directory photos, it has a link count of 2: one link in photos and one in .rep_private.

The file named file2 in photos.1 now has a link count of 1. It is not the same file as the current file2 (notice the different length). The file2 in photos.1 is there for "historical" or recovery purposes only. It represents the previous replication of the directory.

Notice also that file4 in photos.1 has a link count of 2: one for the photos.1 copy and one for the .rep_private copy. There is no file named file4 in the current replication directory named photos.

File1, file3 and file5 share the same disk storage. The storage for file4 is only shared with the .rep_private copy, and this storage will be freed when the next replication occurs. The older version of file2 exists only in photos.1, and its storage will be removed in the next replication.

More About Replication Target Directories

In the previous replication example, source directory /stornext/snfs1/photos on host1 was replicated to target directory /stornext/backup/photos on host2. If the number of copies to keep is more than 1, the previous replication directories are named /stornext/backup/photos.1, /stornext/backup/photos.2, etc.

The default name on the target is the same pathname relative to the file system mount point as the source directory name is relative to the source file system mount point.

Examples

Suppose you apply a replication policy to directory `/stornext/snfs1/a/b/c/d/photos` in file system `/stornext/snfs1`, and replicate to file system `/stornext/backup`. The default target replication directory name would be `/stornext/backup/a/b/c/d/photos`, and previous replication directories would be `stornext/backup/a/b/c/d/photos.1`, etc.

There are other options that can be specified on either the source policy or on the target policy. Since we have been concentrating on the source policy, following are examples of changes there.

When creating or editing a policy, specify the alternative path names in the area of the screen labeled **Pathname on Target** on the **Outbound Replication** tab. When you click the **Override** label, a field appears where you can type some text. Some hints appear above that field, showing special entry values such as `%P` and `%D`.

In all of the following examples, assume that the replication source directory is `/stornext/snfs/photos/ocean` in directory `photos/ocean` relative to the source file system `/stornext/snfs1`. For this example we will replicate to file system `/stornext/backup`. We know that if we do not override the "Pathname on Target" value, the replication target directory name will be `/stornext/backup/photos/ocean`.

- If you enter a string without any of the "%" formatting characters, the replication directory will be the name you specify. For example, if you specify `open/sesame` for Pathname on Target, the replication directory would be `/stornext/backup/open/sesame`.
- `%P` means source pathname relative to the source file system. For example, if you specify `open/sesame/%P` for Pathname on Target, the replication directory would be `/stornext/backup/open/sesame/photos/ocean`
- `%D` means today's date. `%T` means the replication time. For example, if you specify `%D/%T/%P` for Pathname on Target, the replication directory would be `/stornext/backup/2010-02-02/16_30_22/photos/ocean` (on February 2, 2010).

- %H means source hostname. This would be a good value to use when more than one source machine is replicating files to the same target machine and target file system.

There are a lot of ways the "%" characters, and name specifications can be combined.

Note two important facts:

- It is possible to generate target name collisions by specifying the same Pathname on Target for more than one policy. For example, you might choose "daily" for Pathname on Target in two source replication policies. In that case the first policy to replicate would succeed, and the second would fail due to the name collision. Using %H, %P, etc. can help you avoid these collisions.
- Specifying a Pathname on Target is required if you want to replicate into a Storage Manager relation point. This will be discussed further in another section.

Deduplication Overview

Here is the view from 100,000 feet. When StorNext deduplication is enabled, a file is examined and logically split into data segments called BLOBs (binary large objects). Each BLOB has a 128-bit BLOB tag. A file can be reconstructed from the list of BLOBs that make up a file. The data for each BLOB is stored in the blockpool for a machine. We can use the command `snpolicy -report file_pathname` to see the list of BLOB tags for a deduplicated file.

When a deduplicated file is replicated, the BLOBs are replicated from the blockpool on the source machine to the blockpool on the target machine. If the source file system and the target file system are both hosted on the same machine, no data movement is needed. If the same BLOB tag occurs several times (in one file or in many files) only one copy of the data BLOB exists in the blockpool. During replication that one copy must be copied to the target blockpool only once.

This is why deduplicated replication can be more efficient than non-deduplicated replication. With non-deduplicated replication, any change in a file requires that the entire file be recopied from the source to the target. And, if the data is mostly the same in several files (or

exactly the same), non-deduplicated replication still copies each entire file from the source file system to the target.

The following example uses these three files and their corresponding sizes:

f.2m - 2 MB

f.4m - 4 MB

g.4m - 4 MB

The maximum segment size in this example is 1 MB. (That size is artificially low for this example only.)

If we look at the "snpolicy -report" output for the directory containing these files, we see the following:

```
./f.2m
  policy: 1720449      inode: 1720468
  flags: TAG
  mtime: 2010-01-26 14:20:03.590665672 CST
  ingest: 2010-01-26 14:20:03.590665672 CST
  size:      2097152 disk blocks: 4096
  seqno:      4 blk seqno:      2
  offset:      0 length:      1048576 tag:
D03281B0629858844F20BB791A60BD67
  offset:      1048576 length:      1048576 tag:
12665A8E440FC4EF2B0C28B5D5B28159
./f.4m
  policy: 1720449      inode: 1720470
  flags: TAG
  mtime: 2010-01-26 14:22:56.798334104 CST
  ingest: 2010-01-26 14:22:56.798334104 CST
  size:      4194304 disk blocks: 8192
  seqno:      4 blk seqno:      4
  offset:      0 length:      1048576 tag:
D03281B0629858844F20BB791A60BD67
  offset:      1048576 length:      1048576 tag:
12665A8E440FC4EF2B0C28B5D5B28159
  offset:      2097152 length:      1048576 tag:
7F02E08B3D8C35541E80613142552316
  offset:      3145728 length:      1048576 tag:
1FEC787120BEFA7E6685DF18110DF212
```



```
./g.4m
  policy: 1720449      inode: 1720471
    flags: TAG
    mtime: 2010-01-26 14:23:28.957445176 CST
    ingest: 2010-01-26 14:23:28.957445176 CST
    size:      4194304 disk blocks: 8192
    seqno:      5 blk seqno:      4
    offset:      0 length:      1048576 tag:
D03281B0629858844F20BB791A60BD67
    offset:      1048576 length:      1048576 tag:
DF54D6B832121A80FCB91EC0322CD5D3
    offset:      2097152 length:      1048576 tag:
7F02E08B3D8C35541E80613142552316
    offset:      3145728 length:      1048576 tag:
1FEC787120BEFA7E6685DF18110DF212
```

All three files have the same contents in the first megabyte starting at offset 0. The tag for that BLOB is D03281B0629858844F20BB791A60BD67, and that BLOB is stored only once in the blockpool. The second megabyte is the same for files f.2m and f.4m (tag 12665A8E440FC4EF2B0C28B5D5B28159) but file g.4m has a different BLOB in those bytes. The final 2 megabytes of files f.4m and g.4m are the same.

Remember that the above is an artificial example. In actual practice BLOBs do not line up on 1 MByte boundaries and are not all the same length.

Enabling Deduplication

When creating or editing a policy through the StorNext GUI, select the **Deduplication** tab and make sure deduplication is enabled (On). If you use the `snpolicy dumppol` option, you will see `dedup=on` in the output when the policy has deduplication enabled.

Deduplication Modification Time

Note that in the "`snpolicy -dumppol`" output shown earlier we also saw `dedup_age=1m`. This means the file may be deduplicated after it has not changed for at least one minute. If a file is being written its file modification time (mtime) will be updated as the file is being written. Deduplication age specifies how far in the past the modification time must be before a file can be considered for deduplication.

Deduplication and Blockpools

If replication is used, a blockpool is required even if deduplication is not used in any policy on a machine. However, in this situation the blockpool does not store any BLOBs from any file system and can therefore be small: several megabytes is all that is needed.

If you enable deduplication on any policy in the machine, StorNext stores BLOBs in the blockpool and additional space is required. Make sure you have enough space to store file system data if you enable deduplication. You also need space for BLOBs in the blockpool if the machine contains replication target directories for deduplicated replication source directories on other machines.

The current StorNext release supports only one blockpool per machine. Any file system on the machine that needs a blockpool will use that one and only blockpool.

Deduplication and Truncation

Let's look again at the directory in the previous section that has the three files `f.2m`, `f.4m`, and `g.4m`. Using the Linux command `ls -ls` shows this in the directory:

```
total 10240
2048 -rw-r--r-- 1 root root 2097152 Jan 26 14:22 f.2m
4096 -rw-r--r-- 1 root root 4194304 Jan 26 14:22 f.4m
4096 -rw-r--r-- 1 root root 4194304 Jan 26 14:23 g.4m
```

The first column on the left shows the total number of blocks (1024 bytes per block) contained in each file. The column before the date shows the file size in bytes.

StorNext can truncate files that have been deduplicated. By “truncate” we mean that the disk blocks for the file have been freed. If the deduplicated files shown above are truncated, the `ls -ls` command shows this:

```
total 0
0 -rw-r--r-- 1 root root 2097152 Jan 26 14:22 f.2m
0 -rw-r--r-- 1 root root 4194304 Jan 26 14:22 f.4m
0 -rw-r--r-- 1 root root 4194304 Jan 26 14:23 g.4m
```

There are no blocks in any of the three files, although each file retains its correct size.

(As an exercise, in the previous "ls -l" and "ls -ls" examples, what does the line that says "total some_number" tell us?)

When an application or command accesses any of the data in a truncated file, StorNext retrieves the data it needs from the blockpool. This may be the entire file for a small file. For a larger file, a portion of the file would be retrieved: a portion at least large enough to contain the file region required. If you read the entire file, the entire file will be retrieved.

Truncation provides the mechanism by which file system storage space may be reduced. When a file is truncated it takes no space in its file system, but space for its BLOBs is required in the blockpool. If we receive deduplication benefit (that is, if the same BLOB data occurs in more than one place,) then we have less space used in the blockpool than would be in the original file system.

Enabling Deduplication and Truncation

In order to enable truncation, both deduplication and truncation must be enabled in the storage policy. The StorNext GUI contains tabs for both deduplication and truncation which allow you to enable deduplication and truncation respectively.

Before a file is truncated it must pass a "Minimum Idle Time Before Truncation" test. If this minimum age is ten minutes, then ten minutes must elapse after the last file modification or file read before truncation can occur. The default value for the minimum idle time is 365 days.

In the output from "snpolicy -dumpool" the parameters we have been discussing are displayed like this:

```
trunc=on
```

```
trunc_age=365d
```

Storage Manager Truncation

Storage Manager also truncates files. Storage Manager truncation is similar to but not identical with the deduplication-based truncation we have been discussing. Storage Manager truncation will be discussed again when we consider deduplication / replication with Storage Manager.

Replication, Deduplication and Truncation

Consider a directory which is being deduplicated and replicated. We mentioned earlier that in this case data BLOBs move from the blockpool on the source machine to the blockpool on the target machine. When replication happens (the replication namespace realization,) the files appear in the target directory as truncated files. This is true regardless of whether or not the files were truncated in the source directory at replication time.

Let's look again at the example target directories `photos` and `photos.1` after the last replication. If the replication source directory had deduplication enabled, then `"ls -ls"` in target directory `photos` shows the following:

```
total 0
0 -rwxr-xr-x 3 testuser root 1388936 Jan 26 10:11 file1
0 -rw-r--r-- 2 testuser root 1123155 Jan 27 11:20 file2
0 -rw-r--r-- 3 testuser root 1397888 Jan 26 10:12 file3
0 -rwxr-xr-x 3 testuser root 1388965 Jan 26 11:03 file5
```

Target directory `photos.1` contains the following:

```
total 0
0 -rwxr-xr-x 3 testuser root 1388936 Jan 26 10:11 file1
0 -rw-r--r-- 1 testuser root 1430896 Jan 26 10:11 file2
0 -rw-r--r-- 3 testuser root 1397888 Jan 26 10:12 file3
0 -rwxr-xr-x 2 testuser root 1388994 Jan 26 11:02 file4
0 -rwxr-xr-x 3 testuser root 1388965 Jan 26 11:03 file5
```

The file link counts (3, 2, or 1) are the same as in the earlier replication example. The principle is the same: `file1` in `photos` has 3 links. The other two instances are `file1` in `photos.1` and a file underneath the `.rep_private` directory. All the links are to a truncated file: a file whose length is 1388936 bytes, but which contains no blocks. If we read any of the three links, the file would be partially or fully retrieved.

The replicated files appear as truncated files even if deduplication is not explicitly enabled in any policy on the target machine. Remember that this means there must be blockpool space for the replicated BLOBs if a deduplicated directory is replicated onto the machine.

Replication, Deduplication and Storage Manager

Both StorNext replication and StorNext deduplication can be used with Storage Manager. The following discussion assumes you are already familiar with replication and deduplication, and also with Storage Manager.

Here are some interesting possibilities:

- 1 Replicate from a source directory into a target directory where the target directory is within a Storage Manager relation point. Then the replicated files will be stored to tape by Storage Manager. This can be done with deduplicated or non-deduplicated replication.
- 2 Replicate from a source directory that is managed by Storage Manager. This can be done with deduplicated or non-deduplicated replication. It doesn't matter for the source if the target directory is also managed by Storage Manager.
- 3 Use deduplication within a Storage Manager relation point. This means the files will be deduplicated, and the deduplicated data will be stored in the blockpool. In addition, Storage Manager will make tape copies of the files.

Let's consider replicating into a Storage Manager relation point.

Replicating into a Storage Manager Relation Point

To replicate into a relation point, specify a target directory underneath a Storage Manager relation point. Do this with the parameter "Pathname on Target" in the StorNext GUI, or with `rep_realize=...` when configuring a policy with the `snpolicy` command.

Example

Suppose we are replicating to file system `/stornext/backups` on a target machine, and `/stornext/backups/sm1` is a Storage Manager relation point in that file system.

Some possible choices for "Pathname on Target" would be

- `sm1/%P`
- `sm1/mystuff`

- `sm1/%H/%P`

You shouldn't specify something like `/stornext/backups/sm1/mystuff` because "Pathname on Target" is relative to the target file system mount point, which in this case is `/stornext/backups`.

If "Copies to Keep on Target" is more than 1, the rules discussed earlier determine the names for the directories in subsequent replications.

Example

If we replicate the source directory named `photos` into a relation point using the "Pathname on Target" `sm1/%P`, we end up with directories like `/stornext/backups/sm1/photos`, `/stornext/backups/sm1/photos.1` and so on for the replicated directories when we are keeping more than one copy on the target.

The directories `photos` and `photos.1` are in the SM relation point. Let's say we have the two directories `photos` and `photos.1` with the contents that we discussed earlier.

Target directory `/stornext/backups/sm1/photos` contains the following:

```
-rwxr-xr-x 3 testuser root 1388936 Jan 26 10:11 file1
-rw-r--r-- 2 testuser root 1123155 Jan 27 11:20 file2
-rw-r--r-- 3 testuser root 1397888 Jan 26 10:12 file3
-rwxr-xr-x 3 testuser root 1388965 Jan 26 11:03 file5
```

Target directory `/stornext/backups/sm1/photos.1` contains the following:

```
-rwxr-xr-x 3 testuser root 1388936 Jan 26 10:11 file1
-rw-r--r-- 1 testuser root 1430896 Jan 26 10:11 file2
-rw-r--r-- 3 testuser root 1397888 Jan 26 10:12 file3
-rwxr-xr-x 2 testuser root 1388994 Jan 26 11:02 file4
-rwxr-xr-x 3 testuser root 1388965 Jan 26 11:03 file5
```

Question: Will Storage Manager store all the files in `photos` after the most recent replication? The answer is no. In this example, `file2` is a file that was modified since the previous replication. Thus `file2` is the only file that will be stored by Storage Manager after the most recent replication.

When replication occurs we create store candidates for the new or changed files that were included in the most recent replication within a

relation point. In this example, only `file2` will be a store candidate after the latest replication. You can use the `showc` command to see the new Storage Manager store candidates after a replication.

Note: Even if you created a store candidate for every file in the replicated target directory, only the new or changed files would be stored by SM. This is because the other files are links to files that have already been stored by Storage Manager, or at least files that were already on the Storage Manager store candidates list.

Truncation and Deduplication / Replication (with and without SM)

We have already mentioned how deduplication allows files to be truncated. “Truncated” in this case means that the extents have been partially or completely removed from disk, and that the `snpolicyd` daemon must reconstitute the missing extents when a process wants to access them.

We also discussed how we can use the `ls -ls` command to identify truncated files. We looked for files with “0” in the first column of the output of `ls -ls`. The 0 means there are no blocks associated with the file. The file size field in the `ls -l` or `ls -ls` output reflects the real size of the file, and is not changed when the file is truncated.

Example

In the earlier example we this saw this output (for a truncated file) after running `ls -ls`:

```
0 -rwxr-xr-x 3 testuser root 1388936 Jan 26 10:11 file1
```

For an untruncated file, the `ls -ls` output might look something like this:

```
1360 -rwxr-xr-x 3 testuser root 1388936 Jan 26 10:11 file1
```

The 1360 blocks in this file are enough to contain a file of size 1388936 (since $1360 * 1024 = 1392640$). However, we might also see a blocks value that was non-zero but not enough to contain the entire file size. This might indicate the following:

- A sparse file (this will not be discussed here)
- A file with a stub left on the disk

- A file that had been partially retrieved

Both Storage Manager and snpolicyd (replication / deduplication) can truncate files and can retrieve truncated files.

Both Storage Manager and snpolicyd can be configured to leave a stub file on disk when a file is truncated. Using the StorNext GUI you can configure the deduplication stub size on the **Deduplication** tab when creating or editing a replication / deduplication policy. A non-zero stub size must be a multiple of the file system block size.

Both Storage Manager and snpolicyd will retrieve a file when a portion of the file is read that is not already on disk. For Storage Manager there are really three different possibilities for a file's truncation state:

- File is totally truncated. The file has no block in the file system. Reading any byte of the file causes Storage Manager to retrieve the entire file.
- File is truncated, but there is a stub. Reading within the stub causes no retrieval. Reading anything not in the stub causes Storage Manager to retrieve the entire file.
- File is completely on disk.

For a truncated file that was deduplicated by snpolicyd, there can be partial file retrieval from the blockpool. In this situation there is one more possibility in addition to the three previous possibilities:

- Partially retrieved. The file has some data on disk (besides the stub) but the entire file is not on disk.

Example

Suppose you have a 100 GB file that is truncated. If a process reads a few bytes (at the front or even in the middle of the file), several megabytes of file data are retrieved from the blockpool and the process continues. There is no need for the entire file to be retrieved. If more of the file is read, even larger chunks of the file are retrieved.

You can see the snpolicyd state of a file by using the command "snpolicy -report".

Example

Running the command `snpolicy -report /stornext/sn1/dd1/kcm2` gives us output similar to this:


```
/stornext/sn1/dd1/kcm2
  policy: 18      inode: 1704267
  flags: TRUNC TAG
  mtime: 2010-02-05 11:01:25.960012768 CST
  ingest: 2010-02-05 11:01:25.960012768 CST
  size:      1388936 disk blocks: 0
  seqno:      16 blk seqno:      3
  offset:      0 length:      1388936 tag:
0D4093057370DB4FA7EF8162C23AE416
```

The line beginning with "flags:" contains the keyword TRUNC. This tells us that at least part of the file is not on disk in the file system and must be retrieved to be used.

If only snpolicyd is managing a directory, snpolicyd can truncate files when the snpolicyd rules are satisfied. This means that the deduplication has happened and the file is big enough and perhaps old enough. "Large enough" and "old enough" are determined by the deduplication policy parameters.

If only Storage Manager is managing a directory, the Storage Manager truncation rules determine whether and when a file can be truncated. This usually means that all Storage Manager copies have been made and that the file is large enough and old enough. "Large enough" and "old enough" are determined by the Storage Manager policy parameters.

If *both* Storage Manager and snpolicyd are managing a directory, Storage Manager must do the truncation. Storage Manager can only truncate a file when the Storage Manager rules are satisfied and any snpolicyd data copies have been completed.

You will know that both Storage Manager and snpolicyd are managing a directory if:

- The directory is a deduplicated directory and/or a replication source directory, and the directory is a Storage Manager relation point or is within a Storage Manager relation point.
- The directory is a replication target directory within a Storage manager relation point.

The table below summarizes some of the possibilities for snpolicyd managed directories and when truncation is allowed.

Snpolicyd State of the Directory	Directory is in an SM Relation Point	Directory is <i>Not</i> in an SM Relation Point
Non-deduplication Replication Source	SM can truncate when replications are complete	No truncation
Deduplication Replication Source	SM can truncate when deduplication has happened - even before replication	snpolicyd can truncate after deduplication
Deduplication Without Replication	SM can truncate when deduplication has happened	snpolicyd can truncate after deduplication
Target of Deduplication Source	Files are replicated as truncated (0 blocks). However, SM will eventually store each replicated file, causing it to be retrieved by snpolicyd on the target. Retrieved files must be truncated by SM and can only be truncated after all SM copies are made.	Files are replicated as truncated (0 blocks)
Target of Deduplication Source with "Replicate Deduplicated Content" Off	Files are replicated untruncated and are not tagged (deduplicated). Files can be truncated by SM after all SM copies are made (normal SM rules).	Files are replicated untruncated and are not tagged (deduplicated). Not truncatable.
Target of Dedup source with "Replicate Deduplicated Content" off but deduplication is on in the target policy.	Files are replicated untruncated and are not tagged (deduplicated). Files can be truncated by SM when deduplicated by snpolicyd and stored by SM.	Files are replicated untruncated and are not tagged (deduplicated). Files can be truncated by snpolicyd after deduplication.

The following sections summarize some of the facts above (and add some more information) in a "usage case" or "scenario" format.

Replicating From an SM Relation Point and/or Deduplicating the Relation Point

For a new configuration, create the relation point first. Then make it a replication source by applying an snpolicyd policy with outbound replication enabled.

From the command line, you could use the following commands.

Note: These commands assume that the Storage Manager relation point and replication policy have already been configured.

```
faddrelation directory_path -c sm_policy_name  
spolicy -assignpolicy directory_path -inherit  
replication_policy_name
```

Remember that the directory should be empty before using `fsaddrelation`, or else the command will try to unmount the file system (which is often hard to do).

When a file is both an SM relation point and a replication source, the files cannot be truncated by SM until:

- 1 Either all replications have been completed (non-deduplicated replication)
- OR
- 2 All files in the directory have been deduplicated (deduplicated replication)

If a truncated file is both deduplicated and stored by SM, it can be retrieved by either service. By default we retrieve using `snpolicyd` (from the blockpool) and only use the SM copy if there is an error retrieving from the blockpool.

You can use the `fsretrieve` command to force retrieval from Storage Manager instead of from `snpolicyd`.

Adding Source Replication or Deduplication to an Existing SM Relation Point

The following table summarizes the key points you should consider:

When You Are Making a Directory with Existing SM Managed Files Into This:	Then Expect This:
Snpolicyd deduplication policy (no replication)	<ul style="list-style-type: none"> • Untruncated files are deduplicated per snpolicyd policy. • SM truncated files will not be deduplicated until SM retrieval occurs. (snpolicyd will not retrieve the files from SM.) • Once retrieved from SM, files cannot be re-truncated by SM until deduplication is complete. <p>Therefore files may not all be deduplicated.</p>
Snpolicyd deduplication policy that is also a replication source	<p>About the same as above. SM truncated files are not deduplicated or replicated until something causes SM retrieval of the file.</p> <p>Thus there may be some files not deduplicated and not replicated.</p>
Snpolicyd policy that is a replication source with no deduplication	<p>Similar to above. Files are not replicated until something causes SM retrieval. Once retrieved SM will not truncate the file again until each target of the replication policy has its copy.</p> <p>Not all files will be deduplicated unless retrieved.</p>

Adding Target Replication to an SM Relation Point (New or Existing)

When adding targets within an existing SM relation point, the concepts are a little simpler because a new directory is created each time a

replication occurs. There are no existing files other than previously replicated files.

Remember that you must specify a directory within a Storage Manager relation point when you want replicated files to be stored by Storage Manager.

When replication occurs into a directory in a Storage Manager relation point, the replicated files become SM store candidates (unless they are links to previously replicated files). Storage Manager can then store the files based on age and size. Age is determined by the file's modification time in the source directory because the access and modification times are replicated when a file is replicated.

Storage Manager can store replicated files after they have passed the minimum time, regardless of whether or not they have been truncated by `snpolicyd`. Storage Manager retrieves a truncated file from `snpolicyd` in order to store it to SM tape. Deduplicated replicated files are replicated as truncated files, but they are retrieved by `snpolicyd` when the replication is into a Storage Manager relation point.

Note the following implications:

- 1 This means that more file system space will be used when replicating deduplicated files into an SM relation point than is used when replicating deduplicated files into a directory that is not a relation point. In the latter case there is no StorNext process that will cause the file to be retrieved from the blockpool.
- 2 When the file is retrieved it can be re-truncated after all SM copies have been made. Storage Manager will do the truncation. You can configure the SM policy so that it truncates the file immediately after all SM copies have been made.
- 3 This behavior is different than in the case where we add replication / deduplication to a SM relation point. Truncated files are not automatically retrieved from SM tape so that they can be replicated or deduplicated, but deduplicated files from the blockpool are retrieved so that they can be stored by SM.

Adding Storage Manager to an Existing `snpolicyd` Directory

You cannot add a Storage Manager relation point to an existing replication target directory. You would have to create a new directory, add the SM relation point to that directory, and then create or edit a

snpolicyd policy to realize to a directory or a set of directories inside that relation point.

When adding a Storage Manager relation point to any existing directory, one of the following must be true:

- 1 The directory must be empty.
- 2 You must be able to temporarily unmount the file system. (It gets remounted as part of the add relation point process.)

If the directory is empty there is nothing to worry about. If it is not empty you must make sure no process is working in the directory and no files are open. The directory should not be NFS exported or Samba shared.

Once the relation point has been added, Storage Manager makes copies of the files according to the Storage Manager policy settings. As mentioned earlier, Storage Manager retrieves a file from the blockpool if it needs to in order to store the file.

The snpolicyd Debug Log

A log of snpolicyd actions and errors is maintained in directory `/usr/cvfs/debug`. The log file is named `snpolicy.out`. Previous versions of the log file are called `snpolicy.out.1`, `snpolicy.out.2`, and so on.

Various debugging options can be enabled with the `snpolicy` command. For example, the command `snpolicy -debug=/stornext/sn1 -dflags=events,replicate`

turns on debug messages for events processed by snpolicyd and for replication related activity. The `-debug=` option specifies any file system managed by snpolicyd (any file system with replication / deduplication enabled).

You can find the list of possible dflags options by using the following command:

```
snpolicy -helpdebugflags
```

Here is some sample snpolicyd.out log for an ongoing replication:

```
(D) [0209 17:22:20.903918 3398] Sending rep_realize %H/%P
(D) [0209 17:22:20.934098 3398] release_pending_rep_locked@1109
0x14f86e0 ref 1 state started
(D) [0209 17:22:20.934552 18275] release_rep_target_locked
126540130333 ref 0 state sending metadata
(D) [0209 17:22:20.934582 18275] release_rep_target_locked@827
0x14f86e0 ref 1 state started
(D) [0209 17:22:20.934597 18275] process successful replication,
cnt 9/9 space 1996232
(D) [0209 17:22:20.937694 18276] /stornext/sn3: replication reply
for key 1704023 stream 126540130333
(D) [0209 17:22:20.937720 18276] /stornext/sn3: metadata for '/
stornext/sn3/rep5' accepted by target://stornext/sn4@kcm-
rhe15464:
(D) [0209 17:22:20.938490 18276] update_rep_target_file
126540130333 0 => 3
(I) [0209 17:22:23] /stornext/sn3: data replication for '/snfs/
sn3/rep5'
                                completed to target://stornext/sn4@kcm-
rhe15464: in 2.276911 secs
                                9/9 files (Data/Meta) updated
                                1949 Kbytes in 1.741705 secs 1952/1 Kbytes sent/
received
(D) [0209 17:22:23.252035 18276] post_process_pending_replication
for stream 126540130332
(D) [0209 17:22:23.321761 18276] update_rep_target_file
126540130333 3 => 4
(D) [0209 17:22:23.328021 18276] release_rep_target_locked
126540130333 ref 0 state completed
(D) [0209 17:22:23.328088 18276] release_rep_target_locked@827
0x14f86e0 ref 1 state started
(D) [0209 17:22:23.328109 18276] Freed target stream 126540130333
```




Appendix C

High Availability Systems

The StorNext High Availability (HA) feature allows you to configure and operate a redundant server that can quickly assume control of the StorNext file systems and management data in the event of certain software, hardware and network failures on the primary server.

This appendix contains the following topics which provide an in-depth look at HA systems and operation:

- [High Availability Overview](#)
- [HA Internals: HAmon Timers and the ARB Protocol](#)
- [Configuration and Conversion to HA](#)
- [Managing High Availability in the StorNext GUI](#)
- [Configuring Multiple NICs](#)
- [High Availability Operation](#)
- [HA Resets](#)
- [HA Tracing and Log Files](#)
- [FSM failover in HA Environments](#)
- [Replacing an HA System](#)

High Availability Overview

The primary advantage of an HA system is file system availability, because an HA configuration has redundant servers. During operation, if one server fails, failover occurs automatically and operations are resumed on its peer server.

At any point in time, only one of the two servers is allowed to control and update StorNext metadata and databases. The HA feature enforces this rule by monitoring for conditions that might allow conflicts of control that could lead to data corruption.

Before this so-called Split Brain Scenario would occur, the failing server is reset at the hardware level, which causes it to immediately relinquish all control. The redundant server is able to take control without any risk of split-brain data corruption. The HA feature provides this protection without requiring special hardware, and HA resets occur only when necessary according to HA protection rules.

Arbitration block (ARB) updates by the controlling server for a file system provide the most basic level of communication between the HA servers. If updates stop, the controlling server must relinquish control within a fixed amount of time. The server is reset automatically if control has not been released within that time limit.

Starting after the last-observed update of the ARB, the redundant server can assume control safely by waiting the prescribed amount of time. In addition, the ARB has a protocol that ensures that only one server takes control, and the updates of the ARB are the method of keeping control. So, the ARB method of control and the HA method of ensuring release of control combine to protect file system metadata from uncontrolled updates.

Management data protection builds on the same basic HA mechanism through the functions of the special shared file system, which contains all the management data needing protection. To avoid an HA reset when relinquishing control, the shared file system must be unmounted within the fixed-time window after the last update of the ARB. Management data is protected against control conflicts because it cannot be accessed after the file system is unmounted. When the file system is not unmounted within the time window, the automatic HA reset relinquishes all control immediately.

The HA system monitors each file system separately. Individual file systems can be controlled by either server. However, StorNext Storage Manager (SNSM) requires that all managed file systems be collocated with the management processes. So, the shared file system and all managed file systems are run together on one server. Unmanaged file systems can run on either server, and they can fail over to the other server as long as they perform failover according to the HA time rules described above.

When it is necessary to make configuration changes or perform administrative functions that might otherwise trigger an HA reset, snhamgr, the HA Manager Subsystem (patent pending), provides the necessary controls for shutting down one server and operating the other server with HA monitoring turned off. Snhamgr allows the individual servers to be placed in one of several modes that regulate starting StorNext software on each server. The restricted pairing of server modes into allowed cluster states provides the control for preventing Split Brain Scenario. The HA Manager Subsystem uses communicating daemons on each server to collect the status of the cluster at every decision point in the operation of the cluster. This is another one of the levels of communication used in the HA feature.

An occasional delay in accessing the SAN or its disks might trigger an HA reset while the server and File System Manager (FSM) are otherwise functioning correctly. A LAN communication protocol between the servers' File System Portmapper (FSMPM) processes reduces the chance of a server reset by negotiating the reset of HA timers (patent pending) outside of the ARB-update timer-reset system.

When SAN delays are causing undesirable HA resets, the causes of the delays must be investigated and resolved. Quantum support staff can increase the timer duration as a temporary workaround, but this can negatively impact availability by increasing the time required for some failover instances.

The set of features comprising StorNext HA provides a highly automated system that is easy to set up and operate. The system acts autonomously at each server to continue protection in the event of LAN, SAN, disk and software failures.

The timer mechanism operates at a very basic level of the host operating system kernel, and is highly reliable. Protection against Split Brain Scenario is the primary requirement for HA, and this requires the possibility of some unnecessary system resets. But, when communication channels are working, steps are taken to reduce the

number of unnecessary resets and to eliminate them during administrative procedures.

HA Internals: HAmom Timers and the ARB Protocol

Control of StorNext file system metadata is regulated through the ARB dedicated disk block. The protocol for getting and keeping control of the ARB is meant to prevent simultaneous updates from more than one FSM. The protocol depends on timed updates of the ARB, which is called “branding”.

Loss of control of the timing of branding opens the possibility of metadata corruption through split-brain scenario. The extra protection provided by HAmom timers puts an upper limit on the range of timing for ARB brand updates. Brand updates and HAmom timer resets are synchronized. When branding stops, the timer can run out and trigger an HA reset.

When taking control, an FSM uses the same timer value plus a small amount starting from the last time it read a unique brand. This combination of behaviors provides a fail-safe mechanism for preventing split-brain scenario metadata corruption.

FSM Election, Usurpation and Activation

When a client computer needs to initiate or restore access to a file system, it contacts the nameserver-coordinator system to get a LAN port for the controlling FSM. The nameserver-coordinator system will conduct an election if there is no active FSM or the active FSM is no longer healthy.

This measures the connectivity between the possible server computers and the clients. The nameserver-coordinator system uniquely chooses one standby FSM to take control, and sends an activation command to it. At this point, the cvadmin command will display an asterisk next to the FSM to show that the FSM has been given an activation command.

The elected FSM begins a usurpation process for taking control of the file system metadata. It reads the ARB to learn about the last FSM to control the file system. It then watches to see if the brand is being

updated. If the brand is not being updated or if the usurping FSM has more votes than the current controlling FSM has connections, the usurper writes its own brand in the ARB. The FSM then watches the brand for a period of time to see if another FSM overwrites it. The currently active FSM being usurped, if any, will exit if it reads a brand other than its own (checked before write). If the brand stays, the FSM begins a thread to maintain the brand on a regular period, and then the FSM continues the process of activation.

At this point the usurping FSM has not modified any metadata other than the ARB. This is where the HAMon timer interval has its effect. The FSM waits until the interval period plus a small delta expires. The period began when the FSM branded the ARB. The FSM continues to maintain the brand during the delay so that another FSM cannot usurp before activation has completed. The connection count in the ARB is set to a very high value to block a competing usurpation during the activation process.

When an FSM stops, it attempts to quiesce metadata writes. When successful, it includes an indicator in its final ARB brand that tells the next activating FSM that the file system stopped safely so the wait for the HA timer interval can be skipped.

LAN Connectivity Interruptions

When one MDC loses LAN connectivity, clients lose access to that MDC's active FSMs, which triggers elections to find other FSMs to serve those file systems. StorNext attempts to determine which node should have control, based on connectivity, but this effort results in a tie for the HaShared file system because each node gets one vote from itself as a client. In a tie, the activated shared FSM keeps control so long as it keeps branding its ARB.

Managed FSMs are not redundant, so having clients on those file systems does not break the tie. Similarly, unmanaged FSMs can fail over without an HA reset, so clients on those file systems will not break the tie for the shared file system either.

Therefore, a third client that has the shared file system mounted is necessary to break the tie that occurs between the two nodes. The third client makes it possible to determine which of the MDCs has the best connectivity to the LAN.

Note: The third-party client is not necessary for preventing metadata corruption from split brain syndrome. The ARB plus the HAMon timer to back it up does the whole job of protecting the metadata. For more information about HAMon timer, see the following section.

Autonomous Monitoring and HA Resets

When an HA reset is necessary, it occurs before usurpation could complete. This is true because the start of the timer is based on the last update of the ARB brand for both the active and activating FSMs. Brand updating is the only communication between server computers that is necessary for HA protection against split-brain scenario.

Note that there is no communication from an activating FSM to force an HA reset at its peer server computer. The two servers act autonomously when the ARB branding communication stops. The combination of an HA reset when the brand cannot be maintained and the usurpation-branding protocol guarantees protection from split-brain scenario.

Note: There could be a delay between the autonomous HA reset by the active FSM's server and the election of another FSM to take control. These are not synchronized except by the election protocol.

In the original method of HA resets, known as STONITH, a usurping FSM synchronously reset the peer server at the end of usurpation before modifying any metadata. There was no method for the failing FSM's server to reset itself, so the reset would not happen if the usurping FSM failed to brand the ARB. However, the STONITH method always reset its peer the first time any FSM activated on a server regardless of whether it was needed or not.

Both of these strategies result in some unnecessary system resets where an omniscient system would not reset, but the current HA system has fewer of these resets.

Setting the Timer Value

The HAMon timer interval can be changed to work around delays in the access to ARB because of known behavior of a particular SAN deployment. The feature is meant for temporary use only by Quantum

staff. It affects all the monitored FSMs and could add a significant delay to the activation process. Quantum Software Engineering would like to be notified of any long-term need for a non-default timer interval.

For very long HAmom interval values, there are likely to be re-elections while an activating FSM waits for the time to pass before completing activation. An additional usurpation attempt would fail because the ARB brand is being maintained and the connection count is set to a value that blocks additional usurpation attempts.

The optional configuration of this feature is in the following file:

```
<cvfs root>/config/ha_smith_interval
```

The information at the start of the file is as follows:

```
ha_smith_interval=<integer>
```

The file is read once when StorNext starts. The integer value for the HAmom timer interval is expressed in seconds. The value can range from 3 to 1000, and the default is 5 seconds. The timer must be set identically on both servers. This rule is checked on a server that has standby FSMs when a server that has active FSMs communicates its timer value. When there is a discrepancy, all the FSMs on the receiving end of that communication are stopped and prevented from starting until StorNext has been restarted. This status can be observed with the cvadmin tool in the output of its FSMlist command.

In almost all cases of misconfigured timers, the mistake will be obvious shortly after starting the HA cluster's second server. The first server to start StorNext will activate all of its FSMs. The second server should have only standby FSMs. Once the second server detects the error, all of its FSMs will stop. After this, there will be no standby FSMs, so the cluster is protected against split-brain scenario. In the event that a server with active FSMs resets for any reason, that server will have to reboot and restart StorNext to provide started FSMs to serve the file systems.

Negotiated Timer Resets

When an FSM is healthy but cannot maintain its brand of the ARB because of delays in the SAN or LUN, there is the possibility of an undesirable HA reset. To address this problem there is a LAN-based negotiation protocol between FSMPM processes on the two servers for requesting permission to reset HAmom Timers.

The negotiation is initiated by an FSMPM on a server computer with activated FSMs. Every two seconds it sends a list of active FSMs to its

peer FSMPM on the other server to ask which of these standby FSMs are not being activated. Implicit in the response is a promise not to activate the FSMs for two seconds. When the response is received within one second, the first FSMPM resets the timers for those FSMs for which usurpation is not in progress. Obviously, both server computers must be up and running StorNext for this to function.

This can postpone the impending HA reset for a while, but an election could occur if this goes on too long. It is important to quickly investigate the root cause of SAN or LUN delays and then engineer them out of the system as soon as possible.

Primary and Secondary Server Status

Databases and management data for StorNext Storage Manager or the Linux GUI must also be protected against split-brain scenario corruption. Protection is accomplished by tying the startup of processes that modify this data with the activation of the shared file system.

Activating the shared file system leads to setting a Primary status in the local FSMPM, which is read and displayed by the `snhamgr` command. Primary status and the implicit Secondary status of the peer server are distinct from the Active and Standby status of the individual FSMs on the servers.

Unmanaged file systems can be active on either server. When an HA Cluster has no managed file systems and no shared file system, neither server computer has Primary status—they are equals.

File System Types

HA is turned on by default for all StorNext distributions, but has no effect unless FSMs request to be monitored. File system monitoring is controlled by a file-system configuration item named `HaFsType`. Each file system is one of three types: `HaUnmanaged`, `HaManaged` or `HaShared`. The `HaFsType` value is read by FSMs to direct them to set up appropriate HAMon behaviors, and it is read by the FSMPM to control how it starts FSMs.

HaUnmanaged

Each unmanaged-file-system FSM starts an instance of the HAMon timer for itself when it first brands the ARB. Before it changes any metadata, an activating FSM waits for the timer interval plus a small amount of

time to elapse. The interval for a usurping FSM begins with the last time the FSM reads new data in the ARB from a previously active FSM.

Unmanaged FSMs can be active on either server in the HA Cluster. They can be usurped and fail over without a system reset if they exit before the timer expires. The timer interval for an active FSM restarts with each update of the ARB.

HaManaged

Managed-file-system FSMs do not start HAmon timers, and they do not wait the HAmon interval when usurping. The FSMPMs only start Managed FSMs on the Primary server, so there is no risk of split-brain scenario. In the event that a Managed FSM exits without having been stopped by the FSMPM, it is automatically restarted after a ten-second delay and activated. The cvadmin tool's `FSMlist` command displays the blocked FSMs on non-Primary servers. There can be zero or more HaManaged file systems configured.

HaShared

The shared file system is an unmanaged StorNext file system that plays a controlling role in protecting shared resources. It has the same HA behavior as other unmanaged FSMs, but it also sets a flag that triggers an HA reset when the `cvfsioctl` device is closed. This happens when the process exits for any reason. However, if the shared file system has been unmounted from the active server before the FSM exits, the reset-on-close flag gets turned off. This allows ordinary shutdown of CVFS and Linux without resetting the server.

When the HaShared FSM finishes activation, it sets the Primary status in its FSMPM process.

Protected shared data resides on the shared file system. Since only one FSM can activate at one time, the Primary status is able to limit the starting of management processes to a single server, which protects the data against split-brain scenario.

The starting of HaManaged FSMs is also tied to Primary status, which guarantees collocation of the managed file-system FSMs and the management processes. The GUI's data is also shared, and the GUI must be able to manipulate configuration and operational data, which requires that it be collocated with the management processes.

The `ha_peer` and `fsnameservers` File

StorNext HA server software uses peer-to-peer communication between servers and needs to know the peer's IP address. The `fsnameservers` configuration file is not a good source for the address because some installations configure the nameservers outside of the metadata servers. Instead, the following file provides that information:

```
<cvfs root>/config/ha_peer
```

Following are the uses of the peer IP address:

- Negotiating timer resets
- Comparing the HAmom timer value between servers
- HA Manager communications (only on StorNext Storage Manager for Linux)

It is very important to have correct information in the `ha_peer` file, but it is not a requirement that the peer be available for communication. Basic HA functionality operates correctly without IP communication between peers. The file's contents can be changed without restarting StorNext. The new value will be read and used moments after it has changed.

Here are some other points to consider about `fsnameservers`:

- For best practice, the `fsnameservers` file should contain IP addresses, not names.
- All the addresses in the file must be reachable by all members of the StorNext cluster. That is, servers, clients and distributed LAN clients.
- All members of the cluster should have the same nameservers configuration.
- Both or neither of an HA Cluster's MDCs must be included so that a coordinator is always available when either server is running.
- Multiple StorNext Clusters can share coordinators, but every file system name configured on any of the clusters must be unique across all of the clusters.

HA Manager

The HA Manager subsystem collects and reports the operating status of an HA cluster and uses that to control operations. It is part of a Storage Manager installation that has been converted to HA with the `cnvt2ha.sh` script. For manually-configured HA clusters where the

`cnvt2ha.sh` script has not been run, the command-line interface (`snhamgr`) reports a default state that allows non-HA and File System Only HA configurations to operate.

The HA Manager supports non-default HA Cluster functionality such as suspending HA monitoring during administrative tasks. It attempts to communicate with its peer at every decision point, so it is mostly stateless and functions correctly regardless of what transpires between decision points. Following every command, the `snhamgr` command line interface reports the modes and statuses of both servers in the cluster, which provide necessary information for the StorNext control scripts.

HA Manager Modes and Statuses

The HA Manager relies on a set of administrator-setable modes to override the default behaviors of HA. Modes persist across reboots. Following are the modes and descriptions of their purpose:

- 1 **default:** HA monitoring is turned on. When the peer server is not available for communication, it is assumed to be in default mode.
- 2 **single:** HA monitoring is turned off. The peer server must be communicating and in locked mode, or not communicating and certified as peerdown (recommended). This mode is meant for extended production operations without a redundant server such as when one server is being repaired or replaced. When the peer server is about to be restored to service, the operating server can be transitioned from single to default mode without stopping StorNext.
- 3 **config:** HA monitoring is turned off. The peer server must be communicating and in locked mode (recommended), or not communicating and certified as peerdown. The config mode is meant for re-configuration and other non-production service operations. When returning to production service and the default mode, StorNext must be stopped. This ensures that all StorNext processes are started correctly upon returning to default mode.
- 4 **locked:** StorNext is stopped and prevented from starting on the local server. This mode allows the HA Manager to actively query the peer server to ensure that it is stopped when the local peer is operating in single or config mode. Communication with the locked node must continue, so this mode is effective when StorNext is stopped for a short period and the node will not be rebooted. If

communication is lost, the peer node assumes this node is in default mode, which is necessary for avoiding split-brain scenario.

- 5 **peerdown:** The peer server is turned off and must not be communicating with the local server's HA Manager subsystem, so this mode is effective when the server is powered down.

The mode is declared by the `peerdown` command on a working server to give information about the non-working peer server. By setting this mode, the administrator certifies the off status of the peer, which the HA Manager cannot verify by itself. This allows the local peer to be in single or config mode. If the peer starts communicating while this mode is set, the setting is immediately erased, the local mode is set to default to restore HA Monitoring, and StorNext is shut down, which can trigger an HA reset.

The `peerdown` mode is changed to default mode with the `peerup` command. The `peerdown` and `peerup` commands must never be automated because they require external knowledge about the peer server's condition and operator awareness of a requirement to keep the peer server turned off.

- 6 **ha_idle_failed_startup:** A previous attempt to start StorNext with 'service cvfs start' has failed before completion. Attempts to start StorNext are blocked until this status has been cleared by running 'snhamgr clear'.

The HA Manager subsystem collects server statuses along with the server modes to fully measure the operating condition of the HA Cluster. The possible statuses are as follows:

- **stopped:** Running the 'DSM_control status' command has returned a false code.
- **running:** Running the 'DSM_control status' command has returned a true code.
- **primary:** The server's status is running and the FSMPM is in the primary state, which indicates that the HaShared FSM has been activated.

The HA Manager allows the cluster to be in one of the following restricted set of operating states. When a server is in default mode, HA monitoring is turned on.

- **default-default**
- **default-locked**

- **default-peerdown**
- **single-peerdown**
- **single-locked**
- **config-peerdown**
- **config-locked**
- **locked-***

The following states are prohibited and prevented from occurring by the HA Manager, unless there is improper tampering. For example, the last state listed below (peerdown-*), is the case when a node that is designated as peerdown begins communicating with its peer. If any of these is discovered by the HA Manager, it will take action to move the cluster to a valid state, which may trigger an HA reset.

- **single-default**
- **single-single**
- **single-config**
- **config-default**
- **config-single**
- **config-config**
- **peerdown-***

HA Manager Components

The following files and processes are some of the components of the HA Manager Subsystem:

- **snhamgr_daemon**: If the `cnvt2ha.sh` script has been run, this daemon is started after system boot and before StorNext, and immediately attempts to communicate with its peer. It is stopped after StorNext when Linux is shutting down. Otherwise, it should always be running. A watcher process attempts to restart it when it stops abnormally. Its status can be checked with 'service snhamgr status'. It can be restarted with 'service snhamgr start' or 'service snhamgr restart' if it is malfunctioning.
- **snhamgr**: CLI that communicates with the daemon to deliver commands and report status, or to report a default status when the

`cnvt2ha.sh` script has not been run. This is the interface for StorNext control scripts to regulate component starts.

- **`/usr/cvfs/install/.ha_mgr`**: Stored mode value, which allows the single, config, locked, and peerdown modes to persist across reboots.
- **`SNSM_HA_CONFIGURED`**: Environment variable that points to a touch file to indicate that `cnvt2ha.sh` has been run.
- **`/etc/init.d/snhamgr`**: Service control script for starting the `snhamgr_daemon`.
- **`HA_IDLE_FAILED_STARTUP`**: Environment variable that points to a touch file to indicate that a previous run of `'service cvfs start'` failed before completion. This blocks startup attempts to prevent infinitely looping startup attempts.
- **`/usr/cvfs/debug/smithlog`**: When an HA Reset is imminent, a descriptive line is added to the end of this file and the file is `fsync'd` in an attempt to ensure that the information is available for debugging the root cause of the reset. For example, when there is less than one second remaining on the HA Monitor timer, a notice is written in this file. It is likely that all other log files will lose some of the newest information at the time of the reset because it is in buffers that have not been written to disk. The `fsmppm` process writes the file, so the file may not have any diagnostic information in cases where the `fsmppm` itself fails and causes an HA reset.

HA Manager Operation

In addition to the setting of modes, there are some commands provided by the HA Manager to simplify and automate the operation of an HA Cluster.

- 1 **`status`**: Report cluster status. All commands report cluster status on completion. This command does nothing else unless an invalid cluster state is detected, in which case it will take steps to correct the cluster state.
- 2 **`stop`**: Transition the non-Primary server to locked mode, which will stop StorNext if it is running. Then transition the Primary server to config mode, which will turn off HA monitoring. Stop StorNext on the Primary server. Finally, transition both servers to default mode.

- 3 **start:** If either MDC is in config or single mode, transition it to default mode (CVFS stops when transitioning from config mode). If either MDC is in locked mode, transition it to default mode. If the remote MDC is in peerdown mode, then run `peerup`. If the local MDC is stopped, run `'service cvfs start'`. If the remote MDC is accessible and in stopped status, run `'service cvfs start'` on the remote MDC.

Note: Running `service cvfs start` is the preferred method of starting, so Quantum recommends using this command rather than using other methods. Likewise, use `service cvfs stop` to stop StorNext.

- 4 **config:** Transition the peer server to locked mode, then transition the local server to config mode. This command must be run on the Primary server or either server when both servers are stopped.
- 5 **clear:** Erase the `HA_IDLE_FAILED_STARTUP` touch file. This should be used after correcting the condition that caused the failure of the `'service cvfs start'` command.
- 6 **--primary:** Set the Primary status in the FSMPPM on the local server. This is exclusively for use by system scripts. Improper use of this command will violate HA protection rules and could result in split-brain scenario.
- 7 **force smith:** Trigger an immediate HA reset. This is primarily for use by system scripts.

Configuration and Conversion to HA

The following types of StorNext configurations can be run as HA Cluster servers:

1 Windows

The StorNext GUI has a menu item for configuring HA: **Tools > High Availability > Convert**. It automatically inserts the `haFsType HaUnmanaged` configuration item in every file system configuration, and restarts the FSMs to enable HA. This menu item must be selected separately on each server. It is the operator's responsibility

to ensure that HA is running on both servers; there are no built-in tests to ensure that this has been done. The `ha_peer` file must be created manually to contain the IP address of the peer MDC, which is used for negotiated restarts of HA timers to avoid unnecessary HA resets.

For more information about using the Convert menu option, see the StorNext online help.

2 Linux SNFS without GUI support

Each FSM configuration file must be given the `haFsType` `HaUnmanaged` configuration item, and the `ha_peer` file must be given the numerical IP address of its peer, which is used for negotiated restarts of HA timers to avoid unnecessary HA resets.

The FSM configuration files and `fsnameservers` files must be identical on both servers. When these things are done correctly, HA Monitoring is protecting the metadata against split-brain scenario. It is on by default; there is no means for turning it off other than removing the `haFsType` item from the FSM configuration files.

3 Linux Storage Manager with only unmanaged file systems

See [Conversion to HA](#) for more information.

4 Linux Storage Manager with managed and unmanaged file systems

See [Conversion to HA](#) for more information.

Conversion to HA

This section describes what happens in the conversion script.

When a StorNext Storage Manager single-server configuration has been completed, the node is converted to HA by running the `/usr/adic/util/cnvt2ha.sh` script.

Before running the script, add the `haFsType` configuration item to each file system according to its type (`HaUnmanaged`, `HaManaged` or `HaShared`), and enter the peer MDC's IP address in the `/usr/cvfs/config/ha_peer` file.

This script expects there to be one and only one unmanaged file system that is configured with the `haFsType` `HaShared` configuration item. It also expects the `/usr/cvfs/config/license.dat` file to include licenses for both the configured server and its unconfigured redundant

peer. The configured server is running StorNext when the `cnvt2ha.sh` script is run, so it becomes the Primary on completion of the script. (The StorNext GUI, if used to drive the conversion to HA, creates the mergedfile before converting the secondary.)

The following command is invoked to start the conversion process:

```
/usr/adic/util/cnvt2ha.sh primary
```

Output for the operation is displayed to the screen and saved to the `/usr/adic/HA/cnvt2ha.sh.log` file.

The script automatically finds the shared file system by its `haFsType` configuration item and moves its mount point to `/usr/adic/HAM/shared`. It relocates several hundred configuration and database files for Storage Manager to the `/usr/adic/HAM/shared` directory. SNFS configuration items are copied from the `/usr/cvfs/config` directory to the mirror subdirectory of the shared file system. Finally, it creates the following touch file to indicate that conversion has completed:

```
/usr/adic/install/.sasm_ha_configured
```

The existence of that file enables the running of the `sasmgr` service, which starts the HA Manager daemon.

Before the conversion script is run on the secondary, the following file must be copied from the Primary:

```
/usr/cvfs/config/fsnameservers
```

The arguments to the conversion command for the secondary server are as follows:

```
/usr/adic/util/cnvt2ha.sh secondary    <sharedfs name>  
<peer IP address>
```

This gives the redundant peer server enough information to access the shared file system as a client. It then copies the mirrored configuration files into its own configuration directory and sets up the `ha_peer` file. The database and management-components configuration files are rerouted to the `/usr/adic/HAM/shared` shared file system mount point. Finally, the `.sasm_ha_configured` touch file is created, and StorNext is restarted.

SyncHA process

Before the shared file system is up, some configuration files must be available. These are initially “mirrored” to the secondary server by the `cnvt2ha.sh` script, and then maintained across the two server computers by the `syncHA` process, which is run once per minute from `cron`.

On the Primary, the command `stat`’s the mirrored files to see what has changed, and copies these out to the `/usr/adic/HAM/shared/mirror` folder. On the secondary server, the files that have changed are copied in. The list of mirrored files is defined in the `/usr/cvfs/config/filelist` and `/usr/adic/gui/config/filelist` tables as follows.

In the `/usr/cvfs/config` directory:

- `license.dat`
- `fsmlist`
- `fsnameservers`
- `fsroutes`
- `fsports`
- `*.cfg`
- `*.cfgx`
- `*.opt`
- `nss_ctl.xml`
- `snpolicyd.conf`
- `blockpool_settings.txt`
- `blockpool_root`
- `blockpool_config.tpl`
- `blockpool_config.txt`
- `bp_settings`

In the `usr/adic/gui` directory:

- `database/derby_backup.tar`
- `logs/jobs/*`
- `config/host_port.conf`


Note: By default, the syncha process backs up the internal state of the StorNext GUI every minute, which may cause a performance impact to some GUI operations. The file `/usr/adic/gui/config/syncha_interval.conf` can be used to reduce the frequency of the GUI state backups. The contents of the file, if present, should contain an integer value that specifies the minimum number of seconds between GUI state backups. This file is host dependant and applies only to syncha when it is run in cron mode on the primary system.

Managing High Availability in the StorNext GUI

The StorNext Tools menu's **High Availability > Manage** option enables you to view the current status of the file systems on your High Availability (HA) system. Specifically, you can view the primary and secondary node FSM statuses: Active, Standby, or Unknown.

The Manage option is accessible by choosing **High Availability > Manage** from the Tools menu. The **High Availability Manage** screen appears.

Figure 131 High Availability
Manage Screen



The Manage option also enables you to perform the following HA-related actions:

Enter Config Mode: Sets the peer (secondary) node to locked mode and sets the local (primary) node to config mode for administration purposes. The locked mode stops CVFS on the peer, and is designed for automated short-duration stops of the secondary server to make configuration changes and other modifications. This allows the HA Manager to prevent HA resets while making configuration changes or stopping the primary server.

Note: In the event that TCP communication to the secondary server is lost for any reason, the primary server assumes the secondary server is in default mode and transitions the local server out of config mode. For this reason, the locked mode is not appropriate to use for extended secondary-server outages, activities that might include reboots of the secondary server, etc. Best practice is to use Peerdown mode when a server is turned off for an extended period, or to simply keep the primary server in default mode while the secondary server is brought in and out of service in short durations.

- 1 Click **Enter Config Mode**.
- 2 When the confirmation message appears, click **Yes** to proceed or **No** to abort.
- 3 Click **OK** when a message informs you that the cluster was locked.

Exit Config Mode: Starts both nodes of the HA cluster in default mode.

- 1 Click **Exit Config Mode**.
- 2 When the confirmation message appears, click **Yes** to proceed or **No** to abort.
- 3 Click **OK** when a message informs you that the cluster was unlocked.

Configuring Multiple NICs

StorNext supports using a multiple NIC (multihomed) configuration as a solution for adding metadata network redundancy. This section describes this configuration and provides an example.

LAN Configuration

The metadata network connects all StorNext nodes in a StorNext cluster. This critical network infrastructure can be configured for redundancy to enhance the availability of the cluster.

Two or more network segments may be used to transport metadata. In an HA environment two redundant MDCs are configured. Each MDC is connected to each of the metadata networks over a separate interface.

The MDCs run the File System Managers (FSMs) for the cluster. A cluster also needs to designate nodes to serve as the name servers.

These nodes may be an HA MDC pair, but may be separate dedicated nodes.

In the case that there are two name servers and each name server is reachable via two networks, the `fsnameservers` file would contain the four addresses through which the two name servers can be reached.

The StorNext NSS protocol is used by nodes with the cluster to discover the locations of the file system servers and the metadata network address used to reach the service.


Even though an FSM may be reached over multiple addresses, only one FSM address is advertised to the clients for a given file system.

This means that the redundant addresses for an FSM service are used for availability but not load sharing.

In the event that the network containing the address being advertised for the FSM service fails, the backup network address will start being advertised as the location of the FSM for that file system.

Clients which have the file system mounted will automatically reconnect, if their TCP connection was terminated. Note that depending on the nature of the failure, it is possible that existing TCP connections will be maintained. The client TCP mount connection to the FSM will only be re-tried if it is disconnected, not simply because the FSM service begins advertising at a new address.

Example Configuration



Node A:
eth0: 192.168.100.1
eth1: 192.168.200.1

Node B:
eth0: 192.168.100.2
eth1: 192.168.200.2

fsnameservers file on both nodes:
192.168.100.1
192.168.100.2
192.168.200.1
192.168.200.2

High Availability Operation

Most of the information in this section is in regard to GUI-supported configurations of StorNext on Linux servers; that is, those installations having an HaShared FSM. There is very little difference for File System-only installations on Windows or Linux in administrating redundant HA versus non-HA servers.

The supported method for starting StorNext on Linux is the 'service cvfs start' command. This is the method used automatically by Linux when the system enters multi-user mode. The script sets up a failure-detection method that prevents looping starts as described in [HA Manager Components](#) on page 333.

StorNext is automatically started as a service on Windows. If StorNext is started more than once in a three-minute period, StorNext operation is delayed for three minutes. This would allow an administrator to login and stop an infinite cycle of HA resets at startup.

StorNext Server for Windows includes the Advanced File System Configuration tool that automates the configuration of the HaFsType parameter. Ensure that both servers have HA enabled when redundant servers are operating. The ha_peer file must be manually configured.

Note: The spelling is haFsType for XML configuration, and HaFsType for the old .cfg configuration methods. Windows uses the .cfg method exclusively. Linux uses XML exclusively for the StorNext GUI, but the .cfg method is still supported for by-hand configuration.

Windows and Linux SNFS Installations Without the HaShared File System

HA monitoring is turned on by default when FSM configurations include the HaFsType configuration parameter. There is no need to disable HA in almost all cases. The only mechanism for turning it off is to remove the configuration parameter, but this should be done only after the redundant server has been turned off.

Note: Instances of redundant StorNext servers without HA are not supported.

The ha_peer and optional ha_smith_interval files are the only additional configuration items for instances of HA without an HaShared file system. These items must be manually configured. FSM-configuration and FSMlist files must be identical on both servers.

When stopping StorNext on one of these servers, all FSMs will stop. Standby FSMs on the redundant server will activate and resume serving their file systems. No HA reset will occur if every stopping FSM exits before their HAMon timer expires (after the final ARB brand update).

During operation, individual file systems could potentially fail over between servers as the result of a hardware or software failure or because an operator has directed it by the `fail` command in the `cvadmin` tool. The `fail` command can be used for load balancing after the HA cluster has completed startup.

When making file system configuration changes, one of the servers should be stopped and its FSM configurations deleted. This eliminates the possible mistake of asymmetric configurations. After making all the configuration changes on one server and updating the file systems with those new configurations, the configuration files must be copied to the redundant server. Then, the cluster can be operated with redundant servers again.

When updating StorNext software, refer to the release notes and the StorNext Upgrade Guide for current update instructions. These documents address any special considerations for HA according to the scope of the changes in the software release.

Linux SNMS and SNFS Installations with the HaShared File System

The HaShared file system is required for SNMS and GUI-supported installations. The shared file system holds operational information for those components, which must be protected against split-brain corruption. The additional complexity this entails is simplified and automated by the HA Manager Subsystem.

Touch Files that Control StorNext and HA

Environment variables defined in the `/usr/adic/.profile` and `/usr/adic/.cshrc` files reference touch files that are used in StorNext to track state transitions. Some of these are unique to HA configurations. The variables and their values are as follows:

- 1 `ACTIVE_SNFS_SERVER=/usr/adic/install/.active_snfs_server`

The file is created following activation of the HaShared file system to designate the local node as the Primary server.

- 2 `HA_STARTING_UP=/usr/cvfs/install/.ha_starting_up`

The file is created at the start of the `'service cvfs start'` script, and removed at the end of the activation script for the HaShared file system. If it is found before it is created, that causes the creation of the `HA_IDLE_FAILED_STARTUP` file as described in the next item.

- 3 HA_IDLE_FAILED_STARTUP=/usr/cvfs/install/
.ha_idle_failed_startup

When the HA_STARTING_UP file exists as the 'service cvfs start' starts, it is assumed that the previous attempt to start failed before completing, possibly because an HA reset occurred. The HA_IDLE_FAILED_STARTUP file is created to block future attempts to start, and the current script exits. This avoids an infinitely looping series of startup attempts, and allows an administrator to log in and correct problems. The HA Manager reports the existence of this file as a mode, and offers the clear command for removing the file.

- 4 SNSM_HA_CONFIGURED=/usr/adic/install/
.snsm_ha_configured

The file is created by the cnvt2ha.sh script to indicate that the system has been converted to HA. Its existence allows the snhamgr_daemon to run.

- 5 START_SNFS_ONLY=/usr/adic/install/.start_snfs_only

The file is created by running one of the following commands: '/usr/adic/bin/adic_control startonly snfs' or '/usr/cvfs/bin/DSM_control startonly'.

Its existence indicates to the snactivated script that Storage Manager components are not to be started. The file is removed by using any of the following commands: 'DSM_control stop', 'service cvfs stop', or 'adic_control stop snfs'.

Starting a Single StorNext HA Server for Production

The 'service cvfs start' command sets in motion a sequence of events that result in the starting of all the Storage Manager components.

Note: The individual Storage Manager component scripts should not be run by hand. There are safeguards in the control scripts to preserve the HA protections against split-brain scenario in any case, but StorNext can get into certain states that are tricky to reconcile if component scripts are used in the wrong sequence. The shared file system can make that reconciliation more difficult.

The `cvfs` script (indirectly) starts the `DSM_control` script, which starts the `FSMPM`, waits for it, and then repeatedly attempts to mount all of the `cvfs` type file systems. The `FSMPM` reads the FSM configuration files and the `fsmlist` file. It starts the `HaShared` and `HaUnmanaged` FSMs in the `fsmlist`, but delays starting the `HaManaged` FSMs. The sub state of the delayed FSMs can be displayed with the `fsmlist` command in the `cvadmin` tool. Meanwhile, the mounts taking place because of the action of `DSM_control` are triggering elections that are activating the locally started FSMs if they are not already being serviced by active FSMs on the peer server.

When an FSM completes activation, it runs the `snactivated` script. The script for the `HaShared` file system creates the `ACTIVE_SNFS_SERVER` file, and then calls '`snhamgr --primary`' to set the Primary status for this server. That induces the `FSMPM` to start the `HaManaged` FSMs. The `HaShared` activation script waits a limited time for all of the managed file systems to be mounted, and then it calls '`adic control start`' to start the other Storage Manager components. Finally, the `HaShared` activation script removes the startup-failure-detection touch file.

While all this is happening, the `DSM_control` script is monitoring progress and reporting statuses of the mounts and the component startups. It will wait a limited time for completion. When it finishes and exits all the nested scripts and returns to the user, all of the Storage Manager components should be up. But if it times out before that, the background activities should continue bringing up Storage Manager. The results of this can be observed a few moments later.

Starting and Stopping the StorNext HA Cluster

When starting or stopping StorNext HA, it is always helpful to first get the cluster state from the HA Manager as follows:

```
snhamgr status
```

The status output indicates whether one or both servers are stopped, if they are in non-default modes, and if either server has Primary status. The typical first step in stopping an HA cluster is to stop the secondary server and to lock it. This allows the other server to be put in config or single mode to operate with HA monitoring turned off. Then, that server can be stopped without incurring an HA reset. These steps are automated in the following cluster command:

```
snhamgr stop
```

When starting the cluster into production, both servers must be in default mode. The first server to start is likely to have its HaShared FSM activated, which will result in that server becoming Primary. The redundant server becomes Secondary when it starts operation, and its FSM processes wait in Standby until they are elected to usurp control of their file systems. These steps are automated in the following cluster command, which starts the local server, if necessary, to become Primary, followed by starting the Secondary server:

```
snhamgr start
```

StorNext HA also has the ability to stop a Primary server while it is in default mode without incurring an HA reset in most cases. It does this as follows:

- 1 Stop Storage Manager processes, including the database
- 2 Unmount all CVFS file systems on the local server other than the HaShared file system
- 3 Stop all FSMs on the local server other than the HaShared FSM
- 4 Unmount the HaShared file system
- 5 Stop the FSMPPM
- 6 Stop the HaShared FSM

FSMs are elected and activate on the peer server as they are stopped on the local server.

An HA reset can occur if step 4 fails. (That is, if the HaShared file system cannot be unmounted for any reason.) This is the method for protecting Storage Manager management data against split-brain-scenario corruption. All of the at-risk data is contained on the shared file system, so the unmount operation ensures that the local server cannot modify the data.

Upgrading and Changing Configuration

Whenever configuration or software changes are made, a StorNext HA cluster may need to be downgraded to one server with HA disabled. This avoids the possibility of an HA reset being induced by the arbitrary starts and stops of FSMs and other components as changes are made.

After changes have been made on one server, mirrored configuration items must be copied to the peer server. Examples include the following files:

- /usr/cvfs/config/*.cfgx
- /usr/cvfs/config/FSMlist
- /usr/cvfs/config/fsnameservers
- /etc/fstab

Production Single-Server Operation

During extended outages of one server, it might not be productive to incur an HA reset since there is no standby FSM to fail over to. However, reconfiguring the remaining server to non-HA mode is not practical. The single mode solves this dilemma.

Single mode can be entered from default mode, and default mode can be entered from single mode without stopping StorNext. This makes it easy to decommission and replace a failed server. Here are the steps for doing this:

- 1 Power down the decommissioning server (if necessary)
- 2 On the working server, run the following two commands in order:

```
snhamgr peerdown
snhamgr mode=single
```
- 3 Replace the decommissioned server
- 4 Acquire licenses for the new server
- 5 Replace those license file entries on the working server
- 6 Install StorNext on the new server, but do not configure it except for copying in the /usr/cvfs/config/fsnameservers file from the working server
- 7 On the working server, run the following two commands in order:

```
snhamgr mode=default
snhamgr peerup
```
- 8 Run the conversion script on the new server as follows:

```
/usr/adic/util/cnvt2ha.sh secondary <shared fs>
<peer IP addr>
```

Non-production Operation

There is a method for starting the SNFS file systems without starting the Storage Manager management components in the rare case that this is needed. The following two commands accomplish the same goal:

- `adic_control startonly snfs`
- `DSM_control startonly`

These commands create the `/usr/adic/install/start_snfs_only` touch file, which signals to the `snactivated.pl` script not to start the management components. The file exists until StorNext is stopped, and has its effect whenever the FSMs activate.

HA Resets

After a suspected HA Reset, the first place to look is the `/usr/cvfs/debug/smithlog` file, which contains one-line time-stamped descriptions of probable causes for the reset.

There are three methods for producing an HA Reset:

- 1 Expiration of an HA Monitor timer
- 2 Exit of the active HaShared FSM while the shared file system is mounted on the active MDC
- 3 Invocation of the `'snhamgr force smith'` command by a script or manually by an administrator. The `smithlog` file is written by the `fsmrpm` process, so there would not be an entry in the file when an `fsmrpm` exit results in an HA Reset.

HA Resets of the First Kind

The first method of an HA Reset is explained by the following description of the FSM monitoring algorithm (patent pending). The terms *usurp* and *usurpation* refer to the process of taking control of a file system, either with or without contention. It involves the branding of the arbitration block on the metadata disk to take control, and then the timed rebranding of the block to maintain control. The HA Monitor algorithm places an upper bound on the timing of the ARB branding

protocol to prevent two FSMs from simultaneously attempting to control the metadata, even for an instant.

- When an activating HaUnmanaged or HaShared FSM usurps the ARB, create a five-second timer that resets the computer if it expires
- Wait five seconds plus a small delta before completing usurpation
- Immediately after every ARB Brand update (.5 second period), reset the timer
- Delete the timer when the FSM exits

When there is a SAN, LUN, or FSM process failure that delays updates of the ARB, the HA Monitor timer can run out. When it is less than one second from expiring, a one-line message describing this is written to the /usr/cvfs/debug/smithlog file.

If SAN or LUN delays are suspected of occurring with regular frequency, the following test can be run. This will significantly impact performance.

- Increase the timer value (up to 999 seconds) by creating the /usr/cvfs/config/ha_smith_interval file on each MDC with only this line: 'ha_smith_interval=<integer>'. This will allow the delays to run their course without incurring a reset. The value must match on both MDCs.
- Turn on debugging traces with 'cvdbset :ha'
- Display debugging traces with 'cvdb -g -C -D 500'
- Look for the lines like this example 'HAMonCheck PID ##### FS "testfs" status delay = 1'
- When the value grows is more than 1, there are abnormal delays occurring. When a standby FSM is running and the LAN is working, the negotiated timer resets should limit the growth of this value to four. When the value reaches two times the ha_smith_interval (default of $5 * 2 = 10$), an HA Reset occurs.
- Turn off tracing with 'cvdbset - all'

HA Resets of the Second Kind

The second method of HA Reset can occur on shutdown of CVFS if there is an unkillable process or delayed process exit under the HaShared file system mount point. This will keep the file system from being unmounted. The smithlog entry indicates when this has happened, but does not identify the process.

HA Resets of the Third Kind

The third method of HA Reset is the most common. It occurs when the snactivated script for the HaShared FSM experiences an error during startup. The current implementation invokes the 'snhamgr force smith' command to allow the peer MDC an opportunity to start up StorNext if it can. A similar strategy was used in previous releases. In this release, the failure to start will cause the /usr/cvfs/install/.ha_idle_failed_startup touch file to be created, and this will prevent startup of CVFS on this MDC until the file is erased with the 'snhamgr clear' command.

Using HA Manager Modes

The snhamgr rules for mode pairings are easier to understand by following a BAAB strategy for transitioning into and out of config or single mode. In this strategy, B stands for the redundant node, and A stands for the node to be placed into config or single mode. Enter the desired cluster state by transitioning B's mode first, then A's. Reverse this when exiting the cluster state by transitioning A's mode, then B's.

For the configuration-session example, place B in locked mode, then place A in config mode to start a configuration session. At the end of the session, place A in default mode, then place B in default mode.

For the single-server cluster example, shut down Linux and power off B, then designate it peerdown with the 'snhamgr peerdown' command on A, then place A in single mode. At the end of the session, place A in default mode, then designate B as up with the 'snhamgr peerup' command on A, then power on B.

HA Tracing and Log Files

The following log files contain HA related debugging information:

- /usr/cvfs/debug/ha_mgr.out
Log messages from the snhamgr_daemon
- /usr/cvfs/debug/hamgr_cmds_trace

Output from commands run by the `snhamgr_daemon`. Typically, several commands are run simultaneously. Their output becomes intertwined. This is normal.

- `/usr/cvfs/debug/snactivated.<fs name>.log`

Output from the `snactivated.pl` command per file system.

- `/usr/cvfs/debug/nssdbg.out`

Log messages from the `FSMPM daemon`. HA related messages include: the `HAMon` timer interval, anomalies in negotiations for the resetting of `HAMon` timers, setting of Primary status, activations of FSMs etc.

- `/usr/cvfs/data/<fs name>/log/cvlog`

Log messages from FSM processes. HA related messages include: the last write of the ARB after quiescing the metadata writes, waiting the HA interval after branding the ARB, launching of the `snactivated` script etc.

- `/usr/adic/HA/cnvt2ha.sh.log`

Output of the `cnvt2ha.sh` script.

- `/var/log/messages`

Mounts of `cvfs` file systems.

- `/usr/cvfs/debug/smithlog`

When an HA Reset is imminent, a descriptive line is added to the end of this file and the file is sync'd to ensure that the information is available for debugging the root cause of the reset. For example, when there is less than one second remaining on the HA Monitor timer, a notice is written in this file. It is likely that all other log files will lose some of the newest information at the time of the reset because it is in buffers that have not been written to disk.

Single (Singleton) Mode

Single mode (also known as Singleton mode) allows for extended operation without the risk of incurring an HA Reset. In this state HA is disabled, but with the possibility of reduced availability because the

redundant server is missing. Use of the “snhamgr force smith” command produces an error message, and the server continues to run. This and other instances where an HA reset would have occurred under Default mode are still logged in the /usr/cvfs/debug/smithlog diagnostic file.

In Single mode the Secondary must be either “Offline” (peerdown) or “Locked”. When in peerdown mode, the Secondary is truly incommunicado. When locked, Web services are still running on the Secondary.

There is no way in the StorNext GUI to go directly from Singleton/Locked to Default/Default, but it is possible to “Enter Config Mode” and then “Exit Config Mode” to get to Default/Default.

When in Singleton/Peerdown, the “Enter Config Mode” and “Exit Config Mode” sequence transitions the cluster as follows: Single/Peerdown -> Config/Peerdown -> Single/Peerdown.

Between the initial conversions of the Primary and Secondary servers, the GUI sets the cluster to Single/Peerdown. Quantum recommends that conversions be done one right after the other; there is no benefit to remaining in this half-converted state for any length of time. If the Secondary must be replaced (or when it is uninstalled during an upgrade), the StorNext GUI leaves the cluster in Default/Default (Unknown) state.

When leaving Config or Single mode to return to the Default/Default state, it is a best practice to have the same server be the Primary before and after the transition. This allows any configuration changes to be transferred to the Secondary before it activates any FSMs.

FSM failover in HA Environments

When a failover of any file system occurs, the new FSM notices if any clients had a file exclusively opened for writes, and waits up to 35 seconds for those clients to reconnect. In the case of an HA Reset of the Primary MDC, that MDC is not going to reconnect, so the failover to FSMs on the Secondary MDC and the promotion of that MDC to Primary status can be delayed by 35 seconds.

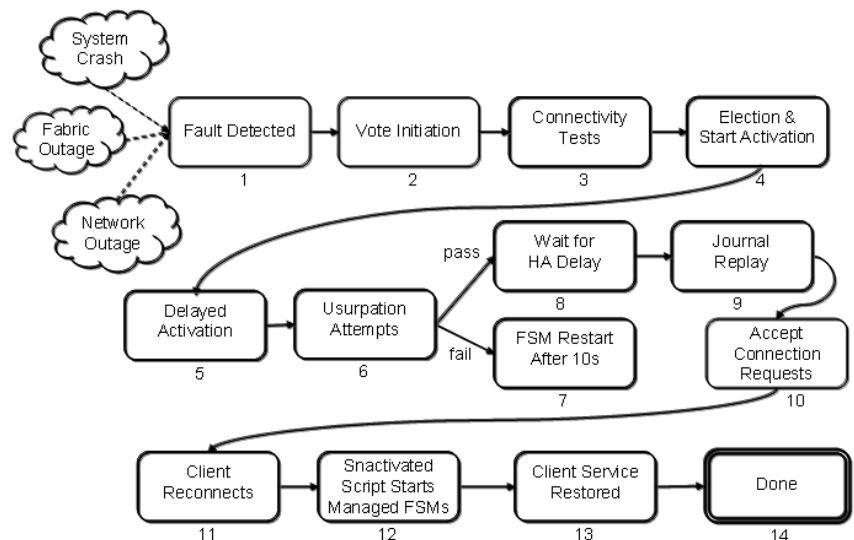
The StorNext system exclusively opens files on the HaShared file system, but assumes that only the Primary MDC does this and waives the delay for that one file system. Quantum advises against running user processes other than StorNext processes on HA MDCs for performance, reliability and availability reasons. In the event that processes running on the Primary MDC have files exclusively open for writes on other file systems, the availability of those file systems to all clients will be delayed by 35 seconds following an HA Reset event.

Failover Timing

The following illustration shows approximate timings of the FSM failover in an HA cluster. The numbers in the notes correspond to the numbers in the illustration.

In this description, both MDCs in an HA Cluster are fully started and the Secondary MDC is ready to assume the Primary role if needed. At time T0, an HA Reset of the Primary occurs.

Figure 132 FSM Failover in an HA Cluster



Not shown in this diagram are the state transitions of the peer MDC when it incurs an HA Reset. The HA Reset is not directly tied to the failover of control to a standby FSM, but rather the detection of a loss of services triggers failovers. The HA Reset may come before or after the loss of services, or not at all. It is only important to know

that by the end of state 8, no FSM on the peer MDC is controlling the arbitration block (ARB). The HA Reset mechanism guarantees that to be true.

The example failures shown here (System Crash, Fabric Outage, Network Outage) can result in a failover. Typically, the loss of heartbeat from the peer MDC's FSMPPM is the first indication that an HA Reset has occurred.

- 1 **Triggering Event:** The loss of heartbeat is detected and triggers an election at approximate time T3.5 seconds. Note that failover of a single unmanaged file system could also be forced with the `cvadmin` command without causing an HA Reset.
- 2 **Vote Initiation:** A quorum-vote election is started where the clients of the file system identify the best-connected MDC having a standby FSM for the file system.
- 3 **Connectivity Tests:** Each live client runs a connectivity test sequence to each server. Connections are tested in less than .5 seconds per server, when successful, and can be repeated up to four times (two seconds) when unsuccessful. At completion of the election, the time is approximately T5.5.
- 4 **Election and Start Activation:** The election is completed, and an activation message is sent to one server's standby FSM.
- 5 **Delayed Activation:** When a server has active FSMs, its FSMPPM process sends a request to the FSMPPM of its peer server to ask if the corresponding Standby FSMs are being activated. If not, the local FSMPPM can reset the HA timer of that file system's active FSM, which reduces the chance of an unnecessary HA Reset. When the peer FSMPPM gives permission, it is constrained from activating the standby FSM for two seconds. Step 5 is for that delay of up to two seconds. The delay completes at approximately T6.5.
- 6 **Usurpation Attempts:** To prevent false takeovers, the ARB is polled to determine whether another FSM is active and must be "usurped". Usurpation is averted if the activating FSM detects activity in the ARB and its vote count does not exceed the active FSM's client-connection count. A typical successful poll after an HA Reset lasts two seconds. When the previously active FSM exits gracefully, the usurpation takes one second.

The activating FSM then performs a sequence of I/Os to "brand" the arbitration block to signal takeover to the peer FSM. An active FSM is required to exit when it sees that its brand has been overwritten.

These operations take two seconds. The HAmom timer is started at this point if the HaFsType is HaShared or HaUnmanaged. This step completes at approximately T9.5.

- 7 FSM Restart:** After five failed attempts to usurp control, an activating FSM exits. The fsmppm restarts a standby FSM ten seconds later.
- 8 Wait for HA Delay:** When an active FSM is configured for HA Monitoring (HaShared or HaUnmanaged), and the ARB brand is not maintained for more than the HA Timer Interval (five seconds by default), the FSM's server computer is reset. After an activating FSM writes its brand, it waits one second longer than the HA Timer while monitoring its brand (HA Delay = six seconds by default), to be certain that the formerly active FSM has not resumed control of the ARB. The delay completes at approximately T13.5.
- 9 Journal Replay:** Any outstanding journal entries are replayed in order to achieve consistent metadata state. The time required for this step can vary due to several factors, but typically completes within seconds. A possible time for completion of this step is T18.5.
- 10 Accept Connection Requests:** The FSM begins to listen for client (re)connects. It waits up to 35 seconds for reconnections from any clients that have files open exclusively for writing, but this delay does not apply to the formerly active FSM's server computer. Approximate time at completion of this step is T20.5.
- 11 Client Reconnects:** The FSM begins servicing reconnects from the live clients. The clients perform a sequence of attribute state synchronization to ensure consistency with the server. Approximate time at completion of this step is T22.5.
- 12 Start Managed FSMs:** When the HaShared FSM reaches this step, it sets the Primary status for the server, which signals the FSMPPM to start the HaManaged FSMs. Those FSMs then proceed through steps 1 through 14, but without the initial 3.5 second delay in step 1, and without the delay in step 8, since they are not HA Monitored. Activation of the HaManaged file systems can complete in seconds, completing at approximately T27.5.
- 13 Client Service Restored:** The clients reinitiate any outstanding RPCs to the server and restore full service to the applications. This runs in parallel with starting HaManaged FSMs.
- 14 Done:** At this point, processes on clients can create, read, write etc. files in StorNext file systems unless Storage Manager Services are

needed. In that case there can be a delay of several minutes as those services are restarted before certain file system operations can be completed.

- 15** The approximate time of 27.5 seconds to complete a failover is variable and could take less or significantly more time.

It is important to note that an HA Reset is possible on the Secondary server if an HaUnmanaged FSM is active there and fails to maintain its brand on the ARB within the timing constraints of the HA system.

The following table presents common timing estimates for failover of all file systems following an HA Reset of the Primary server. Actual performance will vary according to: differences in configurations; file system activities in progress at the time of failover; CPU, SAN and LAN loads, latency and health; and the nature of the conditions that caused the failover. The optimal estimates are for a forced failover at the command line of a single unmanaged file system without an HA Reset.

	Failover Timing Estimates (secs)	
	Optimal	Common
State		
1	0	3.5
2	0	0
3	0.5	2
4	0	0
5	0	1
6	3	3
7	n/a	n/a
8	0	4
9	1	5
10	0.5	2
11	0.5	2
12	0	5
13	0	2
14	n/a	n/a
Total	5.5	27.5

Replacing an HA System

This section describes how to replace an HA server. Before beginning this procedure make sure you have obtained the proper licenses required for the new HA MDC.

Note: This procedure requires a certain level of technical expertise. Do not attempt performing this procedure unless you are confident you can complete the steps successfully.

If you are unsure about your ability to complete these steps, contact the Quantum Technical Assistance Center for help.

Pre-Conversion Steps

- 1 If both HA MDCs are currently up and running, make sure the system you want to replace is designated as the secondary MDC. This can be accomplished by running `service cvfs stop` on the designated machine.
- 2 Run a manual backup to tape from the StorNext GUI.
- 3 Make sure all store/retrieve requests have finished.
- 4 If you are using the Distributed Data Mover (DDM) feature, note the value of the `DISTRIBUTED_MOVING` parameter (either `All` or `Threshold`) in `/usr/adic/TSM/config/fs_sysparm` (or `fs_sysparm_override`).

Use a text editor to set the `DISTRIBUTED_MOVING` value to `None`.
Use the `adic_control restart TSM` command to put this change into effect.
- 5 Unmount all file systems from all clients, and then stop the SNFS processes on each client machine. (On the Linux platform, do this by running `service cvfs stop`).
- 6 Uninstall StorNext from the secondary server, but retain the log files. Do this by running the command `install.stornext -remove`.
- 7 Power down the uninstalled secondary server.

Conversion Steps

- 1 Set the primary node to “Config” mode and the peer node to “Peerdown” mode by running the following commands:

```
snhamgr peerdown  
snhamgr mode=config
```
- 2 Check the StorNext HA Manager (snhamgr) status by running the command `snhamgr status`. The status should look similar to this:

```
LocalMode=config  
LocalStatus=primary  
RemoteMode=peerdown  
RemoteStatus=unknown
```
- 3 Change the `/usr/cvfs/config/ha_peer` file on the primary MDC to the new MDC IP address.
- 4 If the `/usr/cvfs/config/fsnameserver` file includes the old MDC IP address, replace it with the new MDC IP address on the primary MDC and all the clients.
- 5 In the primary MDC’s `/usr/cvfs/config/license.dat` file, remove all the old MDC licenses by commenting out the lines you want removed. Keep only the primary MDC licenses.
- 6 Push those changes to the synchronization mirror directory by running this command: `/usr/adic/util/syncha.sh -primary`
- 7 Install StorNext 4.2.1 build on the NEW secondary server by running this command: `install.stornext`
- 8 Put the new licenses on the NEW secondary servers into `/usr/cvfs/config/license.dat`. The StorNext GUI can be run on the secondary to enter the licenses.

Note: You must restart the StorNext GUI after you create or edit the `license.dat` file.

- 9 In the StorNext GUI, go to the **Tools > High Availability > Convert** screen and convert the secondary MDC to HA.

Post-Conversion Steps

- 1 After the conversion is complete, check the snhamgr status on both MDCs. Run the cvadmin command to verify that all file systems are listed correctly.
- 2 Perform a system backup by running the snbackup command. This process may take substantial time depending on the number of managed files in the system.
- 3 Start and mount StorNext file systems on the clients, and then verify that all clients have full access
- 4 Conduct a failover to confirm that the secondary MDC has converted correctly. Confirm this by testing access to all file systems, moving files to/from tapes, and reviewing GUI configuration information.
- 5 If you conducted a failover to the secondary server, fail back to the original primary server.
- 6 Verify that all clients still have full access.
- 7 If you are using the DDM feature and if you use the secondary server as a DDM mover, make sure the file systems are mounted.
- 8 If you are using DDM, edit fs_sysparm or fs_sysparm_override to use your preferred DDM mode, (All or Threshold).

Use the command `adic_control restart TSM` to put this change into effect.



Appendix D

Web Services API

The Web Services Application Programming Interface feature (WS-API) provides an HTTP-based interface to a subset of the StorNext Storage Manager User Commands. WS-API provides basic control over StorNext Storage Manager systems to track media and drives, and to store/truncate/retrieve files from any computer capable of creating a Web Services connection, which includes Windows, Macintosh, and Linux-based systems, among others.

When first released in StorNext 4.0, WS-API returned the text-format output of the commands. The most recent WS-API release changes ten commands (listed below) to add structured Extensible Markup Language (XML) or JavaScript Object Notation (JSON) optional output formats.

The following subset of Storage Manager Commands supports the -F parameter to select XML, JSON or text output, and their associated WS-API interfaces (in parentheses):

- fscancel (doCancel)
- fsfileinfo (getFileLocation)
- fsmedinfo (getMediaReport)
- fsmedlist (getMediaInfo)
- fsqueue (getSMQueue)
- fsretrieve (doRetrieve)

- `fsrcmcopy` (n/a)
- `fsrcmdiskcopy` (doTruncate)
- `fsstate` (getDriveReport)
- `fsstore` (doStore)
- `vsmove` (doMediaMove)

The new `fsxsd` and `vsxsd` commands provides access to XML schemas describing the output of those commands.

Using the APIs

Accessing commands through WS-API can be done through a Web browser, by constructing an appropriate URL, or programmatically through a machine-generated library that supports the SOAP HTTP Client Protocol. Examples of these are provided in the notes listed below and in an example client application written in C# that is included in the StorNext installation package.

The available WS-API commands and the structure of their arguments and parameters are defined in the Web Service Definition Language file (WSDL) for StorNext, which is available from the StorNext Metadata Controller (MDC). The StorNext WSDL file can be displayed by entering the following URL in the address window of a web browser:

<http://<MDC>:81/axis2/services/stornext?wsdl>

Microsoft Visual Studio provides a menu selection for adding a "Service Reference" to the WSDL, which gleans the interfaces from the WSDL to aid programmers. Visual Studio also provides a `wsdl.exe` command that can generate source code in supported programming languages for compiling into a dynamic-link library (DLL) that supports the SOAP HTTP Client Protocol. Instructions for doing this are included in the example client code provided with StorNext. The sample WS-API client-application code, written in C# for Visual Studio on Windows, is included in StorNext distributions for MDCs. The following pathname shows a typical location for the file:

`/tmp/stornext/phdist/RedHat50AS_26x86_64/examples/SNAPITest.zip`

In order to perform any command on the remote server (except for the `getSNAPIVersion` call), a matching password must be specified with the call. The server verifies this password against the one stored in the `/usr/adic/.snapipassword` file on the server.

Make sure this file is the same on all metadata controllers (MDCs) you will be accessing. This is especially important if you are using virtual IP (vIP) addresses, because you can be communicating to different machines across different API calls or even retries of the same call.

The `.snapipassword` file is stored as clear text, and it should be set to be readable only by root. If this file does not exist, no calls other than `getSNAPIVersion` will succeed.

Using APIs With the High Availability MDC Feature

The StorNext High Availability MDC Feature (HA) uses a pair of redundant MDCs to maximize the availability of StorNext file systems. Under normal operation, one of the MDCs is the Primary server for all Storage Manager processing, and the Secondary server is ready to take over operations if the Primary server stops. This is called an HA Failover.

The Virtual IP Address (vIP) is a feature of StorNext HA which allows using a single IP address to automatically connect with whichever MDC is currently Primary. Using the vIP for WS-API calls simplifies development of client code. However, a client application must handle some transitional behaviors that occur during an HA Failover.

During an HA Failover, there is typically a delay of one to two minutes between the stop of the original Primary MDC and the transition of the Secondary MDC to full Primary MDC status. The former Primary MDC can stop abruptly without closing TCP/IP connections. For a period of time after that, attempts to make new connections to the vIP will fail until the new Primary brings up the vIP interface. Next is a period of time when the Linter database has not fully started. Finally, there is a period when TSM has not fully started. Each of these periods presents a different type of error to the client application.

When using WS-API with a StorNext High Availability (HA) cluster, the following notes apply:

- 1 Configure a Virtual IP Address (vIP). The vIP is a static IP address that is automatically configured as a virtual interface on the Primary MDC. The vIP allows a client application to use a single IP address to locate the WS-API server.

- 2 Set a timeout that allows a client application to recover if it was connected to a server that stopped abruptly. (Since the client software is responsible for setting a timeout, the actual method of specifying a timeout value on each Web service request varies. Check with your client documentation for specific instructions.)

Stores and retrieves to tape may take longer than the timeout, but they will continue to run to completion if they can. Store and retrieve requests can also be submitted multiple times without impacting in-process transfers. When the transfer completes, all the identical requests will complete and return status if there is still a connection to the requestor.

- 3 When a request fails because the vIP server is not available, Linter is not running, or TSM is not fully started, wait a few moments and retry the request.

The WS-API Server does not time-out until the target command completes with either success or failure. The Client application's Microsoft .NET SoapHttpClientProtocol class has a Timeout property with a default of 100 seconds. This is more than enough for status-type requests, but it is too short for tape-transfer-type requests, which are dependent on file size, system configuration and contention for resources. For example, a 100 GB file transferring at 100 MB/s will take about twenty minutes to complete. Setting the SoapHttpClientProtocol Timeout property to -1 is equivalent to infinity.

WS-API APIs

This section provides descriptions for the APIs included with WS-API.

The doCancel API

Given a requestID (which can be retrieved by running the getSMQueue API), running the doCancel API aborts an operation in progress. Running this API is equivalent to running the fscancel command.

The doMediaMove API

Use the doMediaMove API to move media from one archive to another. Running this API is equivalent to running the vsmove command.

The doRetrieve API

Use the doRetrieve API to retrieve the data in files that have been stored and then truncated. Running this API is equivalent to running the fsretrieve command.

The doStore API

Use the doStore API to store files as specified by its respective policy. Running this API is equivalent to running the fsstore command.

The doTruncate API

Use the doTruncate API to truncate files that have been stored, thus freeing their allocated space on disk for reuse. Running this API is equivalent to running the fstruncate command.

The getDriveReport API

Use the getDriveReport API to generate a report about the state of all storage subsystem drive components. Running this API is equivalent to running the fsstate command.

The getFileLocation API

Use the getFileLocation API to generate a report about files known to TSM. Running this API is equivalent to running the fsfileinfo command.

The getMediaInfo API

Use the getMediaInfo API to produce a list of media in a data and/or storage area. Running this API is equivalent to running the fsmedlist command.

The getMediaReport API

Use the getMediaReport API to generate a report on media based on its current status. Running this API is equivalent to running the fsmedinfo and vsmedqry commands.

The getSMQueue API

Use the getSMQueue API to produce a list of currently executing Storage Manager commands (retrieves, stores, etc.). Running this API is equivalent to running the fsqueue command.

Examples

This section contains examples for WS-API URLs, sample XML output, sample JSON output, and sample text output.

Example: WS-API URLs

Following are example URLs which demonstrate how to use the output format parameter.

Note: Two fields in the URL must be modified for local use: MDC_address (or VIP) and Password.

- 1 Run the fsfileinfo command (getFileLocation API) for three files, and return output in structured XML format:

```
http://<MDC_address>:81/axis2/services/stornext/  
getFileLocation?password=<Password>&files=/stornext/dir/  
fileA&files=/stornext/dir/fileB&files=/stornext/dir/  
fileC&format=xml
```

- 2 Run the fsstore command (doStore API) for one file, with the "-f i" option, and return output in structured XML format:

```
http://<MDC_address>:81/axis2/services/stornext/  
doStore?password=<Password>&files=/stornext/dir/  
fileA&retention=i&format=xml
```

- 3 Run the fsrmdiskcopy command (doTruncate API) for one file, and return output in JSON format:

```
http://<MDC_address>:81/axis2/services/stornext/  
doTruncate?password=<Password>&files=/stornext/  
fileA&format=json
```


- 4 Run the `fsretrieve` command (doRetrieve API) for a directory, with the "-a" parameter, and return output in text format:

```
http://<MDC_address>:81/axis2/services/stornext/  
doRetrieve?password=<Password>&directory=/stornext/dir/  
&updateATime=1&format=text
```

Sample XML Output

```
[tester1@smo4 p1]# fsfileinfo -F xml /stornext/dir/fileA
<?xml version="1.0" encoding="UTF-8"?>
<fsfileinfo xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:noNamespaceSchemaLocation="fsfileinfo.xsd">
  <header>
    <commandName>fsfileinfo</commandName>
    <commandLine>fsfileinfo -F xml /stornext/dir/fileA</commandLine>
    <commandDescription>Generate a report about files known to the Tertiary
      Manager</commandDescription>
    <localDateISO>2011-09-07T11:00:20</localDateISO>
    <localDate>2011-09-07</localDate>
    <localTime>11:00:20</localTime>
    <localDayOfWeek>3</localDayOfWeek>
    <gmtDateISO>2011-09-07T16:00:20Z</gmtDateISO>
    <gmtDate>2011-09-07</gmtDate>
    <gmtTime>16:00:20</gmtTime>
    <gmtDayOfWeek>3</gmtDayOfWeek>
  </header>
  <fileInfos>
    <fileInfo>
      <fileName>/stornext/dir/fileA</fileName>
      <storedPathFileName>/stornext/dir/fileA</storedPathFileName>
      <storedPathSameAsFileName>false</storedPathSameAsFileName>
      <lastModificationDateString>03-aug-2011
        15:49:36</lastModificationDateString>
      <lastModificationDate>2011-08-03</lastModificationDate>
      <lastModificationDayOfWeek>3</lastModificationDayOfWeek>
      <lastModificationTime>15:49:36</lastModificationTime>
      <owner>root</owner>
      <location>DISK AND TAPE</location>
      <group>root</group>
      <existingCopies>1</existingCopies>
      <access>644</access>
      <targetCopies>1</targetCopies>
    </fileInfo>
  </fileInfos>
</fsfileinfo>
```

```
<targetStubSize>0</targetStubSize>
  <targetStubScale>1024</targetStubScale>
  <existingStubSize>n/a</existingStubSize>
  <fileSize>1936636</fileSize>
  <store>MINTIME</store>
  <affinity>n/a</affinity>
  <reloc>MINTIME</reloc>
  <class>pool</class>
  <trunc>MINTIME</trunc>
  <cleanDBInfo>NO</cleanDBInfo>
  <medias>
    <media>
      <mediaId>sdisk</mediaId>
      <copy>1</copy>
    </media>
  </medias>
  <checksums>
    <checksum>
      <summary>N</summary>
    </checksum>
  </checksums>
</fileInfo>
</fileInfos>
<statuses>
  <status>
    <statusCode>FS0000</statusCode>
    <statusNumber>0</statusNumber>
    <dayOfMonth>7</dayOfMonth>
    <requestId>2125016624</requestId>
    <commandName>fsfileinfo</commandName>
    <commandStatus>completed</commandStatus>
    <statusText>Command Successful.</statusText>
  </status>
</statuses>
<footer>
  <returnCode>0</returnCode>
  <localDateISOEnd>2011-09-07T11:00:20</localDateISOEnd>
  <localDateEnd>2011-09-07</localDateEnd>
  <localTimeEnd>11:00:20</localTimeEnd>
  <localDayOfWeekEnd>3</localDayOfWeekEnd>
  <gmtDateISOEnd>2011-09-07T16:00:20Z</gmtDateISOEnd>
  <gmtDateEnd>2011-09-07</gmtDateEnd>
  <gmtTimeEnd>16:00:20</gmtTimeEnd>
  <gmtDayOfWeekEnd>3</gmtDayOfWeekEnd>
  <elapsedTimeInSeconds>0.0077</elapsedTimeInSeconds>
</footer>
</fsfileinfo>
```

Sample JSON Output

```
[tester1@smo4 p1]# fsfileinfo -F json /stornext/dir/fileA
{
  "header": {
    "commandName": "fsfileinfo",
    "commandLine": "fsfileinfo -F json /stornext/dir/fileA",
    "commandDescription": "Generate a report about files known to the Tertiary
Manager",
    "localDateISO": "2011-09-07T11:07:36",
    "localDate": "2011-09-07",
    "localTime": "11:07:36",
    "localDayOfWeek": 3,
    "gmtDateISO": "2011-09-07T16:07:36Z",
    "gmtDate": "2011-09-07",
    "gmtTime": "16:07:36",
    "gmtDayOfWeek": 3
  },
  "fileInfos": [
    {
      "fileName": "/stornext/dir/fileA",
      "storedPathFileName": "/stornext/dir/fileA",
      "storedPathSameAsFileName": false,
      "lastModificationDateString": "03-aug-2011 15:49:36",
      "lastModificationDate": "2011-08-03",
      "lastModificationDayOfWeek": 3,
      "lastModificationTime": "15:49:36",
      "owner": "root",
      "location": "DISK AND TAPE",
      "group": "root",
      "existingCopies": 1,
      "access": 644,
      "targetCopies": 1,
      "targetStubSize": 0,
      "targetStubScale": 1024,
      "existingStubSize": "n/a",
      "fileSize": 1936636,
      "store": "MINTIME",
      "affinity": "n/a",
      "reloc": "MINTIME",
      "class": "pool",
      "trunc": "MINTIME",
      "cleanDBInfo": "NO",
      "medias": [
        { "mediaId": "sdisk", "copy": 1 }
      ],
      "checksums": [
        { "summary": "N" }
      ]
    }
  ]
}
```

```
],  
"statuses": [  
  {  
    "statusCode": "FS0000",  
    "statusNumber": 0,  
    "dayOfMonth": 7,  
    "requestId": 2125016625,  
    "commandName": "fsfileinfo",  
    "commandStatus": "completed",  
    "statusText": "Command Successful."  
  }  
],  
"footer": {  
  "returnCode": 0,  
  "localDateISOEnd": "2011-09-07T11:07:36",  
  "localDateEnd": "2011-09-07",  
  "localTimeEnd": "11:07:36",  
  "localDayOfWeekEnd": 3,  
  "gmtDateISOEnd": "2011-09-07T16:07:36Z",  
  "gmtDateEnd": "2011-09-07",  
  "gmtTimeEnd": "16:07:36",  
  "gmtDayOfWeekEnd": 3,  
  "elapsedTimeInSeconds": "0.0011"  
}  
}
```

Sample Text Output

```
[tester1@smo4 pl]# fsfileinfo -F text /stornext/dir/fileA
-----
File Information Report                               Wed Sep  7 11:10:35 2011
Filename:      /stornext/dir/fileA
Stored path:   /stornext/dir/fileA
-----

Last Modification: 03-aug-2011 15:49:36
Owner:            root                      Location:        DISK AND TAPE
Group:            root                      Existing Copies: 1
Access:           644                      Target Copies:   1
Target Stub:      0 (KB)                   Existing Stub:   n/a
File size:        1,936,636                 Store:           MINTIME
Affinity:         n/a                      Reloc:           MINTIME
Class:            pool                     Trunc:           MINTIME
Clean DB Info:    NO

Media:    sdisk(1)
Checksum:      N
FS0000 07 2125016626 fsfileinfo completed: Command Successful.
```




Appendix E

Storage Manager Truncation

Truncation is a StorNext feature that results in removing data blocks from disk. This process frees up space for additional files to be stored on the disk. This appendix contains an overview of how Storage Manager truncation works, and how to perform simple troubleshooting.

Truncation Overview

Truncation operations fall into two categories. The first category is the truncation that is performed as part of the normal StorNext processing. The second category is the “space management” truncation policies that are run only when the disk usage reaches certain key points.

For each file system defined on the MDC, there must be an entry in the `/usr/adic/TSM/config/filesystems` file.

There are five variables specified for each file system:

- 1 Low-water mark (default value is 75%)
- 2 High-water mark (default value is 85%)
- 3 Min-Use mark (default value is 75%)
- 4 Min-Use enable (default is true)

5 Truncation enable (default is true)

If truncation is not enabled on a file system, no files residing within that file system will ever be truncated.

If truncation is enabled on a file system, as files are stored to media they automatically become truncation candidates unless they are marked for immediate truncation. (See below for more details).

Thus, a file can be truncated during one these operations:

- 1 Immediately after store (only if policy class is configured)
- 2 Daily truncation
- 3 LoSpace truncation
- 4 Emergency truncation

Normal Truncation

These truncations are performed as part of the normal processing done by StorNext.

Immediate Truncation

This refers to truncation performed immediately after all copies of a file are stored to media. This is enabled on a policy class basis and can be enabled with this command:

```
fsmodclass <classname> -f i
```

The default is that a stored file becomes a truncation candidate. The file will be dealt with through normal truncation processing.

Immediate Truncation can also be enabled on a file-by-file basis by using the `fschfiat` command: `fschfiat -t i filename...`

Daily Truncation

The `fs_tierman` TSM daemon kicks off policy-based truncations each day after midnight.

In this case the call is: `fspolicy -t -c <class> -m <class-trunc-min-time> -z 1`

This processes each defined policy class within StorNext until all policy classes have been completed. After the `fspolicy` has been run against all policy classes, the daemon waits until the next day to run them again.

Each of these class-based truncation policies truncates eligible candidates until either the min-use mark, if enabled, or the low-water mark is reached or it runs out of truncation candidates. At that time it terminates execution.

An eligible truncation candidate is a file that has not been accessed during the truncation mintime interval.

Space Management

The two main space management cycles are described below. They will continue to run as long as one of the conditions for a particular cycle is met. Both the LOSPACE and "Emergency Space" conditions are handled by the `fs_space` TSM daemon.

LOSPACE Cycle

This cycle is activated when the disk usage of one or more file systems exceeds the percentage full defined by the high-water value. When reached, LOSPACE policies are executed in an attempt to reach the low-water mark on each affected file system.

By default, the policies are executed once in this order on all affected file systems:

- relocation policy
- truncation policy

The high-water and low-water values are displayed by the GUI File System Monitor and can be modified via the StorNext GUI. By default, these values are set to 85% and 75%, respectively.

In contrast to the Emergency policies described in the next section, what's different in the behavior of the LOSPACE policies is that `MINTRUNCTIME` and `MINRELOCTIME` are not ignored. Only files that can be truly relocated and truncated are affected.

First, the relocation policy is executed and it continues until there are no more relocation candidates available at which time it terminates.

The call made to perform the LOSPACE relocation is:

```
fspolicy -r -y <mountpoint> -a <affinity>
```

If the file system usage still exceeds the high-water mark, the truncation policy is executed and it truncates all candidates until no further candidates are available, at which time it terminates.

The call made to perform the LOSPACE truncation is:

```
fspolicy -t -y <mountpoint> -z <mintruncsize>
```

At this time the LOSPACE Space Cycle is complete for this file system. All other affected file systems are then processed in the same manner, first by running the relocation policy and then the truncation policy, if needed.

After all file systems have been processed, if any of them still exceed the high-water mark, a new LOSPACE cycle is started after a one-minute wait.

Thus, the low-water percentage may or may not be reached on any given file system. It depends solely on whether there are enough candidates available for relocation and/or truncation for that file system.

Emergency Cycle

Emergency policies are executed when either of the following conditions is met for a file system:

- 1 When a file system encounters the NOSPACE event, i.e. a file write has failed because of lack of space.
- 2 When the file system usage is greater than 99%.

By default, the policies are executed once in this order:

- 1 emergency truncation policy
- 2 emergency relocation policy
- 3 emergency store policy

The emergency truncation policy finds up to the 3000 largest files that can be truncated, ignoring MINTRUNC TIME, and performs the truncation. This is executed once each time the NOSPACE condition is reached.

The call made to perform this emergency truncation is:

```
fspolicy -t -y <mountpoint> -e
```

If the file system usage has not dropped below 100% after the emergency truncation, the emergency relocation policy is now run.

When the emergency relocation policy is run, it finds all files that can be relocated, ignoring MINRELOCTIME, and performs the relocation. As with the emergency truncation policy, this is executed once each time the EMERGENCY condition is reached.

The call made to perform the emergency truncation is:

```
fspolicy -r -y <mountpoint> -e
```

If the file system usage is still not below 100% after the emergency relocation, an emergency store policy on the file system is performed.

An emergency store means that the request is placed first in the queue, and that any files in the file system which can be stored will be stored regardless of policy. As with the other emergency policies, it is run only once.

The call made to perform the emergency store is:

```
fspolicy -s -y <mountpoint> -e
```

At this point the Emergency Space Cycle is complete.

Disabling Truncation

There are two ways to disable truncation: by using truncation feature locking, and by running commands that disable truncation.

Truncation Feature Locking

Truncation operations can be locked, i.e. prevented from running, by using the `fsschedlock` command.

The feature name for each truncation operation is:

- `mintime`: Daily truncation
- `loSPACE`: LoSpace Cycle

Disable Truncation Commands

Truncation can be disabled for files by using one of the following commands:

```
fschfiat -t e <filename>
```

```
fschdiat -t e <directory name>
```

Running `fschfiat` sets an attribute on the file which will prevent it from ever being truncated. Likewise, running `fschdiat` causes the attribute to be set on files when they are created in the directory, but will not have any effect on files that already exist in the directory.

Common Problems

This section describes some common truncation problems and how to address them.

Files Are Not Truncated as Expected

Even if the truncation mintime requirement is met, files may not be truncated. Files are truncated only to keep the file system below the low-water mark (by default 75% of the capacity). When the daily truncation policies run, the oldest files are truncated first in an effort to bring the file system usage below the low-water mark. Thus, files may remain on disk even though their truncation mintime requirement has been met if the disk space is not required.

You can use the StorNext GUI to adjust the low-water and high-water marks on each file system if more free disk space is desired. A temporary option to free up disk space is to run the following command:

```
fspolicy -t -c <policyclass> -o <goal>
```

The goal argument is a disk usage percentage. For example specifying `"-o 65"` will truncate files until the disk usage either reaches 65% or there are no more valid truncation candidates, i.e. the mintime requirement has been satisfied.

"Old" Files Not Truncating According to Policy Class

Truncation uses the file access time to determine if the truncation mintime requirement has been satisfied. If any application changes the access time of a file, the aging of the file restarts.

An example of this is where image files are listed using the thumbnail mode on an Apple Macintosh. This causes the OS to read the file to present this thumbnail, and the access time of the file gets updated to the current time. This in turn results in StorNext determining this file has not satisfied the truncation mintime requirement.

Emergency truncation ignores mintime so the files could still be truncating if the file system fills up. However, the best solution is to modify the way files are accessed so as to not update the access time. In the above example, this would mean not using the thumbnail view.

Small Files Not Truncating

If a policy class has been configured to use stub files, files with a size that is less than or equal to the stub size will not get truncated.

Miscellaneous Usage Notes

If you ingest lots of data per day relative to the size of the file system, (for example, more than 80%), the file system disk usage can stay at a high level of 90% or even higher. The main reason is that the truncation mintime is a minimum of one day, so that neither the LoSpace nor the daily truncation will truncate any files until the next day.

Also the emergency truncation only works to get the disk usage less than 100% and no lower. If this is the case and the disk usage is a concern, you should consider using immediate truncation on the policy classes within this file system.

Truncation performance is typically from 1500-2000 files per second. This depends on the metadata controller (MDC) hardware configuration and other activity on the MDC.

The rebuild policy checks for truncation candidates.

Scheduling Truncation Manually

Although Quantum recommends scheduling truncation and relocation through the StorNext GUI, you can schedule these policies manually.

Truncation and relocation policies are run automatically only once a day at midnight. (There are also policies that run as file system fill levels warrant, but the time when these occur is not scheduled. See the `filesystems(4)` man page for more information on the space-based policies.)

If there is a real need to have file data truncated at a specific time after last access, then two steps are necessary: setting the class time and scheduling the policy commands.

As an example, assume you want to truncate class data for `class1` after it has remained unaccessed after one hour. The first step would be to set the `mintrunc` time to 1h (or 60m). Next, you must schedule via cron a truncation policy to run every half hour.

(Note that even with policies running every half hour a file may wait up until the time between policies beyond the `trunc` time before truncation occurs. For example, if policies run on the half hour, a file is created at 01:00:01am, and it will not be truncated until 2:30am, because at 2:00am it is 1 second short of being an hour old.)

Setting Truncation

Following are the steps required to set truncation manually:

- 1 Create a new data class and set the truncation time as desired:

```
% fsaddclass -c 1h class1
```

or

```
% fsaddclass -c 60m class1
```
- 2 Create a directory and add a class relationship.

```
% mkdir /stornext/snfs1/relation1
```

```
# fsaddrelation -c class1 /stornext/snfs1/relation1
```
- 3 If a class and relation point had already existed you will just need to modify the truncation time if not correct

```
% fsmodclass -c 60m class1
```

- 4 If you want the files for class1 to be truncated after an hour (or close to it,) set up truncation policies to be run by cron as described in [Setting Up the cron Job](#) on page 383.

Setting Up the cron Job

When setting up the policy cron job it is probably easiest to set up a simple shell script to wrap the policy so that the processing environment is set up correctly. For example, set up a script under TSM: /usr/adic/TSM/util/truncPolicy

The contents of the script may look like:

```
#!/bin/sh
#
. /usr/adic/.profile
/usr/adic/TSM/exec/fspolicy -t -c class1 -o 0
```

Note: The last argument to the policy command '-o 0'. This tells the policy to keep truncating files until it runs out of candidates or the file system reaches 0% full. The filesystems(4) man page indicates the automatic nightly policy only truncates to the Min Use percentage and then quits even if more valid candidates are present. If the desire is to truncate all candidates, the '-o 0' is needed. Truncating all files for a class should be done carefully as there will be an expense to retrieving those files back if needed.

StorNext allows running only one truncation policy at a time. This is due to potential conflicts in candidate management between policies, and also for performance reasons. If you want to run multiple policies “at the same time,” put multiple policies in the truncate script and have them run sequentially.

Be sure to avoid placing an ampersand after the commands, as some will fail because they are locked out by the currently running policy. Also be sure when setting up the cron jobs not to have the scheduled scripts run too closely together. (See [Considerations When Scheduling Truncation and Relocation Policies](#) on page 385 for more considerations on scheduling truncation policies.)

Following is an example of a script with multiple scheduled policies:

```
#!/bin/sh
#
. /usr/adic/.profile
```

```
/usr/adic/TSM/exec/fspolicy -t -c class1 -o 0  
/usr/adic/TSM/exec/fspolicy -t -c class2 -o 0  
/usr/adic/TSM/exec/fspolicy -t -c class3 -o 0
```

The next step is to create the actual cron entry. Run `crontab -e` and set the entry to look something like this:

```
00,30 * * * * /usr/adic/TSM/util/truncPolicy
```

One last thing to note on scheduling 'extra' truncation or relocation policies: There is an expense to running these commands as they get and check their candidate lists, even if no files are actually truncated or relocated. This is especially true for sites where millions of files are resident on disk at one time.

What Happens If the Old Script is Still Running?

The `fspolicy` commands placed in the script have code internally that checks for truncation policies which are already running. If a second truncation policy runs while one is active, that fact will be recognized by the new policy and it will abort reporting a truncation policy is already in progress.

This is true regardless of whether the policy was started by hand, cron job, etc. So in summary, the new script will exit immediately without doing anything.

Note: For this reason, if you want to run multiple truncation policies 'at the same time,' Quantum recommends creating one script with the desired policies running sequentially and then adding that script to cron.

How Can I Tell How Long the Script Has Taken to Run?

Each TSM command logs output to the history log located at `/usr/adic/TSM/logs/history/hist_01` as it runs. This includes time stamps of when the command started and completed. It is possible by parsing this log to determine how much time the truncation policies are using.

Considerations When Scheduling Truncation and Relocation Policies

Because StorNext allows you to set the minimum truncation and relocation times for a class in hours or minutes (as well as days), it may be desirable for sites to schedule policies to run more frequently. By default when StorNext is first installed, a class truncation policy is run for each installed class starting at midnight. Relocation policies are also run at this time for classes with relocation enabled.

Note: These policies are not visible in the output from `fsschedule`, but are run automatically by the daemons that monitor fill levels in managed file systems. If desired, these policies can be locked out by using the `fsschedlock` command.

If there is a need to run these policies more frequently because you cannot wait until midnight to have files truncated (and you do not want to truncate the files immediately upon storing them) then you will need to schedule more of these policies at the desired times. That will currently have to be done by setting up cron jobs to invoke the `fspolicy` commands. (Note the section above on creating the policy scripts and scheduling the cron jobs.)

When scheduling the extra policies take these considerations into account. Mainly these considerations are the time required to run and the affect they have on other processing on the machine. Note that the test machine where these metrics were gathered had 8 CPUs (3 GHz) and 16 GB memory. Also note that the number of candidates referred to below are for files whose blocks are disk resident and do not include files that are already truncated.

- A truncation policy that scans 1 million candidates looking for files to truncate can take over 25 minutes even when no files are actually truncated. The process uses 15% of a CPU and .3 GB of memory. (These numbers will be the same for a relocation policy in that the candidate processing is done in the same way.)
- To scan and truncate these same 1 million files takes 45 minutes, uses 15% of a CPU and use .4 GB of memory. (For relocation the time to relocate is totally dependent on the file sizes as data is actually copied from one location to another.)
- While truncation was being running on a file system writes to that file system were observed to be up to 70% slower and reads were 5% slower.
- The storing of other managed files was observed to be up to 25% slower than normal while truncation was running.

Because of the impacts of candidate processing and truncation only one truncation process is allowed to run simultaneously. Note again the section on cron setup that mentions how to run these policies sequentially. It is easy to see as well that there is an expense to running these policies even if nothing is actually truncated. Care should be taken to only run the policies as often as is absolutely needed so unnecessary impact on other system activity can be avoided. If you have a system in which it is desired that truncation run every 5 minutes then there can't be any more than 111,000 files to truncate every 5 minutes or the system will not keep up. Note that this is assuming no other activity so the real number is probably going to be lower. If the desire is to truncate hourly then the files to truncate maxes out at around 1.33 million files. The recommendation here is to run these policies as infrequently as possible to meet your space requirements.



Appendix F Security

This appendix contains an in-depth overview security in general.

StorNext Security

There are two predominate security models in modern file systems: POSIX and Access Control Lists (ACLs). ACLs are actually “Lists” composed of Access Control Entries. These lists may be quite simple or quite complicated, depending on the user's requirements.

The POSIX model is the older and less flexible of the two, having just three main security groups: “User,” “Group” and “Other,” and three operation categories: “Read,” “Write,” and “Execute”. For a directory, “Execute” translates to the ability to change into that directory, while “Read” and “Write” control directory listings and file creation and deletion.

POSIX permissions are kept in the file's inode information and are read from the file system on Unix/Linux systems by calls to stat().

In order to know what kind of restriction to place on a file or directory, the OS first has to be able to track users and groups so it can later be matched up with its associated information in files and directories. On Windows, all users have two unique Security Identifiers (SIDs): one for their user identification and one for the groups they belong to. On Unix/

Linux and Mac OS X, every user has a User Identifier (UID) and that user is assigned to a group which has its own Group Identifier (GID).

This is the model that's built into StorNext and used by all StorNext clients on all operating systems unless it's overridden by the use of ACLs.

ACLs are currently supported only on Windows and Mac OS X. ACLs give fine-grained control over file access and do things POSIX permissions can't, such as allow for writes to a file while not allowing the file to be deleted. ACLs also offer the benefit of "inheritance", which allows a directory to specify the default set of ACLs for all files created inside of it.

ACLs are kept in the Extended Attributes for a file, which is an internal data structure attached to the file's first inode that contains additional information associated with the file. Only operating systems that know to ask for the extended information with the proper key will understand these ACLs. Currently, only Mac OS X and Windows know to use this information.

The StorNext File System implements both the Unix POSIX model, and on its Windows clients it implements the Windows Security Reference Model (SRM) to a level compatible with Microsoft's NTFS file system. Quantum attempts to marry the two models in a very simplistic way to allow a common user to bridge file objects between Unix and Windows.

StorNext does not implement any of the Unix ACLs models or the NFSv4 ACLs model.

ACLs on Windows

Each mapped drive, file, or folder on Windows contains a Windows Security Descriptor. This descriptor contains the owner, primary group, DACLs, and SACLs. Windows uses the Security Descriptor to control access to each object. Windows Administrators and Users typically use Windows Explorer to view, change, and create ACLs on files. This is done in Explorer by first selecting the file or folder, displaying its properties, and then clicking on the Security tab.

Each file/folder can have zero or more ACLs that specify how a user or group can access or not access the file or folder. The possible controls in each ACE are:

Folders	Files
Full control (all of the following)	Full control (all of the following)
Traverse Folder	Execute File
List Folder	Read Data
Read Attributes	Read Attributes
Read Extended Attributes	Read Extended Attributes
Create Files	Write Data
Create Folders	Append Data
Write Attributes	Write Attributes
Write Extended Attributes	Write Extended Attributes
Delete Subfolders and Files	
Delete	Delete
Read Permissions	Read Permissions
Change Permissions	Change Permissions
Take Ownership	Take Ownership

Each Item can be selected as: Allow, Deny, or not selected. If Full Control is selected as Allow or Deny, all the other attributes are set to either Allow or Deny.

In addition, each ACE on Windows is indicated to apply as follows:

- Folder
 - This folder only
 - This folder, subfolders, and files
 - This folder and subfolders
 - This folder and files
 - Subfolder and files only
 - Subfolder only

- Files only
- File
 - This object only

An individual object can also be set to disallow or allow inheritable ACLs from a parent, parent's parent, etc.

A folder can be created and it can be marked such that all of its ACLs will pass to any children. This process is called *propagation*. Individual ACLs on a folder can be propagated as indicated in the above list. File and sub-folders of a folder can have all or some of the “inherited” ACLs removed.

The propagation/inheritance information is contained in the Windows Security Descriptor. Users and administrators on Windows platforms use this capability extensively.

ACEs are ordered in an ACL. Explicit ACEs come first. An explicit ACE is one that is not inherited. Explicit ACEs which deny come before explicit ACEs which allow. Inherited ACEs are ordered such that the closer the parent, the sooner they appear. Each level of inherited ACEs contain deny before allow.

All file and folder access is determined by matching a user and group to the DACL of the object being accessed. The SACL is not used to perform the access check. The ACEs in the DACL are compared in order with the accessing user and group for the requesting access mode. If a “deny ACE” matches, access is denied. If an “allow ACE” matches all requested access bits, access is allowed. It is possible to have a “deny ACE” inherited after an “allow ACE” which will not take effect. This can happen because explicit ACEs take precedence as do inherited ACEs from a closer parent. See the examples in the Microsoft document “[How Security Descriptors and Access Control Lists Work.](#)”

There is an “everyone ACL” that can be added to objects such that all users are granted certain privileges if they do not match any other ACE.

When a Windows user creates a file on SNFS the Security Descriptor (SD) is kept as an attribute of the file object. The SD contains a primary SID, a group SID and a list of discrete ACLs (also know as the DACL). The SNFS file object also contains the Unix UID, GID and permissions fields. By default SNFS inserts user identifier “nobody” into both UID and GID containers. Then it modifies the Unix mode (permissions) container based on the following two rules.

- 1 If the file object contains the Windows access control entry (ACE) for the everyone SID (which equals S-1-1-0, equivalent to "world" or "all others" on Unix), then it will apply specific permissions using the following guidelines. If the object is a container object (directory) and the FILE_LIST_DIRECTORY access bit is set, mode O+R (4) is set, else it is clear.
 - a If the object is a container object and the FILE_TRAVERSE access bit is set, mode O+X (1); otherwise it is clear.
 - b If the object is a container object and the DELETE bit is set, mode O+W (2) is set; otherwise it is clear.
 - c If the object is a file and the FILE_READ_DATA bit is set, mode O+R (4) is set; otherwise it is clear.
 - d If the object is a file and the FILE_WRITE_DATA bit is set, mode O+W (2) is set; otherwise it is clear.
 - e If the object is a file and the FILE_EXECUTE bit is set, mode O+X (1) is set; otherwise it is clear.
- 2 If there is no everyone ACE, the Unix permissions for the file object will be NONE (-----).

If it is an existing file, when a user changes the Security Descriptor on a file or directory, the change can affect Posix Permissions as well:

If the owner of the file or directory is not being changed, then SNFS checks for a missing DACL or Everyone ACE within the DACL.

If there is no DACL, set the new mode to allow the owner to read/write/execute.

If there is a DACL, scan the ACEs for the "Everyone" SID, either in the Allow or Deny state:

- 1 Check the ACE mask to see if READ_DATA/WRITE_DATA/EXECUTE_FILE is set, and adjust the Other mode of the Posix permissions accordingly.
- 2 The User and Group mode bits are left untouched.
- 3 The Unix*FileCreationOnWindows configuration options are ignored for the Everyone SID

If the owner is changing:

- 1 map the SID owner to unix User/Group ownership via active directory - store this for later application

- If the SID does not have a UID associated with it, map the UID to the value of the MDCs configuration option, `UnixNobodyUidOnWindows`.
 - If the SID does not have a GID associated with it, map the GID to the value of the MDCs configuration option, `UnixNobodyGidOnWindows`.
- 2 Convert the mode bits for the Group and User - apply the `Unix*CreationModeOnWindows` config option masks to these.
 - 3 Apply the Everyone bits per step 1.2 above - again note that the Everyone ACE conversion to Posix Permissions ignores the `Unix*CreationModeOnWindows` configuration options
 - 4 Check to see if the `DOSMODE_READONLY` flag is set, and mask out the User/Group/Owner write bits if it is.
 - 5 If the UID is different from what is currently stored, change it (it is possible to have multiple SIDs mapped to the same UID)
 - 6 If the GID is different from what is currently stored, change it (it is possible to have multiple SIDs mapped to the same GID)

Note: The Standard Posix Permissions Other bits get set via the Everyone ACE regardless of the `UnixFileCreationModeOnWindows` and `UnixDirectoryCreationModeOnWindows` settings.

ACLs on Mac OS X

With Mac OS X 10.3 (Tiger), ACLs were introduced. This ACL implementation is very close to the Windows ACLs implementation.

The `chmod(1)` and `ls(1)` commands have been modified to handle ACLs. There is also a library API for applications, `acl(3)` that allows programs to operate on ACLs.

For a detailed description of Mac OS X ACLs, see "Security Overview: Permissions" from Apples web sites and click on ACLs.

ACLs take precedence over regular UNIX permissions. If no ACE match is found for a user's requested access, UNIX permissions are checked. Therefore, a user may not match any ACE but still have access if UNIX permissions allow.

Each ACE on Mac OS X has the same 13 possible permission bits as a Windows ACE:

Directories	Files
Search Through	Execute File
List Contents	Read Data
Read Attributes	Read Attributes
Read Extended (named) Attributes	Read Extended (named) Attributes
Create Files	Write Data
Create Subdirectories	Append Data
Write Attributes	Write Attributes
Delete Subdirectories and Files	
Delete this Directory	Delete this Directory
Read Permissions (ACL)	Read Permissions (ACL)
Change Permissions (ACL)	Change Permissions (ACL)
Take Ownership	Take Ownership

Inheritance on Mac OS X is similar but does vary from Windows propagation and inheritance. Each ACE applied to a directory can be “propagated” by indicating one of 4 tags:

- 1 file_inherit:** Propagate this ACE to files created in this directory.
- 2 directory_inherit:** Propagate this ACE to subdirectories created in this directory.
- 3 limit_inherit:** After propagating this ACE to a new subdirectory, do not let its subdirectories inherit this ACE.
- 4 only_inherit:** Do not apply this ACE to this directory, just to files and/or directories created below it.

The “limit_inherit” exists in Windows as a check box when creating an ACE on a folder that propagates. The mapping of the 3 remaining

tags to the 7 Windows propagation pull down menu options are as follows:

Windows	Mac OS X
This folder only	(none)
This folder, subfolders, and files	directory_inherit, files_inherit
This folder and subfolders	directory_inherit
Subfolders and files only	files_inherit
Subfolders and files only	files_inherit, directory_inherit, only_inherit
Subfolders only	directory_inherit, only_inherit
Files only	files_inherit, only_inherit

On Mac OS X, propagation/inheritance is typically applied only when a file or directory is created. That is, when an object is created, its parent's list of ACEs is checked and any that apply are "inherited." When an ACE is added to a parent directory, it is not "automatically" propagated to any existing files or directories. Window's has a check box to cause some of this action when creating an ACE. On Mac OS X, the "chmod" command with the "+a1" option can be used to cause children to inherit an ACE. This can be done for large sub-trees with the chmod -R option.

Order of ACE entries is important because some ACEs might explicitly deny while others allow. Local ACEs are entries which are not inherited and by default are inserted before inherited ACEs. ACEs are checked in order for the requesting user/group and the requested access. The first ACE that denies or allows all the requested access stops permission determination. If there is a subsequent opposing deny or allow ACE, it will be ignored.

ACLs can be explicitly ordered with the `chmod` command which can lead to “non-canonical” ordering of ACLs. See Apple documentation for more details.

“Central Control”

StorNext supports cluster-wide central control to restrict the behavior of SNFS cluster nodes (fsm server, file system client and cvadmin client) from a central place. A central control file, `nss_cctl.xml`, is used to specify the desired controls on the cluster nodes. This file resides under `/usr/cvfs/config` on an nss coordinator server.

This control file is in xml format and has a hierarchical structure. The top level element is `snfsControl`. It contains the control element `securityControl` for certain file systems. If you have different controls for different file systems, each file system should have its own control definition. A special virtual file system `#SNFS_ALL#` is used as the default control for file systems not defined in this control file. It is also used to define the cvadmin related controls on clients.

Note: You cannot have a real file system named `#SNFS_ALL#`.

Each file system related control element (i.e., `securityControl`) has a list of `controlEntry` entries. Each `controlEntry` defines the client and the intended controls. A client can be of type `host` or `netgrp`. For the `host` type, `hostName` can be either an IP address or a host name. Both IP V4 and IP V6 are supported.

The `netgrp` entry specifies a group of consecutive IP addresses. `netgrp` has two sub-elements: `network` defines the IP address (either V4 or V6) of the network group, and `maskbits` defines the network mask bits.

Overlap is possible between the IP addresses in the `host` section and the `netgrp` section, and the `host` entries should be defined before `netgrp` entries. In this case, the `netgrp` control is considered to be a generic case, while the controls for individual hosts are considered to be a special case. A special case takes precedence.

(Support for cluster-wide central control debuted in StorNext 4.0.)

Controls

Currently seven controls are supported. Each control has this format:

```
<control value="true|false"/>
```

The “value” can be either “true” or “false”. The control is one of the following controls:

mountReadOnly

Controls whether the client should mount the given file system as read only. Value “true” means the file system is mounted as read only. Value “false” means the file system is mounted as read/write. If this control is not specified, the default is read/write.

mountDlanClient

Controls whether the client can mount the given file system via proxy client. Value “true” means the file system is allowed to mount via proxy client. Value “false” means the file system is not allowed to mount via proxy client. The default is “mount via proxy client not allowed”.

takeOwnership

Controls whether users on a Windows client are allowed to take ownership of files or directories of the file system. Value “true” means Windows clients are allowed to take ownership of files or directories. Value “false” means Windows clients are not allowed to take ownership of files or directories. The default is that “take ownership is not allowed”.

Note: This control only applies to the clients running on Windows platforms.

snfsAdmin

Controls whether cvadmin running on a host is allowed to have super admin privilege to run privileged commands such as starting or stopping a file system. Value “true” means the host is allowed to run privileged commands. Value “false” means the host is not allowed to run privileged commands. If this control is not specified, the default is that super admin privilege is not honored.

snfsAdminConnect

Controls whether cvadmin running on a client is allowed to connect to another fsm host via “-H” option. Value “true” means the client is allowed to connect to another fsm host. Value “false” means the client is not allowed to connect to another fsm host. The default is that “-H” is not allowed.

exec

Controls whether binary files on the file system are allowed to be executed. Value "true" means their execution is allowed. Value "false" means their execution is not allowed. The default is that their execution is allowed.

suid

Controls whether set-user-identifier bit is allowed to take effect. Value "true" means the set-user-identifier bit is honored. Value "false" means the set-user-identifier bit is not honored. The default is that suid bit is honored.

Note: If no match is found for a given client's IP address, then the client has no privilege to access a SNFS cluster. If a file system has been defined, but the client is not defined in that file system's control (securityControl), then the client has no access privilege to the specified file system.

Limitations

Currently only the Linux platform is supported to be a nss coordinator server capable of parsing this xml file. If you have a non-Linux machine as the fsm server, in order to enforce this cluster-wide central control, you must use a Linux machine as your nss coordinator with this central control file in place. The nss coordinator typically can be a very low-end machine as it is not stressed heavily.

Example

Following is an example of a nss_cctl.xml file. In the example this file defines control of file system "snfs1," and also the special virtual file system "#SNFS_ALL#".

```
<snfsControl xmlns="http://www.quantum.com/snfs/
cctl/v1.0">
  <securityControl filesystem="snfs1">
    <controlEntry>
      <client type="host">
        <hostname value="192.168.230.132"/>
      </client>
    <controls>
      <mountReadOnly value="true"/>
    </controls>
  </securityControl>
</snfsControl>
```

```
        <mountDlanClient value="true"/>
        <takeOwnership value="false"/>
        <exec value="true"/>
        <suid value="false"/>
    </controls>
</controlEntry>
<controlEntry>
    <client type="netgrp">
        <network value="192.168.1.0"/>
        <maskbits value="24"/>
    </client>
    <controls>
        <takeOwnership value="true"/>
        <mountReadOnly value="true"/>
    </controls>
</controlEntry>
</securityControl>
<securityControl fileSystem="#SNFS_ALL#">
    <controlEntry>
        <client type="host">
            <hostName value="linux_ludev"/>
        </client>
        <controls>
            <snfsAdmin value="true"/>
            <snfsAdminConnect value="true"/>
        </controls>
    </controlEntry>
</securityControl>
</snfsControl>
```

Cross-Platform Permissions

In a homogenous environment permissions aren't a problem because they are either all POSIX, all Windows ACLs, or all Mac OS X POSIX/ACLs. However, when moving to a heterogeneous environment with, say, Macs and Linux, or Windows and Macs, the interaction between POSIX and ACLs can become complicated.

Config (.cfg) File Options

The StorNext config file has the following options that relate directly or indirectly to security or permissions:

- GlobalSuperUser
- Quotas
- UnixDirectoryCreationModeOnWindows
- UnixFileCreationModeOnWindows
- UnixIdFabricationOnWindows
- UnixNobodyGidOnWindows
- UnixNobodyUidOnWindows
- WindowsSecurity

GlobalSuperUser defines whether or not the global super user (root) privileges on the file system. It allows the administrator to decide if any user with super-user privileges may use those privileges on the file system. When this variable is set to "Yes", any super-user has global access rights on the file system. This may be equated to the `maproot=0` directive in NFS. When the GlobalSuperUser variable is set to "No", a super-user may modify files only where he has access rights as a normal user. This value may be modified for existing file systems.

Quotas has an indirect relationship with security in that it requires a Windows Security Descriptor (SD) to track the owner of a file to correctly maintain their quota allotment. Currently quotas in StorNext File System-only systems work correctly in either all-Windows or all-non-Windows environments. This is because of the way quotas are tracked; when the meta-data server is deciding how an allocation should be charged, it uses either the SD, if one exists, or the UID/GID.

Files created on Windows with WindowsSecurity ON always have an SD. Files created on non-Windows never have an SD. If a file that was created and allocated on a non-Windows platform is simply viewed on Windows, it gets assigned an SD as described above. At that point the quota will be wrong. Subsequent allocations of that file will be charged to the SD and not the UID/GID.

To fix this problem, the UID/GID "space" and SD "space" must be consolidated into one "space".

Note: When you enable or disable quotas using the CLI `cvadmin` command, the change does not persist after rebooting. In order to permanently enable or disable quotas, you must modify the Quotas parameter of the file system config file.

`UnixDirectoryCreationModeOnWindows` controls which initial permissions directories have. Typically this is set to 755, but might be set to 700 to prevent access by anyone other than the owner on Unix systems, and on Windows require the use of ACLs to allow the directory to be accessed by anyone other than the owner.

`UnixFileCreationModeOnWindows` controls which initial permissions files have. Typically this is set to 644, but might be set to 600 to prevent access by anyone other than the owner on Unix systems, and on Windows require the use of ACLs to allow the file to be accessed by anyone other than the owner.

`UnixIdFabricationOnWindows` prevents (when set to “no,”) or allows (when set to “yes”) fabricating a UID/GID for a GUID returned from a Microsoft Active Directory Server. When set to “yes”, the client overrides any UID/GID for that user, and instead fabricates its own UID/GID. Typically this setting is only set to “yes” if you have a Mac OS MDC.

`UnixNobodyGidOnWindows/UnixNobodyUidOnWindows` instructs the client to use this ID on Windows if an ID can't be found using Microsoft Active Directory.

`WindowsSecurity` enables or disables using Windows ACLs on Windows clients. Once turned on (provide a Windows security descriptor is created), it is always on, even if the `.cfg` is changed to “off”. In a Unix/Windows environment, if there isn't a specific Windows- User-to-Unix-User mapping, files created on Windows will be owned by “nobody” on Unix clients.



Appendix G

Troubleshooting

This appendix contains some basic troubleshooting remedies for common error conditions that may occur. Please see if your particular issue is listed, and then try the recommended solution before calling the Quantum Technical Assistance Center.

Another good troubleshooting resource is the Quantum Knowledge Base, which contains articles about issues pertaining to StorNext. Access the Knowledge Base from [Quantum.com](https://www.quantum.com/knowledge-base).

This appendix contains the following troubleshooting topics:

- [Troubleshooting StorNext File System](#)
- [Troubleshooting StorNext Storage Manager](#)
- [Troubleshooting OS Issues](#)
- [Troubleshooting Replication](#)
- [Troubleshooting HA](#)
- [Troubleshooting StorNext Installation and Upgrade Issues](#)
- [Troubleshooting Other Issues](#)

Troubleshooting StorNext File System

This section contains troubleshooting suggestions for issues which pertain to StorNext File System.

Question: How do I rename a standalone (unmanaged) StorNext file system?

Answer: Use the following procedure to change the name of a StorNext file system:

Note: This procedure is only for StorNext file systems that do not have the Tertiary Storage Manager (TSM) component installed.

- 1 Unmount the file system from all the client systems using it.
- 2 Stop the file system in cvadmin.
- 3 Run cvfsck with the following parameters:


```
cvfsck -j file_system_name
cvfsck -n file_system_name
```

 where file_system_name is the actual name of your file system.
 Make sure that cvfsck says that the file system is clean.
- 4 Do one of the following:
 - * If cvfsck detects no file system errors, go to the next step.
 - * If cvfsck detects file system errors, run it in a "fix" mode


```
cvfsck file_system_name
```
- 5 Rename the file_system_name.cfgx file and edit the fsmlist file to reflect the new file system name.
 By default, these files reside in the /usr/cvfs/config directory on UNIX systems and in the C:\SNFS\config folder on Windows systems.
- 6 Update the .cfgx file with the new file system name.
- 7 Run cvfsck in interactive mode to remake icb by typing cvfsck without double quotation marks. You will be asked which file system to check.

Here is a command example:

```
[root@testbox]# cvfsck
StorNext File System File Systems on host testbox:
1) *snfs1
2) *snfs2
The asterisk (*) denotes the file system is active.
Choose a file system by number (1-2) or zero (0) to
exit
```

Run the command against the file system that has a new name.

- 8 During its run, cvfsck senses that there is an icb mismatch and asks if you want to fix the mismatch. Type yes.
- 9 Make adjustments to the /etc/vstab and /etc/fstab files, as well as in the Windows StorNext User Interface to reflect the new file system name on all the systems involved.
- 10 Start the file system and make it active (cvadmin).
- 11 Mount the file system.

Question: What can I do when a StorNext File System client fails to mount a StorNext file system? I receive the following error:

```
'install path'\debug\mount..out
mount.cvfs: Can't mount filesystem 'FILESYSTEMNAME'.
Check system log for details. Invalid argument
```

Answer: This condition occurs when the system cannot resolve the IP address or hostname defined in the fsnameservers file.

Use the following procedure to troubleshoot this problem.

- 1 Find the failure reported in the file install_path/debug/nssdbg.out.

```
ERR NSS: Establish Coordinator failed GetHostByName
of '[HOST01]'
No such file or directory)
INFO NSS: Primary Name Server is 'HOST01' (unknown
IP)
ERR NSS: Establish Coordinator failed GetHostByName
of '[HOST02]'
(No such file or directory)
INFO NSS: Secondary #1 Name Server is '[HOST02]'
(unknown IP)
```

- 2 If it is similar to the events reported above, please check the file `fsnameservers` for proper setup.

The file is located in the following directory, depending upon the product and operating system:

- * For Windows StorNext File System 2.x: `C:\SNFS\config`
- * For Windows StorNext File System 3.x: `C:\Program Files\StorNext\config`
- * For Linux or UNIX 3.x and higher: `/usr/cvfs/config`

For this example, the hostname definition line was set up incorrectly:

```
# Primary FS Name Sserver is rock # HOST01#
# Secondary FS Name Server
#1 is crag
# [HOST02]
```

The format for the `fsnameservers` file should contain the name of the IP address or hostname to use as either a primary or a secondary coordinator. Using the IP addresses is preferred to avoid problems associated with lookup system (DNS or NIS) failures. The format of an `fsnameservers` file line is either:

`IP_address or HOSTNAME`

- 3 Correct the `fsnameservers` file so that it looks like this:

```
10.65.160.42
10.65.160.78
```

- 4 If the same error reoccurs, contact Quantum Technical Support.

Question: I have trouble with StorNext clients connecting to the StorNext metadata controller. What can I do?

Answer: One of the common issues in working with StorNext clients is the inability to connect to the StorNext metadata controllers (MDCs). Usually you can show this problem either by running `cvadmin` on a UNIX-based client or by clicking `Device Mappings>StorNext File System Client Properties` for a Windows-based client, and not being able to see the file systems from the StorNext MDC(s). If file systems are not visible at this level, the client is not connected to the MDC(s).

As described in the StorNext documentation, the MDC(s) and all clients should be on a dedicated and isolated metadata network. The dedicated metadata network should be open to all ports for UDP and

TCP traffic. In addition, the metadata controller(s) and network switches should not have firewalling enabled for the dedicated metadata network.

If the client is still not able to connect to the MDCs through the dedicated metadata network, check for the following:

- Is the hostname or IP address of the correct MDC(s) listed in the `fsnameservers` file (found in `/user/cvfs/config` for UNIX-based clients and `c:\SNFS\config` for Windows-based clients)?
- If the hostname (rather than the IP address) is listed in `fsnameservers`, can the client resolve the hostname (using `nslookup` at the UNIX prompt or at the command prompt on a Windows-based client)?

If the client cannot resolve the hostname, do one of the following:

- Resolve either the DNS setup or hosts file setup
- Enter the IP address of the MDC(s) in the `fsnameservers` file instead of the hostname.
- Can the client ping the metadata controller?

If the client cannot ping the metadata controller, resolve the networking issue to make sure the client is on the same dedicated metadata network and can ping the MDC(s).

- If the client can ping the MDC(s), can the client either telnet, ftp, or ssh from the client to the MDC(s)?

If the client cannot run telnet, ftp or ssh from the client to the MDC(s), it is likely that there is some manner of firewalling set up between the client and the MDC(s). If possible, disable this firewalling.

- If firewalling is set up on the dedicated metadata network and it is not possible to disable it due to some internal policy (the metadata network should be a dedicated and isolated network), the client can specify a range of ports to be used for metadata traffic.

By creating an `fsports` file (located in `/user/cvfs/config` for UNIX-based clients and `c:\SNFS\config` for Windows-based clients), you can specify a range of ports, both UDP and TCP, that can be allowed to pass through the firewall between the client and the MDC(s).

If other clients are having problems connecting to the MDC(s), they must also use their own copy of the fsports file.

Sample fsports File

```
#
# File System Port Restriction File
#
# The fsports file provides a way to constrain the TCP
# and UDP ports used by the SNFS server processes.
# This is usually only necessary when the SNFS
# control network configuration must pass through
# a firewall. Use of the fsports file permits
# firewall 'pin-holing' for improved security.
# If no fsports file is used, then port assignment
# is operating system dependent.
#
# If an fsports file exists in the SNFS 'config'
# directory it
# restricts the TCP and UDP port bindings to the user
# specified
# window. The format of the fsports file consists of two
# lines.
# Comments starting with pound-sign (#) in column one
# are skipped.
#
# MinPort VALUE
# MaxPort VALUE
#
# where VALUE is a number. The MinPort to MaxPort values
# define
# a range of ports that the SNFS server processes can
# use.
#
#
# Example:
#
# Restrict SNFS server processes to port range 22,000 to
# 22,100:
#
# MinPort 22000
# MaxPort 22100
#
```

Question: How much data is reserved for StorNext disk labels, and what is the process for recovering damaged labels?

Answer: StorNext reserves the first 1 MB of the disk for the label.

- For VTOC disk labels, the critical area of the label is the first 1,536 bytes (three 512-byte sectors).

VTOC is the only label type used by StorNext Version 2.6 and earlier, and is the default type used for LUNs with less than 2GB sectors by StorNext Version 2.7.

- For EFI disk labels, the critical area of the label varies with the disk sector size:
 - For 512-byte sectors it is the first 18,432 bytes (36 sectors).
 - EFI is used by StorNext 2.7 for LUNs larger than 2GB sectors.

If a StorNext disk label is ever inadvertently overwritten or otherwise damaged, the only method of recovery is to run the `cvlabel` utility with the original parameters used when the disk was initially labeled. The `nssdbg.out` log file for the system often proves useful in determining what label each disk device on the system had before the problem occurred.

Contact Quantum technical support for assistance recovering damaged disk labels.

Question: `Umount` hangs or fails for StorNext File Systems even though the `fuser` shows nothing. What's going on?

Answer: If a process opens a UNIX domain socket in a StorNext File System and does not close it, `umount` hangs or fails even though `fuser` does not show anyone using the file system.

Use the "`lsof -U`" command to show the UNIX domain socket. The process can be killed with the socket open.

Question: How do I resolve invalid inode errors

Answer: You may receive the error File System FSS 'File System Name[0]': Invalid inode lookup: 0x2a5b9f markers 0x0/0x0 gen 0x0 nextiel 0x0

Deleting an old file system while an NFS client is still mounted leaves legacy data about inodes that no longer exist on that client. The client is out of sync with the file system and calls for inodes that no longer exist. This leaves StorNext users with the concern that they have lost files that

can't be recovered. Because of this issue, the MDC generates alarming messages about metadata corruption.

Checking the "epoch" field of the NFS request and the file system will show that these inodes are all zeros and thus invalid. Code can be changed in the NFS handles so they include a unique identifier such as the "epoch" (microsecond creation time) for the file system.

Question: What happens when a file is moved from one managed directory to another?

Answer: Here are three possible scenarios, which assume that the file data is no longer on disk and only exists on tape:

- Scenario 1: If the managed directories are on the same file system and have the same policy class, then tape is not accessed.
- Scenario 2: If the managed directories are on different file systems and have the same policy class, the data is retrieved from tape so it can be moved to the new file system, but it does not get stored again.
- Scenario 3: If the managed directories have different policy classes, then the data is retrieved, moved, and then gets stored to media associated with the new policy class.

You might receive the following error message if a StorNext file system client system continuously reports restarting the file system and fills up the nssdbg.out file (excerpted from logfile </usr/cvfs/debug/nssdbg.out>):

```
:
[0327 14:40:59] 0x40305960 NOTICE PortMapper: RESTART
FSS service 'stornext-fs1[0]' on host stornext-client.
[0327 14:40:59] 0x40305960 NOTICE PortMapper: Starting
FSS service 'stornext-fs1[0]' on stornext-client.
[0327 14:40:59] 0x40305960 (debug) Portmapper: FSS
'stornext-fs1' (pid 8666) exited with status 2 (unknown)
[0327 14:40:59] 0x40305960 (debug) FSS 'stornext-fs1'
LAUNCHED -> RELAUNCH, next event in 60s
[0327 14:41:59] 0x40305960 (debug) FSS 'stornext-fs1'
RELAUNCH -> LAUNCHED, next event in 60s
[0327 14:41:59] 0x40305960 NOTICE PortMapper: RESTART
FSS service 'stornext-fs1[0]' on host stornext-client.
[0327 14:41:59] 0x40305960 NOTICE PortMapper: Starting
FSS service 'stornext-fs1[0]' on stornext-client.
```



```
[0327 14:41:59] 0x40305960 (debug) Portmapper: FSS
'stornext-fs1' (pid 8667) exited with status 2 (unknown)
[0327 14:41:59] 0x40305960 (debug) FSS 'stornext-fs1'
LAUNCHED -> RELAUNCH, next event in 60s
[0327 14:42:59] 0x40305960 (debug) FSS 'stornext-fs1'
RELAUNCH -> LAUNCHED, next event in 60s
```

:

This error occurs because on the StorNext client system the file `/usr/cvfs/config/fsmlist` was set up and configured. However, the 'fsmlist' file belongs to the server components of StorNext and is set up on the MDC only.

Verify this by running the command `<ls-l/usr/cvfs/config/fsmlist>` on the client: `stornext-fs1`

On the StorNext client, only the client portion of the StorNext product suite is installed. Verify this by running the command `# cvversions`. The following output appears:

Server not installed.

File System Client:

Built in `/scmbld/phxbld/cvfs`

Created on Wed Jun 4 23:11:32 MDT 2008

Built for Linux 2.6.9-67.ELsmp x86_64 -- RHEL 4.6

Client Revision 3.1.2 Build 10

To resolve this issue, delete `/usr/cvfs/config/fsmlist` and then restart the StorNext services.

Before you restart the StorNext services, verify the size of the `/usr/cvfs/debug/nssdbg.out`.

If the output file is considerably large, delete or rename the file and then restart StorNext.

If the problem persists, contact Quantum Technical Support.

Troubleshooting StorNext Storage Manager

This section contains troubleshooting suggestions for issues which pertain to StorNext Storage Manager.

Question: What should I do if StorNext fails to write to a tape drive and varies it to 'off-line' after the drive has been replaced?

Answer: In this situation you might receive the following error message after the write fails and the replaced drive has been placed off-line:

```
fs_resource[6609]: E1004(4)<00000>:{3}: Drive 2
SN:1310012345 has invalid device path and is being
taken off-line.
```

The message indicates that the device configuration was not updated and the replacement tape drive has a different serial number. Compare the serial numbers of the configured tape drives against which tape drives are seen by the operating system.

In this example, the system has two tape drives configured. Run the `-fsconfig` command and check the command output. It should look similar to the following:

Component ID: V0,1

```
-----
-----
```

```
Device pathname: /dev/sg6
Compression: On
User Alias: scsi_archive1_dr1
Component Type: DRIVE
Device serial #: 1310999999
Drive Type: LTO
Drive ID: 1
```

Component ID: V0,2 <--- this is Drive 2

```
-----
-----
```

```
Device pathname: /dev/sg2
Compression: On
User Alias: scsi_archive1_dr2
Component Type: DRIVE
Device serial #: 1310012345 <--- the serial number no
longer matches with the new tape drive
Drive Type: LTO
Drive ID: 2 <--- this is Drive 2
```

Compare the results with the output of the 'fs_scsi -p' command.

```
ADICA0C012345_LLA | Scalar i500 | medium changer | /
dev/sg5
13109999999 | ULTRIUM-TD4 | sequential access | /
dev/sg6 <--- scsi_archive_dr1
13108888888 | ULTRIUM-TD4 | sequential access | /
dev/sg4 <--- device path and serial number is not
known.
```

A new tape drive has a new serial number. If a drive is replaced, the special device file (/dev/sg2) is no longer valid.

To resolve the issue, follow these steps:

- Delete the original removed tape drive using the StorNext GUI. (See the StorNext User Guide for more information about this procedure.)
- Add the replacement tape drive using the StorNext GUI. (Again, consult the StorNext User Guide for more information about this procedure.)

If the issue persists, contact Quantum Technical Support. Refer to the Worldwide Service and Support page for more information.

Troubleshooting OS Issues

This section contains troubleshooting suggestions for issues pertaining to the operating system on which StorNext runs.

Question: When I updated the OS, all connected LUNs were reformatted and data lost. Is there anything I can do to prevent this from happening?

Answer: If you are not careful when performing an operating system update or reload, all attached LUNs can be reformatted and data on those LUNs will be removed. If the updated system includes StorNext, this could cause StorNext to no longer function.

When performing an operating system update or reload, disconnect any fibre-attached media from the system and have only local operating system-required LUNs visible. This will make sure only the required LUNs are affected.

Question: I've discovered that StorNext cannot see all disks when running Red Hat Linux. What should I do?

Answer: StorNext File System cannot see all the drives you want to configure in Red Hat Linux. When Linux is installed, it defaults to only 40 disk devices when it boots up.

To address this limitation, modify the CONFIG_SD_EXTRA_DEVS setting in your Linux config file (or use xconfig under the SCSI Support tab). Then, rebuild the kernel and reboot the system.

If you require assistance rebuilding a Linux kernel, contact the Linux support resources for your installation.

Question: What does the 'heartbeat lost' message from a Solaris client mean?

Answer: On a Solaris client, you may see the following error:

```
fsmpm[3866]: [ID 702911 daemon.warning] NSS: Name Server  
'StorNext hostname' (xxx.xxx.xxx.xxx) heartbeat lost,  
unable to send message.
```

In StorNext, the metadata controller and clients use an Ethernet network to exchange file system metadata. The fsmpm is a portmapper daemon residing on each StorNext File System client and server computer. Its purpose is to register an RPC identifier to the system's portmap daemon. The fsmpm publishes a well-known port where the file system (fsm) daemons register their file system name and port access number. All clients then talk to their local fsmpm to discover access information for their associated service.

Because of the importance of maintaining this connection, a heartbeat is performed over the metadata network, so if this connection is lost, a message is sent indicating a network communication problem to the fsnameserver (xxx.xxx.xxx.xxx).

Portmapper messages are logged in the nssdbg.out log file located in /usr/cvfs/debug.

System administrators should monitor the log files to make sure that connectivity is maintained.

Question: Why does StorNext fail to write to an LTO-4 tape drive and varies media to suspect on my Red Hat 5 and SuSE 10 system?

Answer: StorNext Storage Manager fails to write to a tape drive and marks the medium as 'suspect'.

Note: This is applicable only to Red Hat RHEL 5 and SuSE SLES 10 operating systems and StorNext 3.1.x (not to 3.5.0).

When a medium is marked as 'suspect,' check if the below messages are reported in the TSM log files:

```
Received check condition with no error data. op=0A
Flush residue write to destination failed: errno 0
Unable flush all of residue buffer to destination.
Write error occurred - marking media suspect.
Medium XXXXXX was marked as suspect.
```

If you receive this error message, the default settings of the SCSI generic driver of RHEL 5 and SLES 10 must be adjusted. (For more information about the default settings, visit the following web site: <http://www.linux.org/>.)

Following is a description of the parameters in question:

`allow_dio`: 0 indicates direct I/O disable, 1 indicates enabled

`def_reserved_size`: This is the default buffer size reserved for each file descriptor. Values between 0 and 1048576 are supported.

`allow_dio`: Quantum recommends setting this parameter to 1.

`def_reserved_size`: Quantum recommends setting this parameter to 524288 (= 512kB)

To verify, run these commands:

```
cat /proc/scsi/sg/allow_dio
cat /sys/module/sg/parameters/allow_dio
```

If the above commands return a value of '0', this means direct I/O is disabled. Run these commands:

```
cat /proc/scsi/sg/def_reserved_size
cat /sys/module/sg/parameters/def_reserved_size
```

If the above commands return a value less than '524288', this means the buffer size is not appropriate for LTO-4 tape drives.

Verify if the TSM startup file `/usr/adic/TSM/bin/TSM_control` defines any of the above parameters.

Substitute the settings as seen below or add them to the startup script after the shell declaration (`#!/bin/sh`) and the initial comments.

```
if echo RedHat50AS_26x86 | egrep "RedHat5|SuSE10" > /
dev/null; then
    echo 1 > /proc/scsi/sg/allow_dio
    echo 524288 > /proc/scsi/sg/def_reserved_size
    echo 1 > /sys/module/sg/parameters/allow_dio
    echo 524288 > /sys/module/sg/parameters/
def_reserved_size
fi
```

If the issue persists after making the above changes, contact Quantum Technical Support from the Worldwide Service and Support page.

Troubleshooting Replication

This section contains troubleshooting suggestions for issues which pertain to replication.

Question: After completing the steps to set up replication, I received this message: "Replication disabled on target." What went wrong?

Answer: You will receive this message if you fail to turn on inbound replication. To do this, edit the replication policy named "target" and then click the **Inbound Replication** tab. At the Inbound Replication field, select **On** from the pulldown list of options.

Question: What should I do if something happens to my replication source, such as if the source directory or its contents becomes damaged?

Answer: If there is a failure on the source, the system administrator must reconfigure both the replication source and target hosts. Specifically, the administrator must turn the former replication target into a replication source, and then reconfigure the former source (once it is repaired) as a replication target.

Question: I upgraded to StorNext 4.x from a previous release. How do I replicate files that were previously truncated by Storage Manager in that previous release?

Answer: One solution would be to retrieve the files from the original managed source location, and then replicate and truncate the files.

Troubleshooting HA

This section contains troubleshooting suggestions for issues which pertain to StorNext HA (high availability) systems.

Question: How can I restart a file system without causing an HA failover?

Answer: To be clear, individual file-system failover must be distinguished from HA Reset of an entire MDC. When redundant FSMs are running on both MDCs in an HA Cluster, the active FSM can be on either MDC. In the case of managed file systems, the FSMs are started only on the Primary MDC, so these can be stopped and started at will without causing an HA Reset. Unmanaged file-system FSMs are started on both MDCs, so stopping an active unmanaged FSM will result in a single file system failover to the standby FSM on the peer MDC. An HA Reset occurs only when the failover is putting the file system in danger of data corruption from simultaneous write access to StorNext metadata and databases. This is typically the case for the HaShared file system, so take extra care with its FSM.

The recommended way for making configuration changes and restarting FSMs is to use the 'config' mode, which stops CVFS on one MDC and disables HA Reset on the other. CVFS will be restarted when returning to 'default' mode with both MDCs operating redundantly. Use the following steps to do this at the CLI:

```
snhamgr config
<make configuration changes>
snhamgr start
```

If you are only restarting FSMs without making configuration changes, the following steps will restart an FSM:

To restart an HaManaged FSM, use this cvadmin command:

```
fail <file system name>
```

To restart an HaUnmanaged FSM or the HaShared FSM:

```
snhamgr mode=locked # on the secondary
snhamgr mode=single # on the primary
cvadmin # on the primary
fail <file system name>
select # repeat until you observe the FSM has started and
activated
snhamgr start # on the primary
```

Question: What Conditions Trigger a Failover in StorNext (File System only)

Answer: There could be several reasons why a failover is triggered. See [HA Resets](#) on page 350 in the HA appendix.

Question: What conditions trigger the voting process for StorNext file system failover?

Answer: Either a StorNext File System client or a Node Status Service (NSS) coordinator (the systems listed in the `fsnameservers` file) can initiate a vote.

An SNFS client triggers a vote when its TCP connection to a File System Manager (FSM) is disconnected. In many failure scenarios this loss of TCP connectivity is immediate, so it is often the primary initiator of a vote.

On Windows systems, StorNext provides a configuration option called *Fast Failover* that triggers a vote as a result of a 3 second FSM heartbeat loss. Occasionally, this is necessary because TCP disconnects can be delayed. There is also an NSS heartbeat between members and coordinators every half second. The NSS coordinator triggers a vote if the NSS heartbeat is absent for an FSM server for three seconds. Because the client triggers usually occur first, the coordinator trigger is not commonly seen.

Question: Why does the Primary MDC keep running without the HaShared file system failing over and without an HA Reset when I pull its only Ethernet cable? The HA Cluster appears to be hung.

In this situation the lab configuration is as follows

MDC 1:
Hostname Shasta
10.35.1.110

MDC 2:
Hostname Tahoe
10.35.1.12

Two File Systems:
HaShared type: HAFS
HaManaged type: Reno3

There are no other client computers.

Shasta is the Primary MDC before the Ethernet cable is pulled.

At one point after the Ethernet was pulled, cvadmin on Tahoe showed:

```
Tahoe:/usr/cvfs/config # cvadmin
StorNext Administrator
Enter command(s)
For command help, enter "help" or "?".
```

```
List FSS
```

```
File System Services (* indicates service is in control of FS):
1>*HAFS[0]                located on tahoe:50139 (pid 13326)
```

```
snadmin> select FSM "HAFS"
```

```
Admin Tap Connection to FSM failed: [errno 104]: Connection reset
by peer
```

```
FSM may have too many connections active.
```

```
Cannot select FSS "HAFS"
```

```
snadmin> start reno3
```

```
Start FSS "reno3"
```

```
Cannot start FSS 'reno3' - failed (FSM cannot start on non-Primary
server)
```

```
snadmin> activate reno3
```

```
Activate FSM "reno3"
```

```
Could not find File System Manager for "reno3" on Tahoe.
```

```
Cannot activate FSS reno3
```

Answer: The reason the failover and HA Reset did not occur is because the HaShared FSM on Shasta continues to be active, and this was detected in the ARB block through the SAN by the FSM on Tahoe.

Here's why. When the LAN connection is lost on Shasta, its active HaShared FSM continues to have one client: the Shasta MDC itself. On Tahoe, an election is held when the LAN heartbeats from Shasta's HAFS FSM stop, and Tahoe's FSM gets one vote from the client on Tahoe. The Tahoe FSM is told to activate, but cannot usurp the ARB with a 1-to-1 tie. So, it tries five times, then exits, and a new FSM is started in its place. You can observe this by running the cvadmin command and watching the FSM's PID change every 20 seconds or so.

With the StorNext 3.5.1 and earlier STONITH HA implementation, other FSMs (not the HaShared FSM) would get tie-breaking votes from the

non-MDC clients and activate, and the first of these to activate would trigger the STONITH reset of the disconnected MDC. In StorNext 4.x HA allows HaUnmanaged FSMs to failover without resetting the MDC if possible, and HaManaged FSMs do not fail over because they are only started on the primary MDC.

Starting with StorNext 4.x, HA requires configuring at least one more client (other than the MDCs) of the HaShared file system to break the tie. This allows StorNext to determine which MDC has LAN connectivity, and to elect its HaShared FSM with enough votes to usurp control. When an HA Cluster is configured this way, the disconnected MDC (Shasta) will reset because of the usurpation of the HaShared ARB.

After the reboot, CVFS will restart and attempt to elect its HaShared FSM because it is not getting heartbeats from its peer. However, these activation attempts fail to cause a second reset because the HaShared FSM never has enough votes to have a successful usurpation. (You can watch it repeatedly fail to usurp if you can get on the console and run the `cvadmin` command).

But what about the HaManaged Reno3 FSM? HaManaged FSMs are not started until the HaShared FSM activates and puts the MDC in Primary status. You can observe these blocked HaManaged FSMs with the `cvadmin 'fsmlist'` command, which displays the local FSMPM's internal FSM and process table. A remote FSMPM's table can also be viewed with `'fsmlist on <MDC name or address>'`.

Finally, the message: 'Admin Tap Connection to FSM failed', is an intermittent response that occurred because the timing of the `cvadmin select` command was during the period after the FSM failed the fifth usurpation attempt and before the FSM was restarted (a ten-second delay). At other times, the response will show an activating FSM. Note that the `cvadmin`-displayed asterisk simply indicates that the FSM has been told to activate, not that it has been successful at usurpation and activation.

Question: Using the same configuration above (Shasta and Tahoe), an HA Reset occurs if I pull the fibre connection from Shasta when it is the Primary MDC, but it takes 30-40 seconds. Why does it take so long?

Answer: When the fibre connection is lost, Shasta's FSMs cannot maintain their brands on their ARB blocks, so the HA timers do not get restarted in the *read, write, restart-timer ARB branding* loop. After five

seconds the timers would expire and reset the MDC. However, there is a second method for resetting the timers that uses the LAN.

Every two seconds, the FSMPPM on an MDC with active HA monitored FSMs sends a request to its peer FSMPPM with the list of locally active FSMs. The peer gives permission to reset those FSMs' HA timers if it is not activating them, and promises not to activate them for two seconds. When the reply returns within one second of the request, the timers are reset by the FSMPPM. This allows the cluster to ride through brief periods when LUNs are too busy to maintain the ARB brand, but otherwise are operating correctly.

So why does the reset occur after 30-40 seconds? After this delay, the HBA returns errors to the FSM, and the FSM quits. When the HaShared FSM quits with the file system mounted locally, an HA Reset occurs to protect databases for the Storage Manager etc.

Question: How do I resolve a StorNext GUI login issue in my high availability environment?

Answer: When CVFS is running on only one MDC of an HA Cluster, attempting to connect a browser to the down MDC's GUI produces a single page with a URL to the running MDC. Simply click the URL and login again.

When CVFS is down on both MDCs, the GUIs on both MDCs present a set of four troubleshooting pages. You can start CVFS from the CLI by running the following command: `service cvfs start`

Or, you can use the StorNext GUI's Manage HA page and click the **Enter Config Mode** button, and then click the **Exit Config Mode** button. When the second step has completed, the HA Cluster will be running in Default-Default mode with MDC redundancy.

Troubleshooting StorNext Installation and Upgrade Issues

This section contains troubleshooting suggestions for issues which pertain to installing or upgrading StorNext.

Question: Does a StorNext installation only support a single ACSLS server?

Answer: If there are multiple libraries that will be accessed by a StorNext instance, they are not controlled by the same ACSLS server. StorNext supports multiple ACSLS servers, but only one library on each server.

Troubleshooting Other Issues

This section contains troubleshooting suggestions for general StorNext issues and other issues which do not fall into another category.

Question: How can I find the Product Serial Number?

Answer: The serial number for StorNext Storage Manager is physically located on the front side of the media kit box. In addition, the administrator initially responsible for the software may have received the serial number through email when either he or she requested license information.

Both StorNext Storage Manager and StorNext File System have serial numbers in the format S/N SN000123 (for example, SN02728).

Note: The serial number is not available through the StorNext application.

Question: A system panic caused connectivity loss to metadata. Is there anything I can do to prevent this from happening?

Answer: Quantum testing has determined that there is an extremely small chance of the metadata controller causing a system failure if metadata or journal activity is interrupted by the loss of connectivity to the metadata LUN (one occurrence during a week of testing by unplugging disk devices from the metadata controller every five minutes).

If a metadata write exceeds the space allocated (due to loss of disk), StorNext may stop all kernel activity on the metadata controller to avoid potential data corruption. This is a timing issue when metadata IO activity is attempted during a tear-down of internal data structures as a result of losing disk space.

You can avoid downtime from this system failure by configuring a redundant metadata controller with the High Availability (HA) feature.

You should recover from this particular failure by bringing the metadata controller back online and running the `cvfsck` command to repair the metadata before allowing clients to remount the file system.

Question: My backups failed with database errors. What caused the failure?

Answer: Check the error logs for possible causes of backup failure. Errors in logs might look similar to this:

*

```
Error while executing select statement; code: 1069,
message: HY000 ERR# 1069 : [Relex][Linter ODBC Driver]
SQLExecDirect: native error #1069, STMT: select
count(*).
```

```
*
```

```
2008-10-25-23:00:24: ERR: /usr/adic/DSM/bin/
snmetadump error at line 1490 of snmetadump.c: System
call failed (1): File exists (17): Failed to acquire
snmetadump lock.
```

```
*
```

```
ERR: snmetadump returned exit status 2.
```

To remedy the situation, remove all the lock files under “/usr/adic/TSM/internal/locks/”. .

Question: What should I do after a database (Linter) crash?

Answer: After a planned or unplanned sudden stoppage of the Linter database (due to a system crash, power outage, etc.), the database is left in an inconsistent state. This is due to an uncompleted recovery based on the Incomplete Transaction Log. Thus, tools such as testdb will return an errant or misleading error status.

The database recovery process is automatically initiated the next time the Linter Database starts up. Linter must replay the incomplete transactions, after which the database can be gracefully stopped and checked via testdb if desired.

Note: After replaying the Incomplete Transaction Log, the database SHOULD be in a consistent state, eliminating the need for testdb to be initiated.

Question: snbackup is failing with the log message 'Manual cleanup required'. How do I do manual cleanup? The full log message is similar to this:

```
File /usr/adic/database/metadumps/metadump.<FS Name>.gz exists.
This means that journals were in the process of being applied.
Manual cleanup required for file system: <FS Name>
```

Answer: The metadump.<FS Name>.gz file referenced in the log message is created at the start of backup processing to provide an opportunity to restart the backup processing if it does not complete. Failure to complete could happen because of a power outage, hardware

failure or unknown software bug. After successfully completing backup processing, the .gz file is removed.

Normal backup processing of the metadump file involves three steps:

- 1 'Roll' the restore journal to a new sequence number by closing the current `restore_journal-<FS Name>.<Seq Number>` file, renaming it to `restore_journal-<FS Name>.<Seq Number>.completed`, and creating a new `restore_journal-<FS Name>.<Seq Number>` file.
- 2 Apply the 'completed' restore journals to the metadump to bring it up to date with the time of the last rolling of the restore journal, and rename each `restore_journal-<FS Name>.<Seq Number>.completed` file to `restore_journal-<FS Name>.<Seq Number>.applied` as it finishes being applied.
- 3 Save the metadump and applied restore journals to a Storage Manager relation-point directory that exists to save these files on tape or sdisk.

When a '.gz' file exists, it prevents running the three steps listed above for its related file system to allow the following corrective steps to be taken:

- 1 Change directory to `/usr/adic/database/metadumps`
- 2 Save copies of the metadump file, the .gz file and any `restore_journal` files for the file system
- 3 Take steps to resolve the root cause of the problem if known
- 4 Remove the `metadump.<FS Name>` file, and unzip the `metadump.<FS Name>.gz` file; this should remove the .gz file in the process of unpacking it
- 5 Rename any `restore_journal-<FS Name>.<Seq Number>.applied` files to `restore_journal-<FS Name>.<Seq Number>.completed`
- 6 Run `snbackup`

If the preceding steps fail again and leave another .gz file, send the files saved in step 1 to Quantum for analysis. The following steps will then restore normal backup processing:

- 1 Stop the file system
- 2 Run `cvfsck` on the file system to check for possible metadata inconsistency

- 3 Remove the metadump.<FS Name>.gz file
- 4 Run 'snmetadump -c -d' to create a new metadump
- 5 Start the file system

Note: The preceding steps may take a long time to complete, during which time the file system is not available to clients.



Appendix H

StorNext

Offline Notification

StorNext supports a Windows feature called **StorNext Offline Notification**. This feature should be installed only at sites which are using StorNext Storage Manager.

This appendix provides an overview of the feature and describes how to install the application, configure and uninstall the feature. The following topics are addressed:

- [StorNext Offline Notification Overview](#)
- [Installing the Notification Application](#)
- [Starting the Notification Application](#)
- [Configuring the Notification Application](#)
- [Uninstalling the Notification Application](#)

StorNext Offline Notification Overview

The StorNext Offline Notification feature is a Windows only application which helps users from accidentally retrieving offline StorNext files. Retrieving an offline (i.e., truncated) file often takes several minutes, and the user's application is "blocked" while the retrieval is in progress. This feature prevents unintentional access and the subsequent retrieval of offline files by asking the user if the retrieval is desired before continuing.

What is a StorNext Offline File?

A StorNext *offline file* is a file in the StorNext file system which has been moved from primary storage (such as a disk) to secondary storage by the Storage Manager. Usually this involves a storage device such as tape, which is slower than disk. Files are moved to secondary storage according to administrative policies (see Appendix E: Storage Manager Truncation).

Once this has occurred, the StorNext Storage Manager truncates the file and adds the Windows offline attribute to the file. You can identify an offline file in Windows Explorer by the square black box with a clock icon or X icon. Although the offline file is visible and you can view its properties, the data in the file is not physically present on primary storage (although stub files have some data present). The offline attribute should not be confused with Microsoft Windows® CIFS-related offline files.

Why Block Access to Offline Files?

Moving files from secondary storage to primary storage is often a lengthy process. The application causing the retrieval will be blocked until the file is restored. The file is usually on a tape, which contributes to the length of the delay. Users may wish to avoid these delays for files which are accidentally accessed. Generating a warning not only helps users understand they are accessing an offline StorNext file, but also why accessing the data in the file is taking so long.

Offline Notification Configuration Options

The StorNext Offline Notification application can be configured to work in one of the following three modes (see configuration instructions):

- 1 Notify the user via a pop-up dialog
- 2 Deny access to all offline files
- 3 Allow access to all offline files

The first mode, notify the user, should be used when the user is allowed to decide whether the file should be retrieved.

The second mode should be used if the workstation user never wants to retrieve offline files. In this case the user's application will fail if it tries to access data in an offline file.

The last mode should be used if unrestricted access to offline files is desired.

How the Notification Feature Operates

If your StorNext Offline Notification feature is set to configuration mode #1, you will see a pop-up any time any application on the Windows system tries to read or write an offline StorNext file. Your application will be blocked until either you respond to this pop-up or the pop-up times out after two minutes. If the pop-up times out, the I/O request will be allowed. The pop-up can occur only with files on a StorNext file system. Offline files on other file systems will not be blocked.

Note: Metadata file operations such as looking at a file's properties or renaming a file will not be blocked.

When a user responds to a pop-up, the response is stored in an internal cache. The cache is 1024 entries long. This cache is used for both "allow" and "deny" entries. The cache is checked before generating a pop-up. If a matching entry is found, the previous answer is used. If the cache is full and a new entry is needed, the oldest entry is removed. Files marked as "deny" will eventually time-out. "Allow" entries do not time-out.

If configuration mode #2 is selected, no pop-up will appear; access to the file will automatically be denied. The application immediately receives the error "Access Denied."

If configuration mode #3 is selected, no pop-up will appear; access to the file is automatically allowed. However, the application's I/O request will be blocked by StorNext until the file is retrieved.

For configuration mode #1: If multiple applications are trying to access the same file at the same time, all applications are blocked, but only one pop-up will appear. If the user wants to allow access to the file, he should select the "Yes" button. All blocked applications will continue to be blocked until the file is retrieved to primary storage. If another application tries to access the file before it is retrieved, it will not cause a new pop-up (unless the file has been removed from the Offline Notification's cache.)

If, on the other hand, the user wants to deny access to the file, the user should select the "No" button. All blocked applications will be denied access and all future requests will be denied access until the file is removed from the Offline Notification's cache.

Note: Responding "Yes" or "No" to any new pop-up will not affect previous responses.

In summary, for configuration mode #1 only, one pop-up dialog will appear for each accessed offline StorNext file. The file is no longer offline when the file is restored to primary storage.

In configuration mode #2, access to an offline file is always immediately denied.

In configuration mode #3, access to an offline file is always allowed. The user will experience a delay while the file is being retrieved, and she will not see a popup.

Installing the Notification Application

The Quantum Offline Notification feature can be installed on the following platforms:

- Windows XP (x86)
- Windows 2003 SP1 (x86)
- Windows 2003 SP1 (x64)
- Windows Vista (x86)
- Windows Vista (x64)
- Windows 2008 Server (x64)
- Windows 2008 R2 Server (x64)
- Windows 7 (x86)
- Windows 7 (x64)

Caution: Do not install this feature on a Windows CIFS server or multi-user machines. This feature should be installed only on single user machines.

Installing onto a Standalone Machine

Locate the appropriate installation package for the machine onto which you want to install. Use `SnfsOfflineNotifyInstall32.zip` for 32-bit machines and `SnfsOfflineNotifyInstall64.zip` for 64-bit machines.

Move the appropriate installation package to the destination machine and unzip it. You will find two files: `SnfsOfflineSetup.exe` and `SnfsOfflineSetup.msi`. Use the following steps to install:

- 1 Log onto the machine as a Local or Domain Administrator. In Windows Explorer start the install by double clicking on `SnfsOfflineSetup.exe`.

On Windows XP platforms and higher Windows versions, an alternate method to start the installation is to right-click on `SnfsOfflineSetup.exe` and select "Run as ..." as shown in [Figure 133](#).

If you select this option, you must log in with the credentials for the Administrator account as shown in [Figure 134](#).

Figure 133 Run as Administrator

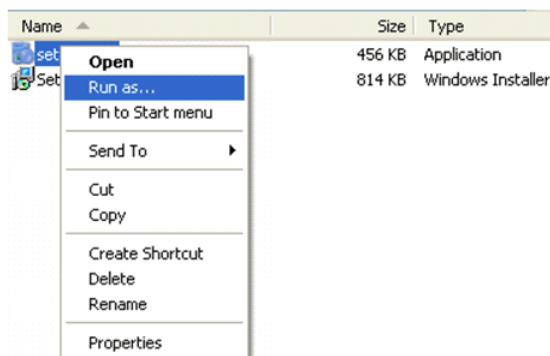
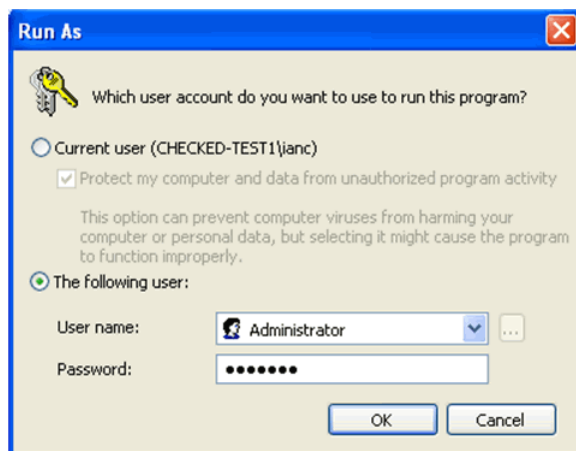



Figure 134 Logging in to the Administrator Account



- 2 The application installer requires .NET updates to function correctly. If you are installing onto a machine without any of the .NET updates, you will be prompted to update to .NET Framework 3.5 as shown in [Figure 135](#).

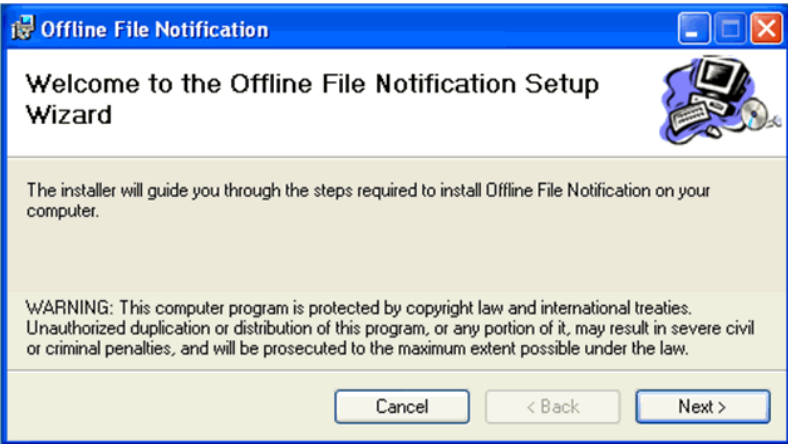
Figure 135 Installing the .NET Framework



- 3 Read the end user license agreement, and then click **Accept** to continue.

- 4 Wait for the .NET updates to be downloaded and installed. After this process is complete, the Offline Notification Setup Wizard launches.

Figure 136 Offline Notification Setup Wizard




- 5 Click **Next** to continue. The Quantum End User License Agreement window appears.

Figure 137 Quantum End User License Agreement



- 6 Read the end user license agreement. If you accept the terms of the agreement, select **I Agree** and click **Next** to continue. The **Select Installation Folder** window appears. (If you do not accept the terms of the agreement, click **Cancel** to stop the installation.)


Figure 138 Select Installation Folder



- 7 On the **Select Installation Folder** window, you can do the following:
 - Change the location where the installation resides by clicking **Browse** and navigating to the desired location
 - Specify whether to install Offline Notification for yourself only, or for everyone who uses the computer

- 8 Click **Next** to continue. The **Confirm Installation** window appears.

Figure 139 Confirm Installation

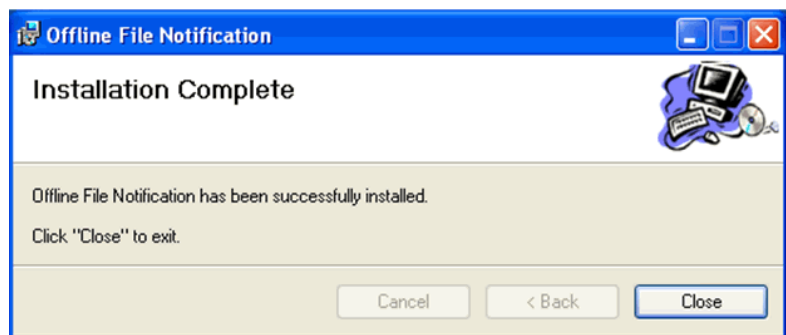


- 9 Before you proceed, make sure you have enough disk space to install the binaries. A minimum of 10MB is required.
- 10 Click **Next** to begin the installation.

Note: This is the last opportunity you will have to cancel the installation, so be certain you want to install before you click **Next**.

After the application is installed, the **Installation Complete** window appears.

Figure 140 Installation Complete



- 11 Click **Close** to exit the installation .

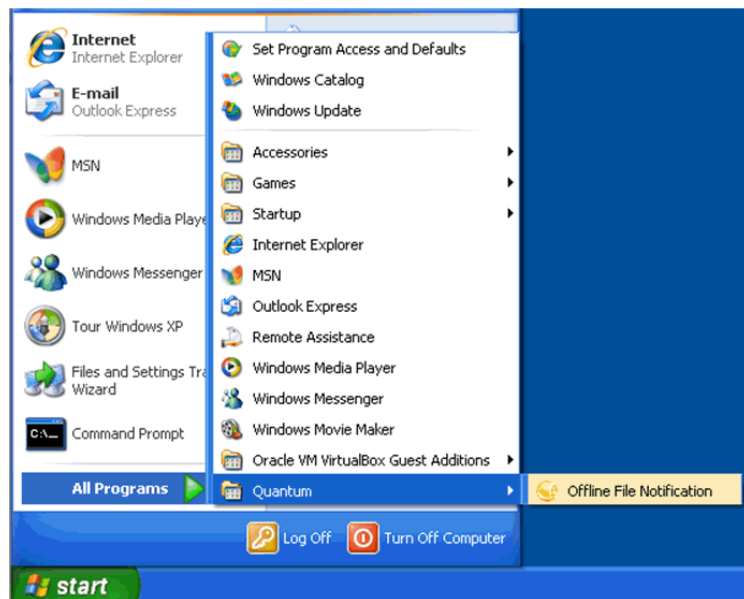
The next time you start your computer, the notification application automatically starts on login.

Starting the Notification Application

The StorNext Offline Notification application starts automatically after you log in, but if necessary you can start it manually.

From the Windows Start menu, choose **All Programs > Quantum > StorNext Offline Notification**. The application starts.

Figure 141 Manual Start



Configuring the Notification Application

When the application is already running, the gold StorNext icon appears in the Windows System Tray (generally found at the lower right corner of the screen).

When you hover the mouse over the gold StorNext icon and right click, three options appear:


- **Preferences:** View and enter application options, view a blocked file list, and view the application log
- **About:** View the StorNext Offline Notification application's current version number and copyright date
- **Exit:** Terminate the application

Setting Application Options

Follow these steps to set preferences for the StorNext Offline Notification application:

- 1 If you have not already done so, move the mouse over the gold StorNext icon in the Windows system tray and right-click. Select **Preferences**. (Alternatively, you can double-click the StorNext icon.)

Figure 142 Application
Options



2 View or adjust the following fields on the **Application Options** tab:

- **Cache Timeout:** The value in the **Cache Timeout** field determines how many minutes a file remains in the "Access Denied" list. The range is 1 to 499,999 minutes, and the default is 5 minutes. However, the cache can hold only 1024 entries total for both allowed and denied entries. If the total entries exceeds 1024, the oldest entry is deleted even if the timeout has not been reached.
- **Default Action:** This drop down list provides three options:
 - **Prompt User:** When this option is selected, users are always prompted with a dialog box whether to open an offline file, which means retrieving the file from offline storage such as tape, or from near-line disk storage. Users also have the option of preventing the file from being retrieved.
 - **Always Deny Access:** When this option is selected, access to offline files is always denied, preventing files from being retrieved from offline storage.
 - **Always Allow Access:** When this option is selected files are always allowed to be retrieved from offline storage without prompting the user.

Viewing Access Denied Files


The **Access Denied List** tab displays a list of files that have been prevented from being retrieved from offline storage. (See [Figure 143](#).) The files listed are automatically prevented from being retrieved as long as they remain in the Access Denied list.

These files are automatically removed from this list after they have not been accessed for a period of time (see the **Cache Timeout** field on the **Application Options** tab to determine the timeout period.) If you need to access any file in this list before the timeout period has expired, you must remove it from the list.

To remove files from the list, select the check box next to the desired files and then click **Remove Files**. Only the files selected will be removed from the access denied list.

Note: The **Remove Files** button is disabled when there are no files in the list.

Figure 143 Access Denied List




When you select the **Prompt User** option from the **Default Action** dropdown list, files are added to the access denied list after an attempt is made to open an offline file and you prevent the file from being retrieved.

Files are also added to the blocked files list whenever the **Default Action** of **Always Deny Access** is selected. This means that every offline file opened is added to the list and no notification is presented to the user.

Viewing the Application Log

The Application Log tab displays any events that have occurred in the system.

Figure 144 Application log



Exiting Application Preferences

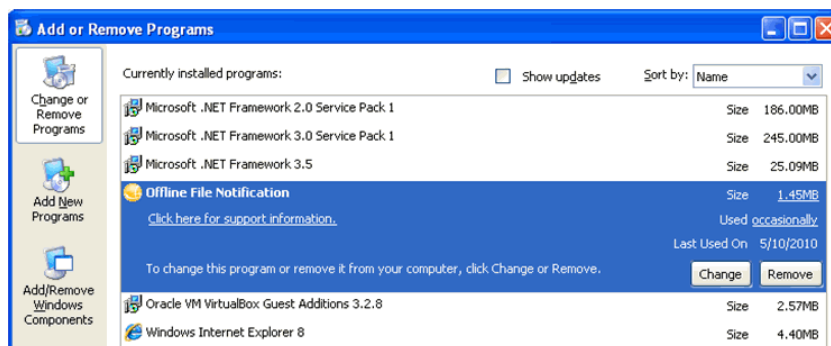
When you are finished viewing or setting preferences for the StorNext Offline Notification application, click **OK** to close the Preferences window.

Uninstalling the Notification Application

Follow these steps to uninstall the StorNext Offline Notification application:

- 1 Access the Control Panel by choosing **Control Panel** from the Windows Start menu.
- 2 When the Control Panel launches, open **Add or Remove Programs** (on Windows XP) or **Programs and Features** (on Windows Vista and later platforms).
- 3 When the **Add or Remove Programs** window appears click **Change or Remove Programs** if it isn't already selected.

Figure 145 Removing the Application



- 4 Locate and select the **StorNext Offline Notification** application. Programs are typically listed alphabetically.
- 5 Click **Remove** to uninstall the application.



Appendix I

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](#)

Note: This appendix shows only the actual RAS messages you might receive. A list of the events or error conditions which might trigger generating a RAS is not provided. For information about triggering events, refer to the SNFS_RAS man page.

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 146 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 <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 147 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	Contact the Quantum Technical Assistance Center .
Flag 2 (02h) - Write warning	
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 148 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	Contact the Quantum Technical Assistance Center .
Flag 14 (0Eh) - Unrecoverable snapped tape	
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	Clean the drive.
Flag 21 (15h) - Clean periodic	
Flag 22 (16h) - Expired cleaning media	<ol style="list-style-type: none"> 1. Remove the media. 2. Add new media.
Flag 23 (17h) - Invalid cleaning media	
Flag 24 (18h) - Retension requested	Contact the Quantum Technical Assistance Center .
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	
Flag 30 (1Eh) - Hardware A	Contact the Quantum Technical Assistance Center .
Flag 31 (1Fh) - Hardware B	
A hardware error has occurred that should be captured and returned for failure analysis.	

Figure 149 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.
Flag 33 (21h) - Eject media The drive has experienced an issue that can be resolved by unloading and reloading media.	<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.
Flag 35 (23h) - Drive Humidity Flag 36 (24h) - Drive Temperature Flag 37 (25h) - Drive Voltage Flag 38 (26h) - Predictive Failure Flag 39 (27h) - Diagnostics Required Flag 40 (28h) - Loader hardware A Flag 41 (29h) - Loader stray tape Flag 42 (2Ah) - Loader Hardware B Flag 43 (2Bh) - Loader door Flag 44 (2Ch) - Loader hardware C Flag 45 (2Dh) - Loader magazine Flag 46 (2Eh) - Loader predictive failure	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 .
Flag 52 (34h) - Tape system area Flag 53 (35h) - Tape system area read failure Flag 54 (36h) - No start of data Flag 55 (37h) - Loading failure	Contact the Quantum Technical Assistance Center .
The problem <u>IS</u> resolved:	Close the service ticket. Refer to Closing Service Tickets .

Figure 150 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">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. 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 151 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:	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 152 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:	<div>1. Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets.</div> <div>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</div> <div>OR</div>
	<div>3. If you are a properly-trained service professional, perform the procedures required for this type of tape library.</div>

Figure 153 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 154 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 155 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 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 156 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 <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

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 157 Configuration Not Supported RAS

Recommended Actions	
Configuration Not Supported	
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. <div><div>In the USA:</div><div>UK, France and Germany:</div><div>EMEA:</div><div>On the Web:</div></div><div><div>1+800-284-5101</div><div>00800 4 QUANTUM</div><div>+44 1256 848 766</div><div>http://www.quantum.com/support</div></div>

Figure 158 Label Validation Failure RAS

Recommended Actions	
Label Validation Failure	
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. <div><div>In the USA:</div><div>UK, France and Germany:</div><div>EMEA:</div><div>On the Web:</div></div><div><div>1+800-284-5101</div><div>00800 4 QUANTUM</div><div>+44 1256 848 766</div><div>http://www.quantum.com/support</div></div>

Figure 159 Connection
Rejected RAS

Recommended Actions	
Connection Rejected	
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. <div><div>In the USA:</div><div>UK, France and Germany:</div><div>EMEA:</div><div>On the Web:</div></div><div><div>1+800-284-5101</div><div>00800 4 QUANTUM</div><div>+44 1256 848 766</div><div>http://www.quantum.com/support</div></div>

Figure 160 File System Failover
RAS

Recommended Actions	
File System Failover	
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. <div><div>In the USA:</div><div>UK, France and Germany:</div><div>EMEA:</div><div>On the Web:</div></div><div><div>1+800-284-5101</div><div>00800 4 QUANTUM</div><div>+44 1256 848 766</div><div>http://www.quantum.com/support</div></div>

Figure 161 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 <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. <div><div>In the USA:</div><div>UK, France and Germany:</div><div>EMEA:</div><div>On the Web:</div></div><div><div>1+800-284-5101</div><div>00800 4 QUANTUM</div><div>+44 1256 848 766</div><div>http://www.quantum.com/support</div></div>

Figure 162 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 <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. <div><div>In the USA:</div><div>UK, France and Germany:</div><div>EMEA:</div><div>On the Web:</div></div><div><div>1+800-284-5101</div><div>00800 4 QUANTUM</div><div>+44 1256 848 766</div><div>http://www.quantum.com/support</div></div>

Figure 163 SNFS License
Required RAS

Recommended Actions	
SNFS License Required	
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:	<div>1. Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets.</div> <div>2. Contact the QuantumTechnical Assistance Center. <div>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</div></div>

Figure 164 SNFS License
Failure RAS

Recommended Actions	
SNFS License Failure	
IF	THEN
You receive a warning that your SNFS license will expire within 48 hours: OR Your SNFS license has expired:	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:	<div>1. Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets.</div> <div>2. Contact the QuantumTechnical Assistance Center. <div>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</div></div>

Figure 165 LUN Mapping
Changed RAS

Recommended Actions	
LUN Mapping Changed	
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. <div><div>In the USA:</div><div>UK, France and Germany:</div><div>EMEA:</div><div>On the Web:</div></div><div><div>1+800-284-5101</div><div>00800 4 QUANTUM</div><div>+44 1256 848 766</div><div>http://www.quantum.com/support</div></div>

Figure 166 Communication
Failure RAS

Recommended Actions	
Communication Failure	
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. <div><div>In the USA:</div><div>UK, France and Germany:</div><div>EMEA:</div><div>On the Web:</div></div><div><div>1+800-284-5101</div><div>00800 4 QUANTUM</div><div>+44 1256 848 766</div><div>http://www.quantum.com/support</div></div>

Figure 167 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 cvfsck 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. <div><div>In the USA:</div><div>UK, France and Germany:</div><div>EMEA:</div><div>On the Web:</div></div><div><div>1+800-284-5101</div><div>00800 4 QUANTUM</div><div>+44 1256 848 766</div><div>http://www.quantum.com/support</div></div>

Figure 168 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 snmetadump for the affected file system as soon as possible. Note: This condition could occur if cvfsck or cvupdatefs was recently run, or if a Restore Journal error occured 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. <div><div>In the USA:</div><div>UK, France and Germany:</div><div>EMEA:</div><div>On the Web:</div></div><div><div>1+800-284-5101</div><div>00800 4 QUANTUM</div><div>+44 1256 848 766</div><div>http://www.quantum.com/support</div></div>

Figure 169 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. <div><div>In the USA:</div><div>UK, France and Germany:</div><div>EMEA:</div><div>On the Web:</div></div><div><div>1+800-284-5101</div><div>00800 4 QUANTUM</div><div>+44 1256 848 766</div><div>http://www.quantum.com/support</div></div>

Figure 170 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. <div><div>In the USA:</div><div>UK, France and Germany:</div><div>EMEA:</div><div>On the Web:</div></div><div><div>1+800-284-5101</div><div>00800 4 QUANTUM</div><div>+44 1256 848 766</div><div>http://www.quantum.com/support</div></div>

Figure 171 File Processing Failure RAS

Recommended Actions	
File Processing Failure	
IF	THEN
A failure occurred while trying to process an internal file.	See the error details for more complete information about the failure. Possible reasons for the failure: <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 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.
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. <div><div>In the USA:</div><div>UK, France and Germany:</div><div>EMEA:</div><div>On the Web:</div></div><div><div>1+800-284-5101</div><div>00800 4 QUANTUM</div><div>+44 1256 848 766</div><div>http://www.quantum.com/support</div></div>

Figure 172 Missing LUNs RAS

Recommended Actions	
Missing LUNs	
IF	THEN
A client fails to mount because a LUN is missing:	Check the system logs to determine the root cause. Run the cvsadmin "disks" and "paths" commands, and then check for missing LUNs.
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. <div><div>In the USA:</div><div>UK, France and Germany:</div><div>EMEA:</div><div>On the Web:</div></div><div><div>1+800-284-5101</div><div>00800 4 QUANTUM</div><div>+44 1256 848 766</div><div>http://www.quantum.com/support</div></div>

Figure 173 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 <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. <div><div>In the USA:</div><div>UK, France and Germany:</div><div>EMEA:</div><div>On the Web:</div></div><div><div>1+800-284-5101</div><div>00800 4 QUANTUM</div><div>+44 1256 848 766</div><div>http://www.quantum.com/support</div></div>

Figure 174 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 cvfs_config(4) and mount_cvfs(1) 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 <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. <div><div>In the USA:</div><div>UK, France and Germany:</div><div>EMEA:</div><div>On the Web:</div></div><div><div>1+800-284-5101</div><div>00800 4 QUANTUM</div><div>+44 1256 848 766</div><div>http://www.quantum.com/support</div></div>

Figure 175 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/adic/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/adic/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:</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/adic/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/adic/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/adic/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/adic/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 <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. <div> <div>In the USA:</div> <div>1+800-284-5101</div> </div> <div> <div>UK, France and Germany:</div> <div>00800 4 QUANTUM</div> </div> <div> <div>EMEA:</div> <div>+44 1256 848 766</div> </div> <div> <div>On the Web:</div> <div>http://www.quantum.com/support</div> </div>

Figure 176 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. <div><div>In the USA:</div><div>UK, France and Germany:</div><div>EMEA:</div><div>On the Web:</div></div><div><div>1+800-284-5101</div><div>00800 4 QUANTUM</div><div>+44 1256 848 766</div><div>http://www.quantum.com/support</div></div>

Figure 177 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 QuantumTechnical Assistance Center. <div><div>In the USA:</div><div>UK, France and Germany:</div><div>EMEA:</div><div>On the Web:</div></div><div><div>1+800-284-5101</div><div>00800 4 QUANTUM</div><div>+44 1256 848 766</div><div>http://www.quantum.com/support</div></div>

Figure 178 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 fsmrpm 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

Figure 179 SNFS 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 <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 180 Port Failure

Recommended Actions	
SNFS Port Failure	
IF	THEN
SNFS cannot resolve the IP address of the coordinator <name server> after 600 seconds.	Check the name in the /sname servers file.
The name server <name server> (<IP address>) heartbeat is lost for 600 seconds.	The name server may not be reachable. Verify that the name server machine is running, can be reached on the network, and that SNFS is running on it.
The problem is <u>NOT</u> resolved:	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 181 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. <div> <div>In the USA:</div> <div>1+800-284-5101</div> <div>UK, France and Germany:</div> <div>00800 4 QUANTUM</div> <div>EMEA:</div> <div>+44 1256 848 766</div> <div>On the Web:</div> <div>http://www.quantum.com/support</div> </div>

Figure 182 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"> 1. Capture the StorNext system state. Refer to Capturing a System State. 2. Download the captured state. Refer to Downloading a System State Capture. 3. 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.
You receive a message that the FSM is delayed or the file system is not responding:	If you suspect a software bug, contact the Quantum Technical Assistance Center.
A health check operation is launched with invalid command arguments:	<p>Verify that the FSM process for the specified file system is running on the metadata controller. Also check the health of the metadata network.</p> <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"> 1. Locate the 'filelist' file under /usr/adic that contains the health check entry for the failed operation. 2. 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"> 1. 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 Problem. 2. 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"> 1. Try restarting the TSM software. (On the command line, run "tasmstop; tasmstart".) If restarting succeeds, rerun the Health Check using the StorNext GUI. 2. 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"> 1. Restart the system software. 2. Run the Health Check service using the StorNext GUI. 3. 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"> 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 183 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"> • If extending storage disk space, use the fsdiskcfg command to refresh the new capacity of the device. <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"> 1. If this is a new or temporary problem, run the fscshstate command to change the tape drive state. 2. If the problem persists, restart StorNext. 3. If the problem still persists, contact the __OEM_VENDOR_NAME__ Technical Assistance Center.
A tape drive is off-line:	<ol style="list-style-type: none"> 1. If this is a new or temporary problem, run the fscshstate command to change the tape drive state. 2. If the problem keeps occurring, contact the __OEM_VENDOR_NAME__ 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"> 1. Verify that the device is powered on. 2. Verify that the device has not registered any errors. 3. Verify that the device is physically connected by checking all cabling and connections. 4. 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 <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 __OEM_VENDOR_NAME__ Technical Assistance Center: In the USA: __OEM_SUPPORT_USA_PHONE__ UK, France and Germany: __OEM_SUPPORT_UK_FR_DE_PHONE__ EMEA: __OEM_SUPPORT_EMEA_PHONE__ On the Web: __OEM_SUPPORT_WEB__

Figure 184 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 185 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>vsarchivevary</code> to change the archive state to online.2. Run <code>Zsetstate</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">◦ Run <code>tsmascop; tsmasart</code> to reinitialize software communication pathways.◦ 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 186 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 187 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:	<div>1. Modify the ticket according to the troubleshooting steps taken. Refer to Analyzing Service Tickets.</div> <div>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</div> <div>OR</div> <div>If you are a properly-trained service professional, perform the procedures required for this type of tape library.</div>

Figure 188 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>snbackup -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>snbackup -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>ts_syssparm</code> file:	The <code>/usr/adc/TSN/config/ts_syssparm</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>snbackup</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>snbackup</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>snbackup -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/tsmlist</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>snbackup -s</code> or the <code>snbkreport</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 189 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"> 1. Stop StorNext. 2. Change the <code>/usr/adiac/DST/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. 3. 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"> 1. Restart the StorNext software system, which will correct the internal records of the device path and serial number. 2. 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"> 1. Use the StorNext GUI to delete and then re-add the drive. <p>OR</p> <ol style="list-style-type: none"> 2. Contact the Quantum Technical Assistance Center.
A configured tape drive has no detectable SCSI device path:	<ol style="list-style-type: none"> 1. Verify that the device is powered on. 2. Verify that the device has not registered any errors. 3. Verify that the device is physically connected by checking all cabling and connections. 4. 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>CfgDisk</code> element:	<ol style="list-style-type: none"> 1. Use the StorNext GUI to delete and then re-add the drive. <p>OR</p> <ol style="list-style-type: none"> 2. Contact the Quantum Technical Assistance Center.
<p>You receive one of the following "Invalid Media Configuration" messages:</p> <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. 	<p>Contact the Quantum Technical Assistance Center for help further diagnosing the invalid media configuration.</p>
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>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 190 Invalid
Configuration RAS part 2

A media ID that exists in TSM does not exist in MSM.	<div>1. Visually locate the medium in a specific archive.</div> <div>2. Run the <code>vsaudit</code> command and verify that the medium now appears as a New Medium.</div> <div>3. Run the <code>vazeclassify</code> command, converting the medium from the ADDBLANK to the DATA media class.</div> <div>If the medium does not physically exist in any archive or if the above commands fail, contact the Quantum Technical Assistance Center.</div>
The SNAPI configuration file is not valid.	Correct the errors in the configuration file <code>/usr/adic/SNAPI/config/snapi.cfg</code> .
The problem <u>IS</u> resolved.	Close the service ticket. Refer to Closing Service Tickets .
The problem has <u>NOT</u> been resolved.	<div>1. Modify the ticket according to the troubleshooting steps taken.</div> <div>Refer to Analyzing Service Tickets.</div> <div>2. Contact the Quantum Technical Assistance Center.</div> <div><div>In the USA:</div><div>UK, France and Germany:</div><div>EMEA:</div><div>On the Web:</div></div> <div><div>1-800-284-5101</div><div>00800 4 QUANTUM</div><div>+44 1256 848 766</div><div>http://www.quantum.com/support</div></div>

