

Quantum[®]

User's Guide

StorNext 4.7.x



Quantum StorNext 4.7.x User's Guide, 6-67947-01, Rev. D, January 2016, 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

© 2016 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.

StorNext utilizes open-source and third-party software. An enumeration of these open-source and third-party modules, as well as their associated licenses/attribution, can be viewed at www.quantum.com/opensource. Further inquiries can be sent to ip@quantum.com.



Contents

Chapter 1	Introduction	1
	Purpose of This Guide	1
	StorNext Gateway Terminology	1
	About StorNext File System.	3
	About StorNext Storage Manager.	3
	About StorNext LAN Clients	3
	StorNext Features	4
	How This Guide is Organized.	4
	Notes, Cautions, and Warnings	5
	Document Conventions	7
Chapter 2	StorNext GUI Overview	9
	Accessing the StorNext GUI	9
	StorNext Browser Support	9
	Internet Explorer 9 Security Settings.	10
	The StorNext Home Page	13
	StorNext Monitors	14
	StorNext Home Page Dropdown Menus	17

Chapter 3	The Configuration Wizard	23
	High Availability Systems	24
	Step 1: Welcome	24
	Step 2: Licenses	25
	Enabling and Installing Licenses	26
	How to Obtain Licenses	27
	Installing StorNext Licenses on a Windows MDC	28
	Installing StorNext Licenses from the GUI for non-HA Installations 28	
	Updating Licenses	29
	Installing StorNext Licenses from the GUI for HA Installations	31
	Updating Licenses	32
	Installing Optional StorNext Add-on Software Features and Clients from the StorNext MDC GUI	37
	Step 3: System (M660, M440 and M330 Metadata Appliances and Pro Foundation Only).	37
	Configuration Network Settings	38
	Entering Date and Time Information.	44
	Applying Your Changes.	46
	Step 4: Name Servers.	47
	Multiple fsnameservers Hosts and Redundant Metadata Networks	48
	Entering Name Servers	50
	Deleting a Name Server.	52
	Configuring a Foreign Server	52
	Step 5: File Systems.	54
	File Systems and Stripe Groups Limitations.	59
	Allocation Session Reservation	59
	Editing a File System and Performing Other Actions.	60
	Step 6: Storage Destinations	62
	Adding a New Library	62
	Viewing an Existing Library.	63
	Editing a Library	65
	Deleting a Library	66
	Performing Other Library Actions	66
	Storage Disk Overview	67
	Adding a New Storage Disk	68
	Viewing an Existing Storage Disks.	69
	Editing a Storage Disk.	70

Deleting a Storage Disk	71
Setting up Wide Area Storage Destinations	71
Adding a New Data Replication Target	71
Editing a Data Replication Host	72
Deleting a Data Replication Target	73
Adding a New Mount Point	73
Enabling Data Deduplication	74
Step 7: Storage Policies	74
Adding a Storage Manager Storage Policy	76
Enhanced Control of Tape Drive Allocation	79
Adding a Replication Storage Policy	82
Viewing a Storage Policy	82
Running a Storage Policy	83
Editing a Storage Policy	83
Deleting a Storage Policy	84
Step 8: Email Server	84
Adding an Email Server	85
Step 9: Email Notification	86
Adding an Email Recipient	86
Viewing Email Recipient Information	87
Editing an Email Recipient	88
Deleting an Email Recipient	88
Step 10: Done	89

Chapter 4

File System Tasks	91
Label Disks	92
Labeling a Device	92
Unlabeling a Device	94
Understanding Resource Allocation	95
About File System Expansion	95
About Stripe Group Movement	95
Expansion and Movement Steps	96
Using Resource Allocation From the Command Line	97
Checking the File System	97
Adding a Stripe Group Without Moving	97
Adding and Moving a Data Stripe Group	98
Moving a Metadata/Journal Stripe Group	99
Check File System	101

Viewing and Deleting a Check Report	103
File System Check Output Files	104
Affinities	105
Allocation Strategy	106
Example Use Cases	106
Adding a New Affinity	107
Deleting an Affinity	109
Migrate Data	109
How to Migrate Metadata and Journal Data	111
How to Migrate User Data	113
Truncation Parameters	115
Manage Quotas	117
Quota Limits	120
Quota Types	121
Renaming a Standalone (unmanaged) StorNext File System	122

Chapter 5

Storage Manager Tasks	125
Storage Components	126
Setting Devices Online and Offline	127
Additional Options for Tape Drives	127
Drive Pools	127
Viewing Drive Pool Information	128
Adding a Drive Pool	129
Editing a Drive Pool	130
Deleting a Drive Pool	130
Media Actions	131
Viewing Media Information	131
Filtering Media	132
Performing Media Actions	132
Storage Exclusions	142
Accessing Storage Exclusions	142
Adding an Exclusion Pattern	143
Truncation Exclusions	145
Accessing Truncation Exclusions	146
Adding an Exclusion Pattern	147
Tape Consolidation	149

Setting Tape Consolidation Parameters.	150
Scheduling Tape Cleaning and Defragmentation	152
Library Operator Interface	152
Software Requests.	154
Scheduler	155
Viewing a Schedule.	156
Adding a Schedule	157
Editing an Existing Schedule.	159
Deleting an Existing Schedule.	159
Alternate Retrieval Location.	160
Distributed Data Mover (DDM)	162
Distributed Data Mover Overview.	162
Installing the DDM Feature on Clients.	166
Accessing Distributed Data Mover	167
Enabling DDM.	168
Managing DDM Hosts.	168
Host Priority	171
Distributed Data Mover Reporting	171
Active Vault Policy	172
Introduction	172
Overview.	172
Configuring Active Vault.	174
Configuring External and Internal Vaults	175
Access and Retrieval of Media from Vaults	177
Command Syntax	179
Tools.	180
Usage Tips	181

Chapter 6	Replication and Deduplication	183
------------------	--------------------------------------	------------

Replication Overview	184
Replication Configuration Overview	184
Replication Process Overview	185
Files Excluded From Replication	186
Replication Terms and Concepts	187
Namespace Realization	187
Blockpool	188
Blackout Period	188
Replication Source Policy and Replication Source Directory	188

Replication Target Directory	188
Replication Schedule	189
Replication Copies	189
Bandwidth Throttling	189
Multilink	189
Virtual IP (vIP)	190
Some Replication Scenarios	190
Scenario 1: Simplest Replication	191
Scenario 2: Replicating Multiple Copies in the Same Target File System	191
Scenario 3: Replicating to Multiple Target Hosts / File Systems	193
Additional Replication Possibilities	194
Non-Supported Replication Between Source and Target	197
“Chained” Replication	197
Configuring Replication	198
Step 1: Create Source and Target File Systems	199
Step 2: Setting up the Blockpool	204
Step 3: Creating Replication Targets	206
Step 4: Create a Replication Storage Policy	207
Configuration Steps Summary	213
Scheduling Replication Blackouts (Optional)	214
Optional HA and Multilink Configuration	216
Running Replication Manually (Optional)	220
Replication Statuses and Reporting	221
Replication Reports	221
Replication Administration	221
StorNext Jobs	222
Replication Target Relocating Procedures	223
Replication Target Relocating Basics	223
Target Relocating Procedures	227
Perform Cross-Mount Initial Replication	228
Relocating the Target File System to the Target Site	231
Troubleshooting Replication	234
Data Deduplication Overview	234
How Deduplication Works	235
Deduplication and Replication	236
Setting Up Deduplication	236
Step 1: Creating a Deduplication-Enabled File System	236

- Step 2: Specifying the Blockpool 237
- Step 3: Creating a Deduplication-Enabled Storage Policy . . . 237
- Data Deduplication Functions 239
 - Deduplication Administration 239
 - Deduplication Reports 239
- Replication / Deduplication Removal Procedures 239
 - Assumptions 240
 - Removal Procedures 241
 - Collect and Understand Replication/Deduplication Configurations 241
 - Obtain Information from StorNext GUI 242
 - Obtain Information from Command Line 242
 - Replication Removal on a Target Host 245
 - Replication Removal on a Source Host 250

Chapter 7

Tools Menu Functions 253

- User Accounts 255
 - Adding a New User 255
 - Viewing an Existing User Profile 259
 - Modifying an Existing User 259
 - Deleting an Existing User 260
 - Enable or Disable User Accounts 261
- Client Download 262
- System Control 264
 - Starting or Stopping StorNext File System 265
 - Starting or Stopping StorNext Storage Manager 265
 - Refreshing System Status 265
 - Specifying Boot Options 265
- Lattus Certificates 266
 - HTTPS Default CA ROOT Certificate File or Path 266
 - How to Update Expired CA Root Certificates 267
 - New 271
 - View 274
 - Import 277
 - Convert 278
 - Download 279
 - Delete 280
 - Refresh 281
 - HTTPS Configuration 281

Creating a Single Public/Private SSL Certificate for use on Both a StorNext MDC and a Lattus System.	282
Lattus Certificates	284
HTTPS Default CA ROOT Certificate File or Path	284
How to Update Expired CA Root Certificates	285
New.	289
View.	292
Import.	295
Convert.	296
Download	297
Delete	298
Refresh	299
HTTPS Configuration	299
Creating a Single Public/Private SSL Certificate for use on Both a StorNext MDC and a Lattus System.	300
File and Directory Actions	301
Store Files	302
Change File Version	304
Recover Files	305
Recover Directories	306
Retrieve Files	309
Retrieve Directory	310
Truncate Files	310
Move Files	311
Modify File Attributes	312
View File Information	314
Assign Affinities	314
File Systems	316
Storage Manager.	316
Replication and Deduplication.	317
HA.	318
Upgrade Firmware.	318
Upgrade Considerations	319
Obtain the Firmware Upgrade Files.	319
Upgrade Times	320
Upgrade Procedure.	326
Deleting Uploaded Files	329
StorNext 5 Release 5.1.1 Post-Upgrade Required - Critical Fix329	

Chapter 8	Service Menu Functions	333
	Health Check	333
	Running a Health Check	334
	Viewing the Health Check Results	335
	Viewing Health Check Histories	336
	Capture State	336
	Creating a Capture State Log	337
	Deleting a Previous System State Capture	338
	Creating a Capture State for an HA Secondary Node	338
	Capture DSET	338
	Creating a Capture DSET File	339
	Downloading a Capture DSET File	339
	Deleting a Previous Capture DSET File	340
	System Backup	340
	Admin Alerts	341
	Tickets	343
	Viewing Ticket Information	343
	Changing the Display View	345
	Using Filter Options	345
	Editing Ticket Information	346
	Closing Tickets	347
	Deleting Tickets	347
	Logging	348
	Enabling Logging	348
Chapter 9	Converting to HA	351
	HA Overview	351
	HA Terms and Concepts	352
	Failover	352
	Primary Node	352
	Secondary Node	352
	Virtual IP (vIP)	353
	Virtual Netmask	353
	HA Reset	353
	Preparing for HA Conversion	354
	Pre-Conversion Steps	354

HA and LAN Clients	355
StorNext LAN Clients in HA Environments	355
Converting to HA	356
HA Conversion Procedure	357
Managing HA	361
HA Statuses and Reporting	363
Troubleshooting HA	363

Chapter 10

StorNext Reports	365
Report Navigation Controls	366
StorNext Logs	367
Administrative Logs and Alerts	368
How to Access StorNext Log Files	369
The Jobs Report	370
Viewing Detailed Job Information	371
Exiting the Jobs Report screen	372
The Files Report	373
The Drives Reports	375
The Media Report	377
The Relations Report	381
The File Systems Report	381
The SAN Devices Report	382
The Tape Consolidation Report	384
The SAN and LAN Clients Report	385
The LAN Client Performance Report	387
Replication Deduplication Reports	389
Policy Activity Report	389
Policy Summary Report	391
The Distributed Data Mover Report	394
The Hardware Status Report	396
The Gateway Metrics Report	400
Changing the Graph Display	403

Changing to Detail Mode 404
 Changing Displayed Columns 408
 Viewing and Changing Gateway Metrics Settings. 409

Chapter 11

Wide Area Storage (Lattus) 411
 Audience. 412
 Overview. 412
 Lattus Media versus Storage Manager Media 412
 Wide Area Storage Features in the StorNext (GUI). 413
 Configuring Lattus Object Storage 420
 The **fsobjcfg** Command 420
 How to Route Backup Files to Lattus for Easier Recovery . . . 422
 The **MAX_STORE_SIZE** System Parameter 422
 Setting Up Lattus Object Storage Destinations 424
 Viewing Lattus Object Storage Destinations 425
 Adding a New Lattus Object Storage Destination 425
 Editing a Lattus Object Storage Destination 432
 Deleting a Lattus Object Storage Destination 432
 Performing Other Lattus Object Storage Destination Actions . . . 433
 Changing the Current State of Lattus Object Storage Destinations, Controllers, and I/O Paths 434
 Special Considerations for Multi-Geo Configurations 434
 HTTPS Support for Lattus Object Storage 435
 The **FS_OBJSTORAGE_CACERT** System Parameter 435
 The **FS_OBJSTORAGE_CAPATH** System Parameter 436
 The **FS_OBJSTORAGE_SSL_VERIFY_PEERHOST** System Parameter 436
 Configuring HTTPS on DDM Hosts 436
 HTTPS Support for Q-Cloud Archive™ Object Storage 438
 Changes to Existing CLI Commands. 438
 The **fsfileinfo** Command 438
 The **fsmedinfo** Command 439
 The **fsmedread** Command 439
 The **fsmedscan** Command 439
 The **dm_info** Command 439
 The **dm_util** Command 439
 Other Changes and Considerations 440

Wide Area Storage Segment Size	440
snpolicyd	441
Storage Manager	442

Chapter 12	Customer Assistance	443
	StorNext Upgrades	443
	Contacts	443
	Comments	444
	Getting More Information or Help	444
	Worldwide End-User Product Warranty	445

Appendix A	Operating Guidelines	447
	The Reserved Space Parameter	447
	Linux Configuration File Format	448
	Gateway Server/Client Network and Memory Tuning	449
	Gateway Server and Client Network Tuning	449
	Gateway Server Memory Tuning	450
	Configuring LDAP	451
	Using LDAP	451
	UNIX File and Directory Modes	452
	LDAP Refresh Timeout	452
	Setting Up Restrictive ACLs	453
	Default Single-Path I/O Retry Behavior	453
	Event Handles for fsm.exe on a Windows Metadata Server	453
	FSBlockSize, Metadata Disk Size, and JournalSize Settings	454
	Disk Naming Requirements	456
	Changing StorNext's Default Session Timeout Interval	457
	Configuring a Data Partition for Use with Spectra Logic T-series Tape Storage Libraries	458
	Configuring the Partition	458
	Basic Secure Sockets Layer (SSL) Guidelines	460
	Example of a server.pem File	464
	Example of a public.pem File	465
	Name Limitations	465

Ports Used By StorNext 467

Log Rolling and Disk Space Health Check 469

 Log Rolling 469

 Disk Space Health Check 472

General Operating Guidelines and Limitations 473

Appendix B Additional Replication and Deduplication Information 493

Replication Configuration File 493

Replication Terminology and Conventions. 494

Copies in Replication Versus Copies and Versions in Storage Manager 494

Replication Target Directories 497

 Number of Replication Copies. 497

 Isolating a Replication Target Directory. 498

 Final Recommendation For Target Directories. 500

StorNext snpolicyd Policies 500

Replication Copies = 2 (Detail) 503

More About Replication Target Directories 507

Deduplication Overview. 509

 Enabling Deduplication. 511

 Deduplication Modification Time 511

 Deduplication and Blockpools. 511

Deduplication and Truncation 512

 Enabling Deduplication and Truncation 513

 Storage Manager Truncation 513

Replication, Deduplication and Truncation 513

Replication, Deduplication and Storage Manager 514

 Replicating into a Storage Manager Relation Point. 515

 Truncation and Deduplication / Replication (with and without SM). 517

The snpolicyd Debug Log 525

Appendix C High Availability Systems 527

High Availability Overview	528
HA Internals: HAmom Timers and the ARB Protocol	530
Primary and Secondary Server Status	534
File System Types	534
The <code>ha_peer</code> and <code>fsnameservers</code> File	536
HA Manager	536
Configuration and Conversion to HA.	543
Conversion to HA	544
SyncHA process.	546
Managing High Availability in the StorNext GUI	547
Configuring Multiple NICs	550
LAN Configuration	550
High Availability Operation.	551
Windows and Linux SNFS Installations Without the HaShared File System	552
Linux SNMS and SNFS Installations with the HaShared File System	553
HA Resets	559
HA Resets of the First Kind	559
HA Resets of the Second Kind	560
HA Resets of the Third Kind	560
Using HA Manager Modes	561
HA Tracing and Log Files	561
Single (Singleton) Mode	562
FSM Failover In HA Environments	563
Failover Timing	563
Replacing a HA System	567
Moving a HA Shared File System to a New Raid	568

Appendix D

Web Services API	573
Using the APIs	574
Using APIs With the High Availability MDC Feature	575
WS-API APIs.	576
The <code>doCancel</code> API	576
The <code>doMediaMove</code> API	577

<hr/>		
Appendix F	Security	621
	StorNext Security	621
	ACLs on Windows	622
	ACLs on Mac OS X	626
	“Central Control”	629
	Limitations	632
	Example	632
	Cross-Platform Permissions	633
	Config (.cfg) File Options	633
<hr/>		
Appendix G	Troubleshooting	637
	Troubleshooting StorNext File System	637
	Troubleshooting OS Issues	644
	Troubleshooting Replication	647
	Troubleshooting HA	649
	Troubleshooting StorNext Installation and Upgrade Issues	655
	Troubleshooting Other Issues	655
<hr/>		
Appendix H	StorNext	
Offline Notification	659	
	StorNext Offline Notification Overview	660
	What is a StorNext Offline File?	660
	Why Block Access to Offline Files?	660
	Offline Notification Configuration Options	661
	How the Notification Feature Operates	661
	Installing the Notification Application	663
	Installing onto a Standalone Machine	663
	Starting the Notification Application	669
	Configuring the Notification Application	670
	Setting Application Options	670
	Viewing Access Denied Files	672
	Viewing the Application Log	673
	Exiting Application Preferences	673
	Uninstalling the Notification Application	674

Appendix I	RAS Messages	675
Appendix J	Repairing and Replacing StorNext Metadata Servers	677
	Replacing an MDC in a non-HA environment (non-backup/restore method)	678
	Replacing an MDC in non-HA environment (backup/restore method)	681
	Replacing an MDC in an HA environment	684



Figures

Figure 1	StorNext Login Window	11
Figure 2	System Control	12
Figure 3	StorNext Home Page.	14
Figure 4	Configuration > Configuration Wizard Welcome Screen .	25
Figure 5	Configuration > Licenses Screen	26
Figure 6	Configuration > Licenses > Agreement Screen	29
Figure 7	Configuration > Licenses Screen	30
Figure 8	Configuration > Licenses > Agreement Screen	31
Figure 9	Configuration > Licenses Screen	33
Figure 10	Configuration > System Screen	38
Figure 11	Configuration > System > Date & Time Screen	44
Figure 12	Configuration > System > Date & Time Screen (Manual Entry) ⁴⁶	
Figure 13	Name Servers Screen	51
Figure 14	StorNext Foreign Servers.	54
Figure 15	Configuration > File System Screen	55
Figure 16	Configuration > File System > New Screen	56
Figure 17	Configuration > File System > New Screen 2.	57

Figure 18	Configuration > File System > New Screen 3.	58
Figure 19	Configuration > File System > New Screen 3.	59
Figure 20	Storage Destinations > Library Screen	62
Figure 21	Storage Destinations > Library > New Screen.	63
Figure 22	Library Details Screen	64
Figure 23	Edit Library Screen	65
Figure 24	Configuration > Storage Destinations > Storage Disk Screen 68	
Figure 25	Storage Destinations > Storage Disk > New Screen.	69
Figure 26	View Storage Disk Screen	70
Figure 27	Configuration > Storage Destinations > Replication Targets / New Screen72	
Figure 28	Configuration > Storage Destinations > Replication / Deduplication Screen (Blockpool)74	
Figure 29	Configuration > Storage Policies Screen	76
Figure 30	Storage Policies > New Screen	77
Figure 31	Storage Policies > New > General Tab.	78
Figure 32	Storage Policies > New > Relocation Tab	79
Figure 33	Storage Policies > New > Steering Tab	80
Figure 34	Storage Policies > New > Schedule Tab.	81
Figure 35	Storage Policies > New > Associated Directories Tab.	82
Figure 36	View Storage Policies Screen	83
Figure 37	Configuration > Email Server Screen	85
Figure 38	Configuration > Email Notifications Screen	86
Figure 39	Configuration > Email Notifications New Screen	87
Figure 40	Configuration > Configuration Wizard Done Screen	89
Figure 41	Label Disks Screen	93
Figure 42	Check File System Screen	103
Figure 43	Check File System Report	104
Figure 44	Affinities Screen	108

Figure 45	New Affinity Screen	108
Figure 46	Migrate Data page	110
Figure 47	Metadata/Journal Data migration page	112
Figure 48	Metadata/Journal Data migration (new page)	112
Figure 49	User Data migration page.	114
Figure 50	Truncation Parameters Screen	116
Figure 51	Manage Quotas Screen.	117
Figure 52	Edit Quotas Screen	119
Figure 53	New Quotas Screen	119
Figure 54	Storage Components Screen	126
Figure 55	Drive Pools Screen	128
Figure 56	New Drive Pool Screen	129
Figure 57	Media Actions Screen	131
Figure 58	Storage Exclusions Screen.	143
Figure 59	Exclusion Screen	143
Figure 60	Truncation Exclusions Screen	147
Figure 61	Exclusion Screen	147
Figure 62	Tape Consolidation Screen	150
Figure 63	Library Operator Interface Screen	153
Figure 64	Software Requests Screen.	155
Figure 65	Scheduler Screen	157
Figure 66	Scheduler > New Screen	158
Figure 67	Alternate Retrieval Location Screen	161
Figure 68	Configuration > Distributed Data Mover Screen	167
Figure 69	DDM Screen New Host	169
Figure 70	Replication Process	186
Figure 71	Replication scenario 1.	191
Figure 72	Replication Scenario 2	192
Figure 73	Replication Scenario 3	194

Figure 74	Replicating From One Source to Multiple Targets	195
Figure 75	Replicating From Multiple Sources to One Target	196
Figure 76	Non-Supported Replication From Source to Target	197
Figure 77	Configuration > File System > New Screen	200
Figure 78	Configuration > File System > New Screen 2.	201
Figure 79	Configuration > File System > New Screen 3.	202
Figure 80	Configuration > Storage Destinations > Deduplication Screen (Blockpool)	205
Figure 81	Storage Destinations > Replication Targets Screen	206
Figure 82	Configuration > Storage Policies > New Screen	208
Figure 83	Configuration > Storage Policies > New / Source Directories Screen	209
Figure 84	Storage Policies > New > Outbound Replication Tab	210
Figure 85	Outbound Replication Tab > Replication Schedule.	211
Figure 86	Configuration > Storage Policies Screen (Select "target")	212
Figure 87	Storage Policies > Edit > target > Inbound Replication Tab	213
Figure 88	Configuration > Storage Policies (Run Policy).	214
Figure 89	Storage Policies > New > Blackout Tab	215
Figure 90	Tools > HA Convert Screen.	217
Figure 91	Tools > Replication > Bandwidth Screen	219
Figure 92	Tools > Replication/Deduplication > Administration Screen	222
Figure 93	Deduplication	235
Figure 94	Replication/Deduplication Policy Screen	238
Figure 95	User Accounts Page	256
Figure 96	New User Page	257
Figure 97	View User Page.	259
Figure 98	Edit User Page	260
Figure 99	Client Download Page	263

Figure 100	Client Download Link	263
Figure 101	System Control Page.	264
Figure 102	Lattus Certificates Page	269
Figure 103	Creating a New self-signed Lattus Certificate	272
Figure 104	Viewing a Lattus Certificate	275
Figure 105	Lattus Certificates Page	287
Figure 106	Creating a New self-signed Lattus Certificate	290
Figure 107	Viewing a Lattus Certificate	293
Figure 108	File and Directory Action Page	303
Figure 109	Change File Version Page	304
Figure 110	Recover Files Browse Page	305
Figure 111	Recover Directories Browse Page	307
Figure 112	Retrieve Files.	309
Figure 113	Retrieve Directory Page.	310
Figure 114	Truncate Files Page.	311
Figure 115	Move Files Page	312
Figure 116	Modify File Attributes Page	313
Figure 117	View File Info Page	314
Figure 118	Assign Affinities Page	315
Figure 119	Firmware Upgrade Page	326
Figure 120	Health Check Screen	334
Figure 121	Health Check > View Selected Screen	335
Figure 122	Health Check > View History Screen.	336
Figure 123	Capture State Screen	337
Figure 124	Capture DSET Screen	339
Figure 125	Backup Screen	341
Figure 126	Admin Alerts Screen	342
Figure 127	Tickets Screen	343
Figure 128	Tickets > View Ticket Screen	344

Figure 129 Tickets > Filter Options Screen	346
Figure 130 Tickets > Edit Ticket Screen	347
Figure 131 Logging Screen	349
Figure 132 Tools > HA > Convert (primary node not yet converted)	358
Figure 133 Tools > HA > Convert (primary node previously converted). 359	
Figure 134 Manage HA Screen	362
Figure 135 Reports > Logs Screen	369
Figure 136 Jobs Report.	370
Figure 137 Files Report.	373
Figure 138 StorNext File Browser	374
Figure 139 File Info Screen	375
Figure 140 Drives Report	376
Figure 141 Drive Information Report	377
Figure 142 Media Report	378
Figure 143 Media Information Report	380
Figure 144 Relations Report	381
Figure 145 File Systems Report.	382
Figure 146 SAN Devices Report	383
Figure 147 Tape Consolidation Report	384
Figure 148 SAN and LAN Clients Report.	386
Figure 149 LAN Client Performance Report	388
Figure 150 Replication/Deduplication Policy Activity Report.	390
Figure 151 Replication/Deduplication Policy Summary Report.	392
Figure 152 Replication/Deduplication Policy Details Report.	393
Figure 153 Replication/Deduplication Policy Completion Report	394
Figure 154 Distributed Data Mover Report.	395
Figure 155 Hardware Status Report System Board Tab.	396
Figure 156 Hardware Status Report Network Ports Tab	397

Figure 157 Hardware Status Report Fibre Channel Ports Tab 398

Figure 158 Hardware Status Report Storage Arrays Tab 399

Figure 159 Gateway Metrics Report (Summary) 402

Figure 160 Gateway Metrics Report (Detail) 405

Figure 161 Gateway Metrics Report (Client Detail) 407

Figure 162 Gateway Metrics Report (File System Detail) 408

Figure 163 Gateway Metrics Settings 410

Figure 164 Home Page 414

Figure 165 Configuration > Licenses Page 415

Figure 166 Configuration > Storage Manager Policies > Edit Page . 416

Figure 167 Tools > Storage Manager > Distributed Data Mover Page. .
417

Figure 168 Tools > Storage Manager > Storage Components Page 418

Figure 169 Tools > Storage Manager > Media Actions (Media Selection)
Page419

Figure 170 Tools > Storage Manager > Media Actions (Available Media
Actions list) Page419

Figure 171 Configuration > Storage Destinations > Lattus Object
Storage424

Figure 172 High Availability Manage Screen 548

Figure 173 FSM Failover in an HA Cluster 564

Figure 174 Run as Administrator 664

Figure 175 Logging in to the Administrator Account 664

Figure 176 Installing the .NET Framework 665

Figure 177 Offline Notification Setup Wizard 666

Figure 178 Quantum End User License Agreement 666

Figure 179 Select Installation Folder 667

Figure 180 Confirm Installation 668

Figure 181 Installation Complete 668

Figure 182 Manual Start 669

Contents

Figure 183 Application Options	671
Figure 184 Access Denied List	672
Figure 185 Application log	673
Figure 186 Removing the Application	674



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.

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.

StorNext Gateway Terminology

For the purposes of this document, we will use the following terminology:

StorNext Gateway Term	Description	Historical Customer-configured Gateway Equivalent Terminology
StorNext Gateway	A StorNext Gateway is a StorNext SAN Client which allows LAN-based client connectivity to a StorNext File System.	Gateway Server; Server; LAN server, LAN-based server; DLC Gateway server; Clustered Gateway; DLC Gateway; DLS
StorNext LAN Client	A LAN-connected computer attached to a StorNext Gateway that has shared access to a StorNext SAN.	StorNext DLC
StorNext Gateway Metrics	A performance reporting and monitoring software module for StorNext Gateways.	N/A, newly created for StorNext Gateway

How the StorNext Gateway license is enabled depends on the current configuration:

- The StorNext G300 Gateway Appliance and the StorNext M660 Metadata Appliance have a “per Gateway” license model. This license allows you to add clients without having to purchase additional individual licenses.
- For new customers with no existing StorNext components, the license comes from the factory pre-installed and enabled for use with the StorNext G300 Gateway Appliance and the StorNext M660 Metadata Appliance.
- For customers with existing customer-configured MDCs, if you choose to install the StorNext G300 Gateway Appliance or the StorNext M660 Metadata Appliance with the Gateway feature enabled in the same StorNext configuration as a customer-configured DLC gateway, you will be limited to the existing client DLC license count.

Note: The Gateway license is located on the StorNext G300 Gateway Appliance and the StorNext M660 Metadata Appliance. To determine whether existing StorNext Gateway licenses are enabled, click the **Connected Licensed Gateways** link at the bottom of the StorNext license screen on the associated MDC.

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.

Note: The maximum supported file system size is 1 Billion files per File System. *As of StorNext 4.3.0, the database will handle up to 1 Billion files per MDC, and 1 Billion files per file system.*

Note: The maximum supported filename is 255 bytes. The maximum supported path length is 1023 bytes. In Linux, paths may be longer than 1023 bytes, but such paths are not compatible with certain StorNext features including Storage Manager, Directory Quotas, and Replication.

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 StorNext LAN Clients

In addition to supporting StorNext clients attached via fibre channel, StorNext also supports LAN clients. Unlike a direct-attached StorNext SAN client, a LAN client connects across a LAN through a gateway

server, which includes the StorNext G300 and StorNext M660. Gateway servers are themselves directly-connected StorNext SAN clients. The StorNext M660 is a Metadata Controller (MDC) which can also be licensed to function as a gateway server. Gateway servers process requests from LAN clients in addition to running applications.

For more information about StorNext licensing, see [Step 2: Licenses](#) on page 25, and the *StorNext Licensing Guide*.

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

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

StorNext Features

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 25.

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, Wide Area Storage \(Lattus\)](#)
- [Chapter 12, 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](#)
- [Appendix J, Repairing and Replacing StorNext Metadata Servers](#)

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.

Description	Definition	Consequences
WARNING:	Highly recommends you to take or avoid a specified action.	Failure to take or avoid this action WILL result in loss of data.

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.

StorNext Browser Support

StorNext browser requirements are listed in the *StorNext Compatibility Guide* posted here (click the “Select a StorNext Version” menu to view the desired documents): <http://www.quantum.com/sndocs>

Note: 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.

Internet Explorer 9 Security Settings

Some Internet Explorer 9 default security settings could prevent StorNext from operating properly with this browser. Follow the procedure below to enable security options in Internet Explorer 9:

- 1 Launch Internet Explorer 9.
- 2 On the **Tools** menu, click **Internet Options**.
- 3 Click the **Advanced** tab, and then enable the following:
 - **Run ActiveX Controls and Plugins**
 - **Script ActiveX Controls Marked Safe**
 - **Active Scripting**
 - **File Download**
 - **Don't Prompt for Client Certificate...**
- 4 Click the **Security** tab, and then make sure:
 - The option **Do not save encrypted pages to disk** is not checked.
 - The option **Display mixed content** is checked.
- 5 Click **OK**.
- 6 Close Internet Explorer 9 and then open it again to continue.

Use this procedure to access the StorNext GUI.

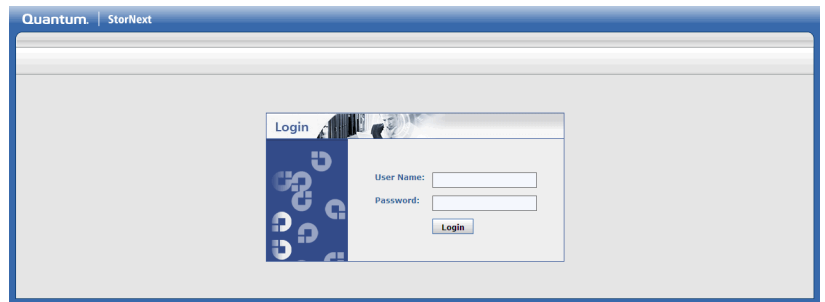
- 1 Open a Web browser.
- 2 In the browser's **Address** field, type the full address of the machine and its port number, and then press **Enter**. For example: `http://<machine name>:<port number>`. Use the name of the machine and port number you copied when you installed the StorNext software.

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

Note: The StorNext GUI may be inaccessible in a Web browser, with one of the following error messages displayed:
For Firefox: Unable to connect. Firefox can't establish a connection to the server.
For Internet Explorer: Internet Explorer cannot display the web page.
If you encounter either of these conditions, restart the StorNext GUI on the MDC server by performing the following:
Open a root UNIX shell window on the MDC, then run the command `service stornext_web restart`. The `service` command returns before the service is ready to be accessed by a browser. Wait a few moments before trying to connect, and then retry if that fails.

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

Figure 1 StorNext Login Window

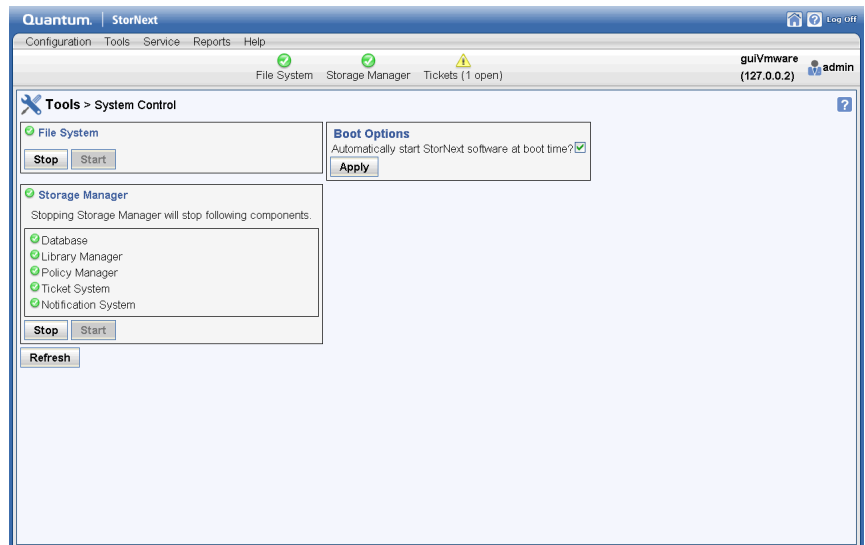


- 3 In the **User ID** field, type admin.
- 4 In the **Password** field, type password.

- 5 Click **Login**. The StorNext **Home** page appears (refer to [The StorNext Home Page](#) on page 13).

Note: If the StorNext File System and Storage Manager components are not currently started, the StorNext **Tools > System Control** screen appears. 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. Click the home (house) icon in the upper right corner to go to the StorNext **Home** Page.

Figure 2 System Control



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 login sessions.

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.

Note: The capacity indicators on the StorNext home page provide approximations 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.

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 LAN Clients for which you are licensed. For more information about LAN Clients, see [About StorNext LAN Clients](#) on page 3.
- 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
- **System:** Manage the cluster network interfaces between HA nodes, NTP, DNS, hostnames, gateway, domain and set bonding options. (This option is only visible on StorNext M660, M440 and M330 Metadata Appliances.)
- **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
- **Firmware Upgrade:** Download StorNext upgrades. (This option is only visible on M660, M440 and M330 Metadata Appliances.)
- **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
 - **Manage Quotas:** Limit the amount of disk storage consumed on a per user, or per group basis across an entire file system, or within a designated directory hierarchy.
- **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
- **Capture DSET:** Obtain and preserve detailed information about the current state of hardware. (This option is only visible on StorNext M660, M440 and M330 Metadata Appliances.)
- **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 LAN clients, and client performance
- **LAN Client Performance:** View information about 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
- **Gateway Metrics:** View performance statistics for gateways, clients and file systems

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. The Gateways tab provides information about your gateway(s), including server name, serial number, version and gateway license number.



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 with a green check mark. 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: System \(M660, M440 and M330 Metadata Appliances and Pro Foundation Only\)](#): Configure network settings.
(This step only applies to StorNextM660, M440 and M330 Metadata Appliances.)
- [Step 4: Name Servers](#): Specify and order the machines acting as StorNext name servers
- [Step 5: File Systems](#): Add a StorNext file system
- [Step 6: Storage Destinations](#): Add a library, storage disks, and other storage destinations. Additionally, you can add Wide Area Storage Destinations (see [Setting Up Lattus Object Storage Destinations](#) on page 424)
- [Step 7: Storage Policies](#): Add a Storage Manager or replication storage policy
- [Step 8: Email Server](#): Specify an email server to handle StorNext notifications
- [Step 9: Email Notification](#): Add email notifications recipients
- [Step 10: 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 354.


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.

[Step 4: Name Servers](#) on page 47

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.

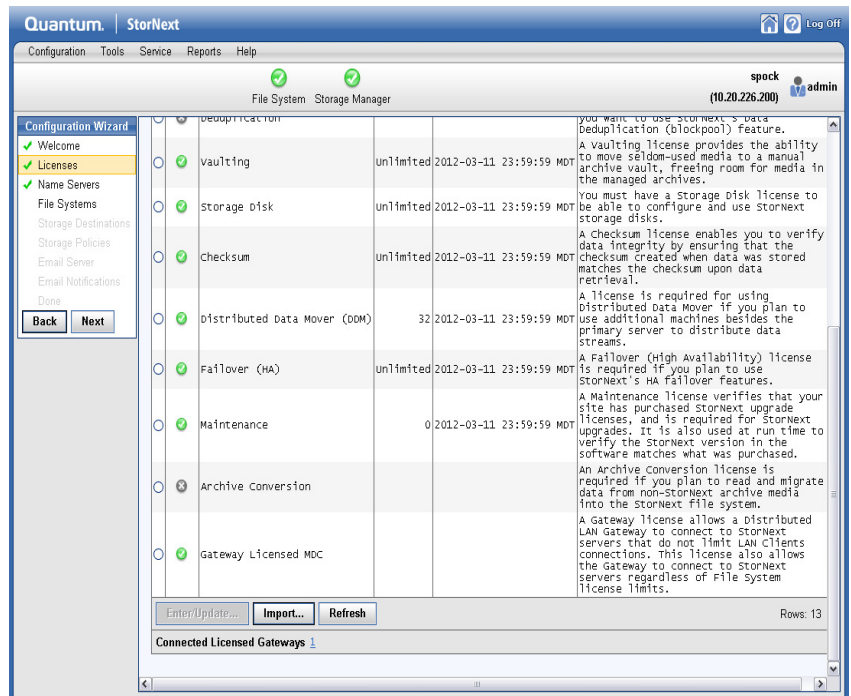
Figure 4 Configuration > Configuration Wizard Welcome Screen



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



Enabling and Installing Licenses

Some StorNext licenses come pre-enabled on your system, and others must be purchased and then enabled by installing a license key file on your system. You must do this to obtain your correct maintenance expiration date. If your system is using a 30-day an auto-generated license, you should enter permanent licenses as soon as it is convenient. The following are the steps you need to take to obtain and install your StorNext licenses.

Note: 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.

How to Obtain Licenses

StorNext feature licenses must be enabled by installing a license file on the system. If your system is using a 30-day auto-generated license, you should enter permanent licenses as soon as it is convenient, so that the maintenance expiration date for the licenses is set correctly on the system. (You can determine the current license limits for many StorNext features, such as SAN clients, LAN clients, DDM, and Storage Manager licenses, in the StorNext GUI on the **Configuration > Licenses** screen.)

If you need license keys for StorNext features or capacity, enter the required information about your system at:

<http://www.quantum.com/ServiceandSupport/License/StorNext/Index.aspx>

Note: If you cannot access the web page or need additional help filling out the form, contact Quantum support at <http://www.quantum.com/ServiceandSupport>

In order to receive a `license.dat` file, you will need the following:

- The System Serial Number. (Use this in the “Product Information, Serial number of the Original Media” section of the form)

System serial numbers are alpha-numeric. (Example:
SV1728CKH02059)

- If you don't have the StorNext serial number, you can use the StorNext serial number instead. For new installations, the serial number is delivered via email. If you are adding to an existing StorNext installation and cannot locate the serial number, you can find it in the `license.dat` file. The file can be found on the MDCserver node at `/usr/cvfs/config/license.dat` for UNIX systems, and `c:\Program Files\StorNext\config\license.dat` for Windows

systems. Open your current `license.dat` file and locate the serial number.

Example:

```
# Serial Number:      Q8574321
```

- The `cvfsid` string for MDCserver node (if you have an HA system, you will need the `cvfsid` string for both nodes).

Example:

```
ECF4BCEECC0E linux 0 xcellis13
```

- A list of StorNext features already licensed and enabled on your system.
- A list of purchased StorNext add-on features for which you wish to enable licenses.

After the Quantum Technical Assistance Center receives the above information, a representative will send you a `license.dat` file (which contains license keys for the products/features you specified) within one business day. Save the file to a temporary location to apply to the system.

Installing StorNext Licenses on a Windows MDC

See the help files installed with StorNext.

The help files are located in the **Start** menu at **All Programs > StorNext File System > StorNext Help**.

Installing StorNext Licenses from the GUI for non-HA Installations

Note: Use these steps to install the license keys on a single-node system, or one that has not yet been configured as HA. For HA systems, see [Installing StorNext Licenses from the GUI for HA Installations](#) on page 31.

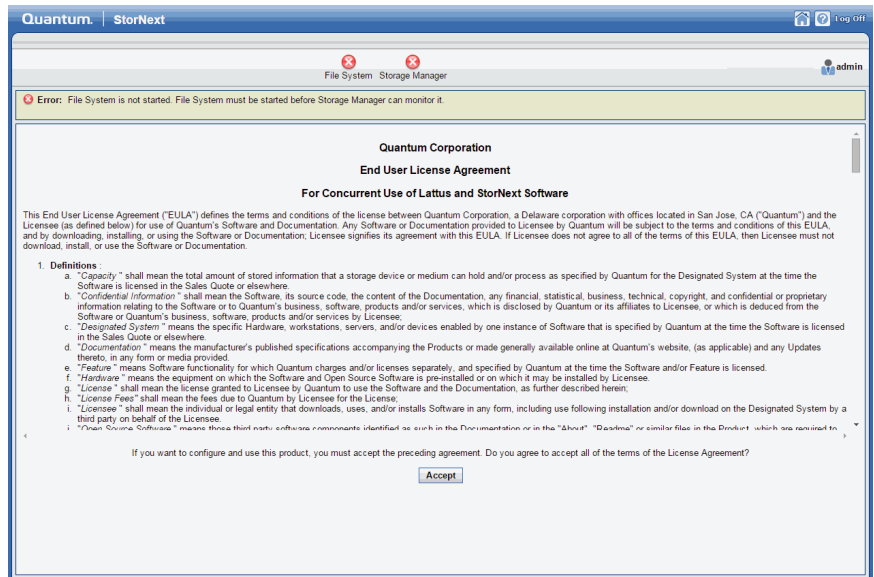
To install the license keys on a single-node system, or one that has not yet been configured as HA.

Do one of the following, depending on how you are applying the StorNext license:

If you are logging into StorNext for the first time:

- 1 Read the end-user license agreement carefully, and then click **Accept**. The **Configuration > Licenses Entry** screen appears. Continue to [Step 3](#).

Figure 6 Configuration > Licenses > Agreement Screen



Updating Licenses

You will need to update a license if the license expires or if your configuration changes (for example, if you add additional clients or increase capacity).

To update a license, select the desired product/feature and then click **Enter/Update**. When the **Configuration > Licenses > Enter/Update** screen appears, copy the license string you received from Quantum into the **[License/Feature Name] License** field, and then click **Apply**.

If you are updating StorNext licenses:

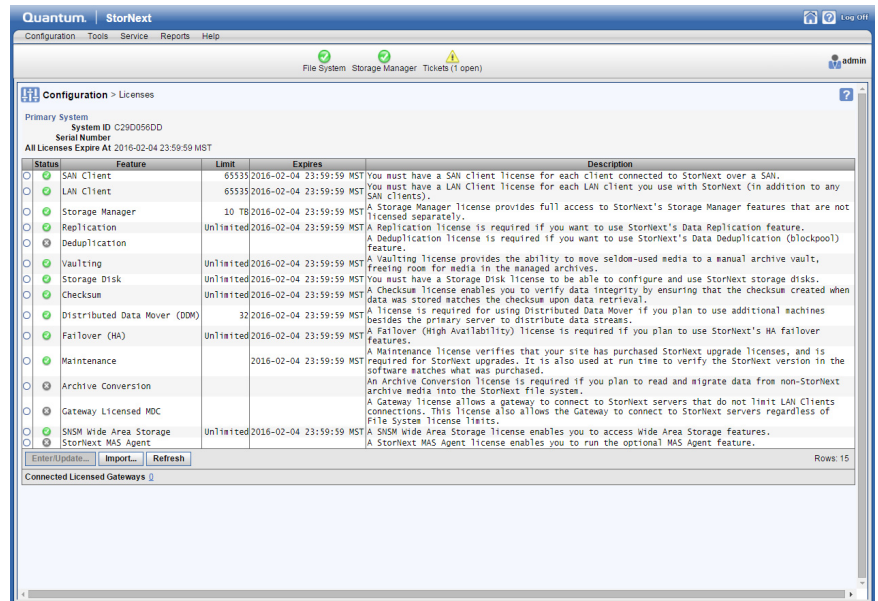
- 1 Stop StorNext services by navigating to **Tools > System Control** option.

Caution: You can update **SAN Client** licenses, **LAN Client** licenses, and **Disk** licenses (**Certified** and **Uncertified**) without restarting the StorNext services. All other license types will require a restart of the StorNext services.

2 Navigate to **Configuration > Licenses > Enter/Update**.

Continue the installation

Figure 7 Configuration > Licenses Screen



3 To import the license key from a text file, click **Import**. When the **Import License File** window appears, click **Browse** and navigate to the file's location. Click **Close** to continue.

A message at the top of the screen informs you whether the information was successfully imported and copied from the file into the StorNext license file. The status indicator changes to enabled (a green check mark icon) and you are now ready to activate your licenses.

- 4 After you install your licenses, you should restart StorNext services by using the **Tools > System Control** option. StorNext may not recognize some feature licenses until you restart services.

Installing StorNext Licenses from the GUI for HA Installations

Note: Use these steps to install the license keys on an MDC configured for HA. For non-HA systems, see [Installing StorNext Licenses from the GUI for non-HA Installations](#) on page 28.

Do one of the following, depending on how you are applying the StorNext license:

If you are logging into StorNext for the first time:

- 1 Read the end-user license agreement carefully, and then click **Accept**. The **Configuration > Licenses Entry** screen appears. Continue to [Step 5](#).

Figure 8 Configuration > Licenses > Agreement Screen



Updating Licenses

You will need to update a license if the license expires or if your configuration changes (for example, if you add additional clients or increase capacity).

To update a license, select the desired product/feature and then click **Enter/Update**. When the **Configuration > Licenses > Enter/Update** screen appears, copy the license string you received from Quantum into the **[License/Feature Name] License** field, and then click **Apply**.

If you are updating StorNext licenses:

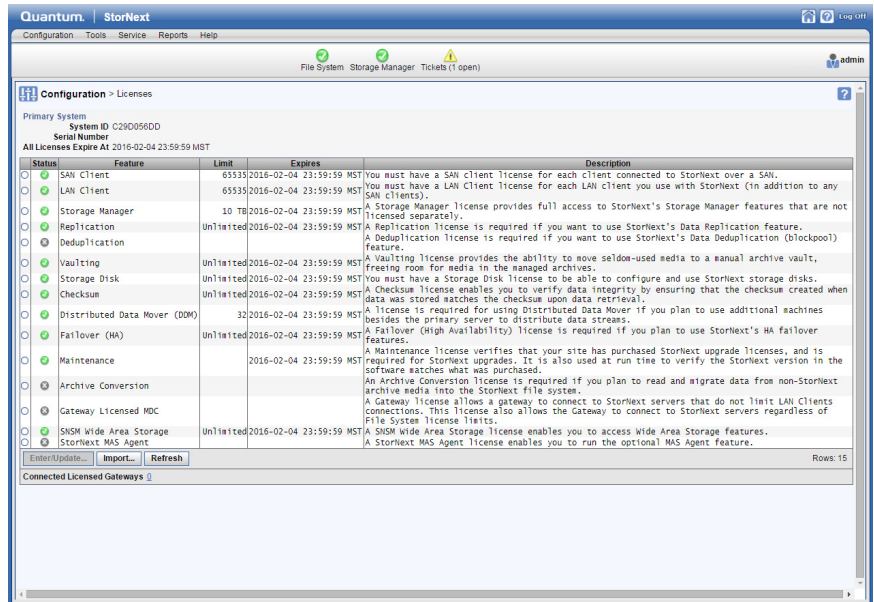
- 1 Place the system into config mode by navigating to **Tools > High Availability > Manage**.
- 2 Select **Enter Config Mode** button.

This will pop up a confirmation to attempt to lock the HA cluster configuration.

Caution: You can update **SAN Client** licenses, **LAN Client** licenses, and **Disk** licenses (**Certified** and **Uncertified**) without restarting the StorNext services. All other license types will require a restart of the StorNext services.

- 3 Choose “yes” when prompted.
See the appendix **High Availability Systems** in the *StorNext User’s Guide* (appropriate for your release of StorNext) for further information about HA systems.
- 4 Navigate to **Configuration > Licenses > Enter/Update**.

Figure 9 Configuration > Licenses Screen



5 To import the license key from a text file, click **Import**. When the **Import License File** window appears, click **Browse** and navigate to the file's location. Click **Close** to continue.

A message at the top of the screen informs you whether the information was successfully validated and copied into the StorNext license file. The status indicator changes to enabled (a green check mark icon) and you are now ready to activate your licenses.

6 Navigate back to **Tools > High Availability > Manage** and press **Exit Config Mode**. You will be prompted for confirmation that you want to unlock the cluster. Click **Yes** to exit config mode.

For File System installations, or in instances where there is no available StorNext GUI, StorNext license keys for the system can be installed manually. This section describes how to install licenses from the command line in those situations.

Note: Licenses must be installed “by hand” on RPM-only installations because there is no GUI to lead you through the steps.

Before installing a **license.dat** file ensure that exactly one key is present in the file for each licensed feature, and that a Maintenance license key

is present. The new license file must be complete before proceeding to the next steps. If you are updating a license file, ensure that old keys are removed so that they do not interfere with the updated license keys.

If you are installing a license file for HA MDCs, make sure that the keys for both MDCs are present in the **license.dat** file before proceeding.

Depending on the configuration of the system you are updating, you will use a different set of installation steps to install the license keys for StorNext and components from the command line. Here are the possible installation scenarios:

- [Installing StorNext License Keys on a Single-Node MDC](#) on page 34
- [Installing StorNext License Keys on HA Systems \(with a Shared File System\)](#) on page 35
- [Installing StorNext License Keys on a Failover System \(without a Shared File System\)](#) on page 36

Installing StorNext License Keys on a Single-Node MDC

- 1 From the command line of the MDC, type the following command:

```
service cvfs stop
```

Caution: The command **service cvfs stop** will stop all StorNext services on the MDC.

- 2 Navigate to the temporary location where you stored the new **license.dat** file.
- 3 Copy the **license.dat** file to the **/usr/cvfs/config/** directory on the MDC.
- 4 Type the following command on the MDC:

```
service cvfs start
```

Installing StorNext License Keys on HA Systems (with a Shared File System)

Note: Stop the StorNext service on the **secondary** node of the HA system **before** stopping the StorNext service on the **primary** node.

- 1 From the command line of the **secondary** node of an HA MDC (both File System and Storage Manager), type:

```
service cvfs stop
```

Caution: The command **service cvfs stop** will stop all StorNext services on the MDC.

- 2 From the command line of the **primary** node of an HA MDC, type:

```
service cvfs stop
```

Caution: The command **service cvfs stop** will stop all StorNext services on the MDC.

- 3 Navigate to the temporary location where you stored the new **license.dat** file.
- 4 Copy the **license.dat** file to the **/usr/cvfs/config/** directory on the **primary** node of the MDC.

Note: Start the StorNext service on the **primary** node of the HA system **before** starting the StorNext service on the **secondary** node.

- 5 From the command line of the **primary** node of an HA MDC, type:

```
service cvfs start
```

- 6 From the command line of the **secondary** node of an HA MDC, type:

```
service cvfs start
```

Note: Wait up to 2 minutes for the **license.dat** file to propagate to the secondary node of the MDC. After 2 minutes, display the file and verify the license has been propagated on the secondary node.

Installing StorNext License Keys on a Failover System (without a Shared File System)

Note: Stop the StorNext service on the **secondary** node of the HA system **before** stopping the StorNext service on the **primary** node.

- 1 From the command line of the **secondary** node of an HA MDC, type:

```
service cvfs stop
```

Caution: The command **service cvfs stop** will stop all StorNext services on the MDC.

- 2 From the command line of the **primary** node of an HA MDC, type:

```
service cvfs stop
```

Caution: The command **service cvfs stop** will stop all StorNext services on the MDC.

- 3 Navigate to the temporary location where you stored the new **license.dat** file.

- 4 Copy the **license.dat** file to the **/usr/cvfs/config/** directory on the **primary** node of the MDC.
- 5 Copy the **license.dat** file to the **/usr/cvfs/config/** directory on the **secondary** node of the MDC.

Note: Start the StorNext service on the **primary** node of the HA system **before** starting the StorNext service on the **secondary** node.

- 6 From the command line of the **primary** node of an HA MDC, type:

```
service cvfs start
```

- 7 From the command line of the **secondary** node of an HA MDC, type:

```
service cvfs start
```

Installing Optional StorNext Add-on Software Features and Clients from the StorNext MDC GUI

Client software installers and Optional Add-on Software features can be downloaded from the StorNext GUI. Downloads are available on the **Tools > Client Download** screen.

Step 3: System (M660, M440 and M330 Metadata Appliances and Pro Foundation Only)

This step describes how to enter or view network configuration information for the primary and secondary nodes of the StorNext M660, M440 or M330 Metadata Appliance and Artico. The system date and time can also set during this procedure.

The system comes pre-configured for its environment during system installation. IP addresses for the primary and secondary nodes were entered, and the system was configured to operate as High Availability (HA) pair of nodes. Here is some general information about the System Network page.

Configuration Network Settings

General Settings

Follow these steps to view or enter network configuration for the secondary node of the HA system.

- 1 After reviewing the StorNext Licenses, click the **Next** button. (Alternatively, choose **System** from the **Configuration** menu on the left side of the screen.)

The **Configuration > System** screen displays pre-configured information about the primary and secondary nodes configured during the initial StorNext installation.

Figure 10 Configuration > System Screen

The screenshot displays the Quantum StorNext 6 M440 Configuration > System > Network page. The interface includes a navigation bar with 'Configuration', 'Tools', 'Service', 'Reports', and 'Help'. The main content area is divided into several sections:

- General Settings:** Fields for Primary Hostname (bstm440-1a), Secondary Hostname (bstm440-1b), Default Gateway (10.65.160.1), Domain Suffix Search List (mdh.quantum.com), Primary DNS IP Address (10.65.162.1), Secondary DNS IP Address (10.65.162.2), and Tertiary DNS IP Address.
- Bonding:** A table showing bonding configurations for eth1 (1GbE), eth2 (1GbE), and eth3 (1GbE). The 'Not Bonded' row shows 'bond0 (Metadata)' with checkboxes for each interface.
- Configuration:** A table for configuring network interfaces.

	Alias	Primary IP	Secondary IP	Netmask	Gateway	Jumbo Frame	Bonding Mode
bond0 (Metadata)	bond0 2	[192.168.44.100]	[192.168.44.101]	[255.255.255.0]	[192.168.44.1]	<input type="checkbox"/>	[Round Robin]
eth1 (SysMgmt)	eth1.1	[10.65.166.55]	[10.65.166.56]	[255.255.224.0]	[10.65.160.1]	<input type="checkbox"/>	
- Status:** A table showing the status of network interfaces.

	Link Speed	Ports	Primary Status			Secondary Status		
			Link	Carrier	Link	Carrier		
bond0 (Metadata)	1GbE	eth2 eth3	Up	Down	Up	Down	Down	
eth1 (SysMgmt)	1GbE	eth1	Up	Up	Up	Up	Up	

At the bottom, there are 'Apply' and 'Reset' buttons, and a note: '* Required Field'.

IP addresses are assigned to the Primary and Secondary Nodes during system installation.

- 2 Verify that the IP addresses are correct.
- 3 Enter Secondary and Tertiary DNS IP addresses if necessary.

Bonding Settings

- 4 Modify the bonding settings as needed. In this section you specify the bonding mode used by the network interface. Bonds are used by StorNext M660, M440 and M330 Metadata Appliances and Pro Foundation systems to increase the size of the data pipes available for network traffic. By default, Bond0 is used for metadata traffic, and Bond1 is used for system management traffic.

To break the bonds, click the radio button for “Not Bonded” underneath the appropriate ethernet port listed in the **Bonding** section of the screen.

Note: If you break the bonds for Bond 0 (M660, M440 and M330 and Pro Foundation) and Bond 1 (M660-only), the first ethernet port of the bond pair will allow network traffic, but the second ethernet port will be disabled. (Bond 0 bonds ethernet ports 2 and 3, and Bond 1 bonds ethernet ports 1 and 4.) For example, if you break Bond 0, port 2 will be used to handle metadata traffic, but port 3 will be disabled. Likewise, if you break Bond 1, port 1 will be used to handle system management traffic, but port 4 will be disabled. For this reason, and since bonding increases the amount of data for the bonded pair, it is recommended to leave the bond settings alone for Bond 0 and Bond 1.

Bonding rules:

- a You need two or more ports for a bond.
- b Each bond needs to be on a separate network.
For example:
bond0 - 10.60.224.10
bond1 - 10.60.225.10
bond2 - 10.60.226.10
- c You cannot bond ports with different ethernet speeds. For example, you cannot bond a 10GbE port with a 1GbE port.
- d There is a maximum number of $n/2$ bonds available. For example, if there are 9 ports available, there can be 9/2, or 4 bonds maximum available.

Configuration Settings

5 Enter the appropriate information into the following fields:

Note: An IP address, Netmask and Gateway is required for all ethernet ports which will be used in your environment.

- **The first column:** This lists the current network bond and ethernet ports available for configuration.
- **Alias:** For faster ethernet ports (such as the 10GbE port) the alias can segregate network traffic with different subnets by creating additional virtual IP addresses for that port on the system. Click the plus sign (+) next to the ethernet port for which you wish to add an alias. By default, the initial ethernet port is assigned the alias of 1. For example, ethernet port 5 is defined as "eth5:1". Its first alias would be defined as "eth5:2". You must define an IP address for both the Primary and Secondary, and a Netmask and Gateway for each alias.

Note: Up to 10 aliases may be defined for every ethernet port.

Note: For systems that have had VLAN tagging manually configured, the following apply:

- An upgrade should not impact any configured VLAN devices in the system.
- VLAN devices appear as network devices in the GUI. For example, a VLAN device appears as bond2.345 where 345 is the assigned VLAN tag.
- You can view the settings in the GUI, but should avoid changing any VLAN device settings using the GUI.

-
- **Primary IP:** This is the IP address of the Primary Node.
 - **Secondary IP:** This is the IP address of the Secondary Node.
 - **Netmask:** This is the netmask for this Bond or Ethernet connection.
 - **Gateway:** This is the Gateway used for this Bond or Ethernet connection.

- **External Host:** Select an interface from the Configuration Table on the **Configuration > System > Network** page by clicking the **External Host** radio-button for that interface. The selected interface IP address will be added to the file /etc/hosts as the external host IP address of the host (the IP address will be used for the FQDN entry for the host in the /etc/hosts file). When the **Configuration > System > Network** page is loaded, the **External Host** radio-button is set to the current external host IP.

Note: You can only select an external host whose IP address' network matches that of the default gateway IP address' network.

- **Jumbo Frames:** Use this option for networks which support this type of ethernet payload.

Note: **DO NOT** enable the Jumbo Frames option on the MDC node(s) without first ensuring that the network is configured to use Jumbo Frame payloads. You will typically get better total throughput with Jumbo Frames enabled. However, this option is dependent on the existing network infrastructure of StorNext network. Enabling Jumbo Frames within a network would require reconfiguring the network switches upstream from the MDC and all Ethernet ports connected to this network would need to be capable of supporting Jumbo Frames.

- **Bonding Mode:** The current bonding mode is displayed in this column. There are two modes supported in the M660, M440 and M330 and Pro Foundation - Round Robin, and LACP.
 - **Round Robin:** In this mode the Ethernet frames are sent via the bonded Ethernet ports with a valid MII link in a round-robin fashion starting with the first slave device and then the rest of the devices. This applies only to the traffic sent from the M660, M440 and M330 or Pro Foundation. Your Ethernet switch needs to aggregate the ports, so the connected ports are treated as a logical port. The frame reception is completely dependent on your Ethernet switches' transmission algorithm. The bonding mechanism does not balance the frame reception.
 - **LACP:** LACP stands for "Link Aggregation Control Protocol." This mode is based on the 802.3ad IEEE standard for aggregating the Ethernet ports. (There is IP and MAC based

LACP.) If the bonding algorithm is set to LACP, your Ethernet switch ports must be configured in a 802.3ad-based Link Aggregation Group (LAG), in LACP mode. The frame reception and transmission is controlled by the LACP between the bonded ports and your Ethernet switch ports.

Note: In order to maintain network connectivity to your system, you must configure the switch that is connected to your system to use the same bonding mode. The best time for you to change the bonding mode on your switch will be during the next reboot of your system, after you have saved the new network settings. Changing the bonding mode on your switch before saving these settings and rebooting may result in the loss of network connectivity to your system.

Routing

The **Routing** section allows you to manage static routes for a given alias. Enter the appropriate information into the following fields:

Parameter	Description	Notes
The first column	Defines the interface for which a static route is defined. The packets for the destination of the static route are sent through this interface.	

Parameter	Description	Notes
Destination IP	Defines the destination IP address of the route. Enter a network address (for routing packets to a network), or a host IP address.	The destination defined with the combination of Destination IP address and Destination Netmask may be a network address or a host address.
Destination Netmask	Enter the netmask address to use as the netmask for the Destination IP address.	<p>The combination will be a host address if the Destination IP address is a host IP address and the Destination Netmask is 255.255.255.255, and the defined route will be a host route.</p> <p>The combination will be a network address if the Destination IP address is a host or network IP address and the Destination Netmask is not 255.255.255.255, and the defined route will be a network route.</p>
Destination Gateway	Enter the IP address of the gateway to send the packets to. The IP address reflects the destination defined with the combination of the Destination IP and the Destination Netmask .	

Status

- 6 Review your network configuration settings and make changes as necessary.

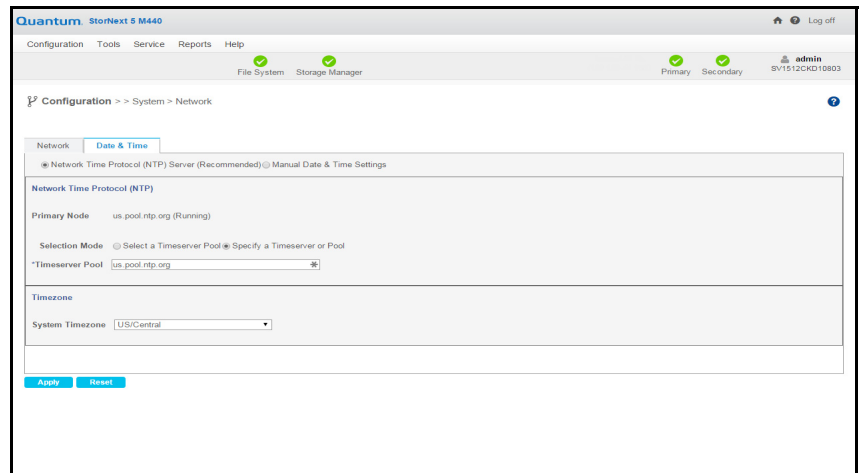
Note: If you click **Apply** after changing network settings, you will be required to reboot your system, which could take 30 minutes or longer to complete. If you plan on updating the date & time for the system, Quantum recommends you finish configuring those settings, described in the next section, prior to clicking **Apply**.

Entering Date and Time Information

Follow these steps to set the system date and time.

- 1 From the **Configuration > System** screen, click the **Date & Time** tab. The **Configuration > System > Date & Time** screen appears.

Figure 11 Configuration > System > Date & Time Screen



- 2 Choose one of the following:
 - **Network Time Protocol (NTP) Server (Recommended)**
 - **Manual Date & Time Settings**

This setting determines whether the date and time are automatically retrieved from a specified server, or entered manually.

Setting the Date and Time via NTP

1 After you choose **Network Time Protocol (NTP) Server (Recommended)**, choose one of the following:

- **Select a Timeserver Pool**
- **Specify a Timeserver or Pool**

Note: Quantum strongly recommends using an NTP server to keep the time synchronized between both nodes of the HA pair and the metadata on the shared file system, alleviating data mismatches due to incorrect time stamps. If date and time between the two StorNext MDC nodes becomes inconsistent, it will also cause HA to become unreliable.

Consistent time synchronization is also critical for StorNext Pro Solutions, and when using StorNext with Quantum QCloud, and Quantum Lattus S3 products. An out-of-sync system will produce errors similar to “The difference between the request time and the current time is too large”. In order to avoid this situation, it is critical that all the StorNext MDC nodes and other components that capture time stamps and operate in the same network are synchronized to the same time.

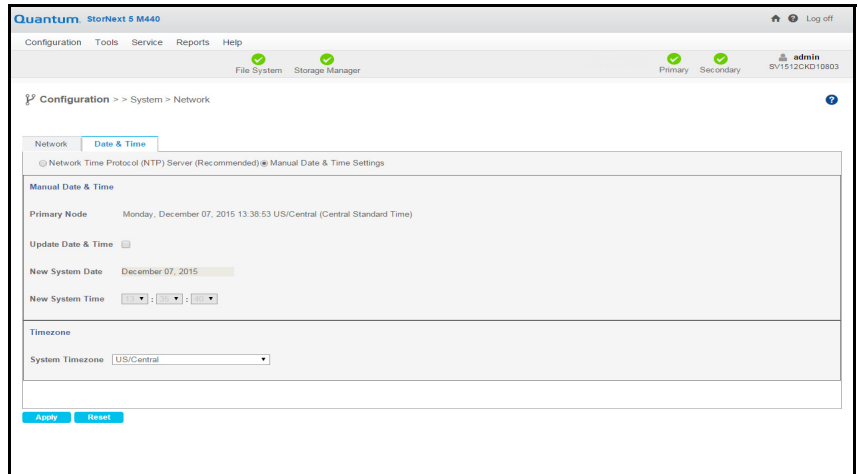
If an Internet connection is available, Quantum strongly recommends you use a public NTP server pool. If there is no outside connection available, set up an internal NTP server on another node and use that server as the NTP server used for the StorNext MDC nodes. That way the time and date stamps will be consistent for data.

- 2 If you chose **Select a Timeserver Pool**, select the desired server from the dropdown list next to **Timeserver or Pool**.
- 3 If you chose **Specify a Timeserver Pool**, enter the IP address or URL of the time server or pool you want to use.
- 4 At the **System Timezone** field, select the area corresponding the location of the M660, M440 and M330 and Pro Foundation.

Setting the Date and Time Manually

- 1 After you choose **Manual Date & Time Settings**, the following screen appears:

Figure 12 Configuration > System > Date & Time Screen (Manual Entry)



- 2 Enter the following fields:

- **Update Date & Time**
- **New System Date**
- **New System Time**

For detailed information about what to enter in these fields, see the StorNext online help.

- 3 At the **System Timezone** field, select the area corresponding to the location of the M660, M440 and M330 or Pro Foundation.

Applying Your Changes

Clicking **Apply** on any tab applies the settings you specified on all three tabs, so there is no need to click Apply on all three tabs. You can simply click Apply once when you are finished.

When you are satisfied with your network configuration information, click **Apply** to save and apply the information. (To clear all fields of information, click **Reset**.)

Note: A system reboot could take 30 minutes or longer, in order for all system services to function correctly. So, after the network configuration has been saved, please close your browser window and wait for at least 30 minutes before attempting to log back in. If the IP address you used to log in to this system will be changing, you will temporarily lose your connection to the Web-based interface. Because of this, you may not receive a confirmation page, letting you know that your settings have been saved.

After you click **Apply**, a message warns you that both nodes must be rebooted for your changes to become active. When you are ready to proceed, click **Yes**.

Step 4: 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).

Note: On Linux systems the NetworkManager service should be turned off because it can interfere with the StorNext nameserver and Linux network devices.

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).

Note: The `fsnameservers` file must be the same for all StorNext clients.

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 anMDC 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 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 `fsnameservers` 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:

- The `fsnameservers` file must be the same for all StorNext MDCs.
- Metadata controllers needn't be name servers.
- Each additional `fsnameservers` entry adds additional heartbeats from every file system host.
- If multiple metadata networks service an individual file system, each network must have an `fsnameservers` interface. Each `fsnameservers` host must have network interface(s) on every metadata network, and each interface must be listed in the `fsnameservers` file.
- At maximum heartbeat rate, a host sends a heartbeat message to every `fsnameservers` 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 fsnameservers 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 fsnameservers addresses would produce an aggregate of about $500*4$ or 2000 Mbits or 2Gbits of heartbeat reply messages per second. If the 4 fsnameservers 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.

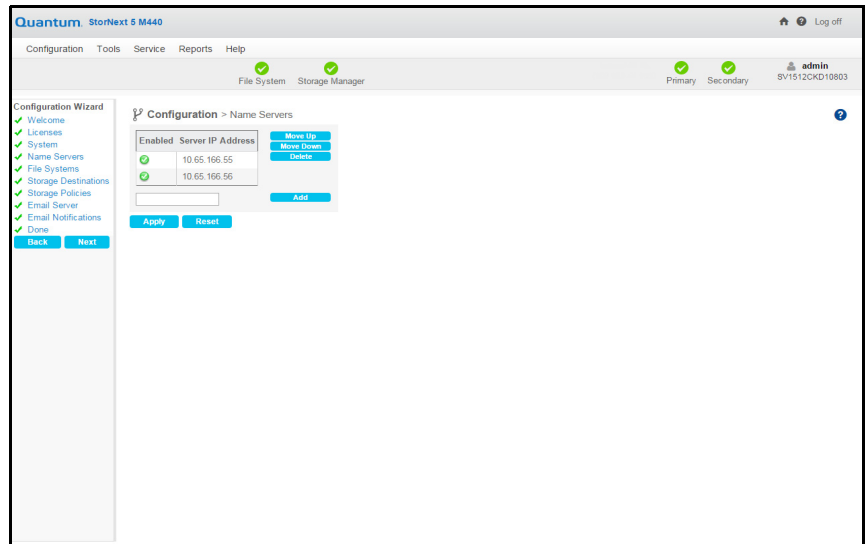
- 7 The metadata network carries more than just name server traffic. All metadata operations such as open, allocate space, and so on use the metadata network. File system data is often carried on the metadata network when LAN clients and servers are configured. Network capacity must include all uses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

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 13 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 server appears in the list of Server IP addresses.
- 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.

Note: If the file system fails to mount and returns a “device not connected” error after changing `fsnameservers` in the StorNext GUI, then stop and start the StorNext services manually from the command line.

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**.

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, which is created in the `cvfs` config directory. (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 LAN client mount option.

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

Figure 14 StorNext Foreign Servers



Step 5: File Systems

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 91.

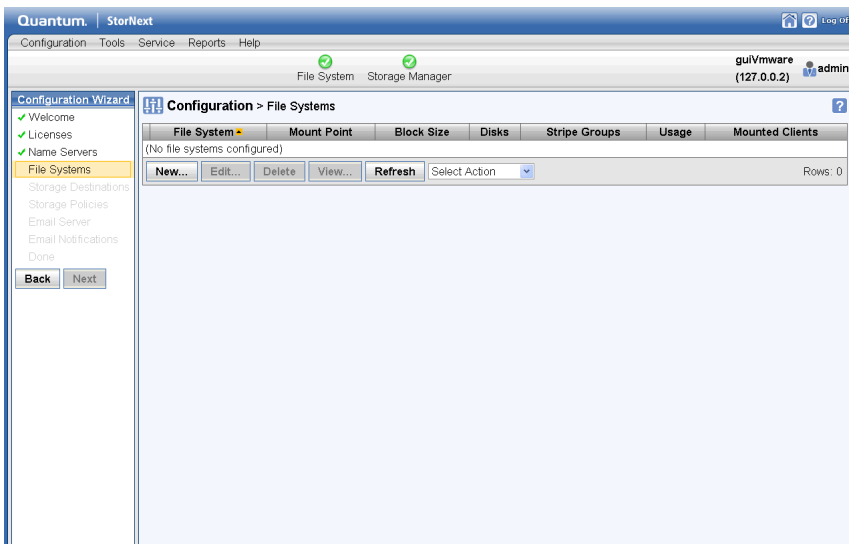
Note: The maximum supported file system size is 1 Billion files per File System.

Note: For managed file systems only, the maximum recommended directory capacity is 50,000 files per single directory. This recommendation does not apply to un-managed file systems.

- 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.)

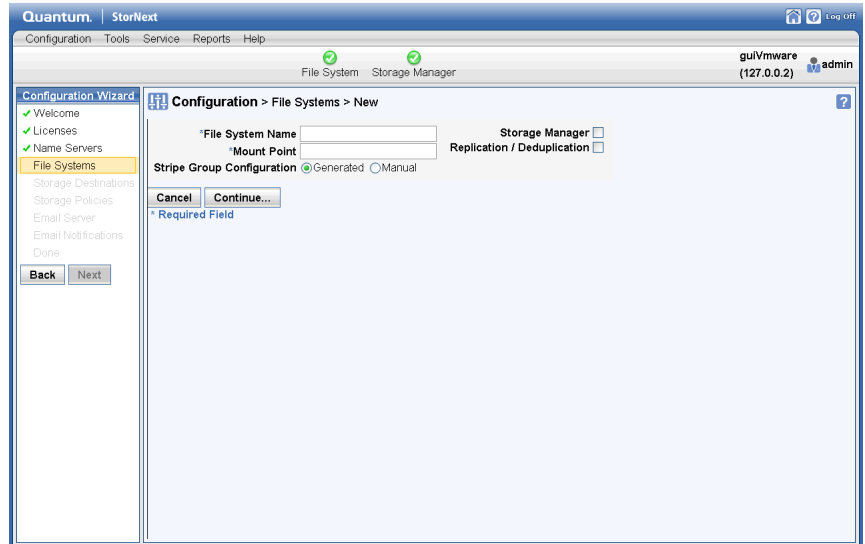
From this screen you can view, add, edit, or delete a file system. For information on these procedures, see the online help.

Figure 15 Configuration > File System Screen



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

Figure 16 Configuration > File System > New Screen



3 Enter the following fields:

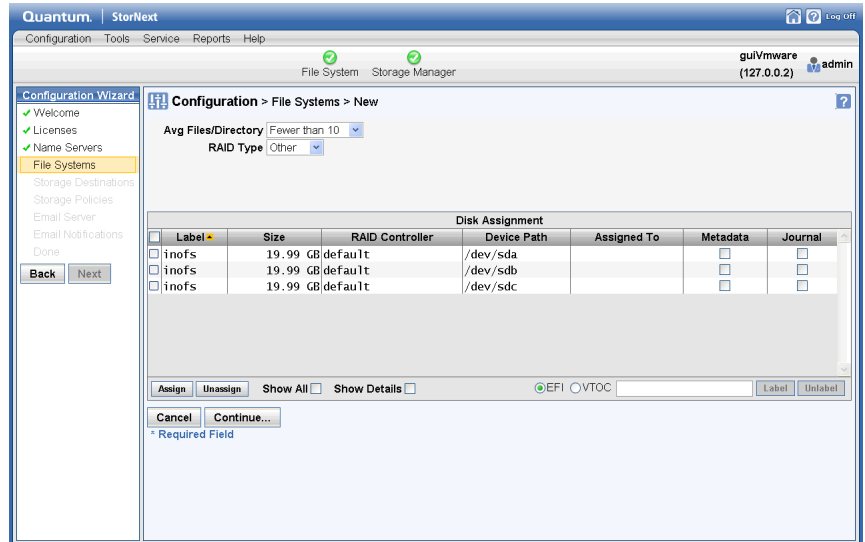
- **File System Name**
- **Mount Point**
- **Storage Manager**
- **Replication/Deduplication**
- **Stripe Group Configuration: Generated or Manual**

(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 17 Configuration > File System > New Screen 2



- 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.

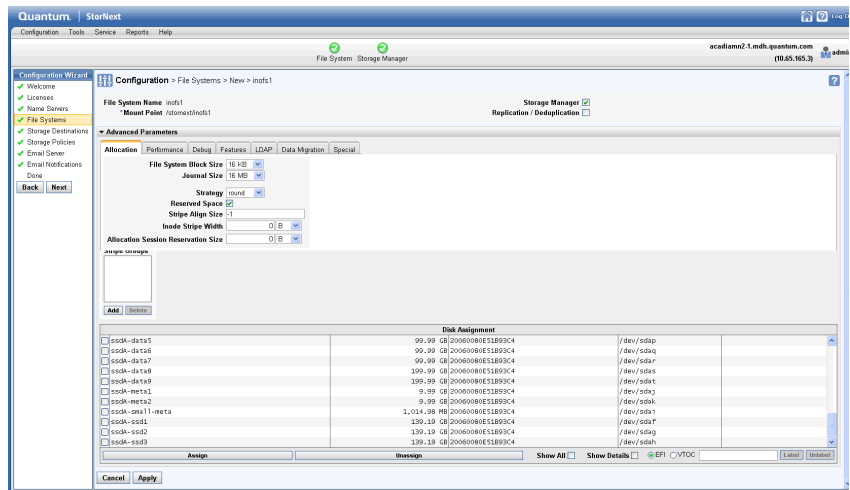
This value helps StorNext calculate capacity. You should select the option that most closely corresponds to your projected usage. You are not limited by the option you select. For example, if you indicate that there will typically be fewer than 10 files per directory, there is nothing to prevent the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

- At the **RAID Type** field, select the type of RAID the file system uses. (If you are using a StorNext G300 Appliance, the default value is **Quantum Disk**.)
- Select one or more disks to assign to the file system.
- 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.
- Specify the label type by choosing **EFI** or **VTOC**. For more information on label types, see [Label Disks](#) on page 92.
- Enter a disk label name in the text field to the right of the **EFI** and **VTOC** buttons. 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.

- 11 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.)
- 12 Click **Continue**.

Figure 18 Configuration > File System > New Screen 3



- 13 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.
- 14 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

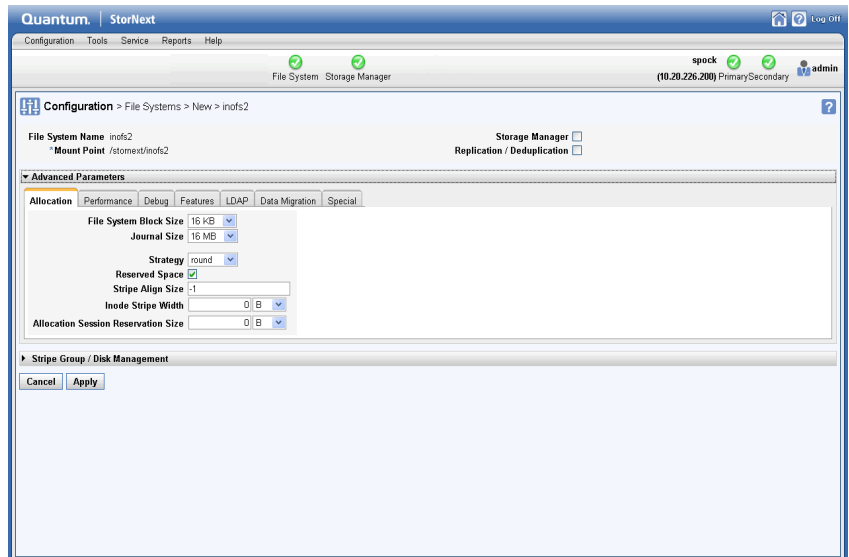
If you chose Manual Configuration, you must enter the fields on the **Advanced Parameters** tabs and the **Stripe Group/Disk Management** fields. (Click the arrow to the left of these headings to display the tabs and fields.)

File Systems and Stripe Groups Limitations

The following are limitations regarding file systems and stripe groups:

- 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.

Figure 19 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 / Start: Stop or start the file system.	Update: Applies configuration changes to the file system. Note: This option may trigger a temporary stall as the file system is updated. If a managed file system is modified a new metadata dump may be generated.
Start and Activate: Starts and activates the file system in one step, saving you the time of starting and activating separately.	Metadump: This option will generate a new metadata dump. Metadata dumps are stored during StorNext backups to provide a means of recovering the file system metadata in the event of a failure of the metadata storage. Generating new metadata dumps are necessary in the event that StorNext backups cannot succeed due to an invalid metadata dump. Generating a metadata dump will stop the file system temporarily. Note: A full backup should be scheduled any time new Metadumps have been generated to ensure that all metadata is protected.
Start and Mount: Starts and mounts the file system in one step, saving you the time of starting and mounting separately.	Check: Initiates a check of the file system. You should perform a check if you plan to expand and migrate the file system. This operation could take a significant amount of time depending on the size of the file system, so plan accordingly.

<p>Activate: Activate the file system.</p>	<p>Expand: Expands the file system in preparation for migration. This option can be used to add stripe groups to a file system. Ensure that all needed disks are visible before starting this process.</p> <p>Note: This option may trigger a temporary stall as the file system is updated. If a managed file system is modified a new metadata dump may be generated.</p> <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>
<p>Mount / Unmount: Mount or unmount the file system.</p>	<p>Migrate: Migrates the file system after expansion. This option is used to migrate data or metadata off of existing storage. If migrating metadata the metadata will be migrated without changing the metadata layout. If a data stripe group is being migrated, data will be moved from a selected stripe group to other stripe groups.</p> <p>Note: The GUI is not able to down a stripe group when LUNs are unavailable. Mark stripe groups down in the GUI before taking the stripe group's disks offline. If that is not possible, set the stripe group down directly through the FSM configuration file and restart the FSM. See the snfs_config man page or the <i>MAN Pages Reference Guide</i> for details.</p>
<p>Make: Make the file system.</p>	

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

Step 6: Storage Destinations

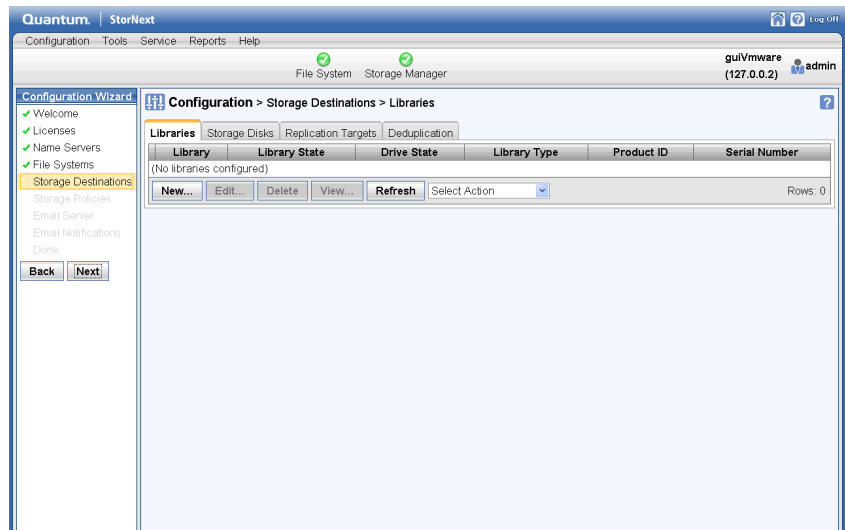
After you have created at least one file system, the Configuration menu's Storage Destinations option allows you to add, edit, or delete libraries and storage disks. You can also enter or edit targets for data replication, and specify a blockpool host file system for data deduplication. Additionally, you can add Wide Area Storage Destinations (see [Setting Up Lattus Object Storage Destinations](#) on page 424).

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 20 Storage Destinations
> Library Screen



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

Figure 21 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.

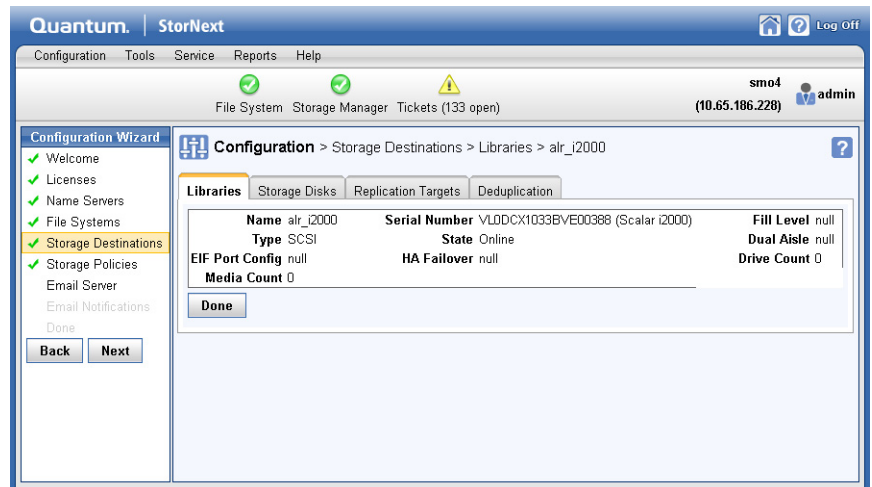
Viewing an Existing Library

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 20](#).)

- 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 22 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

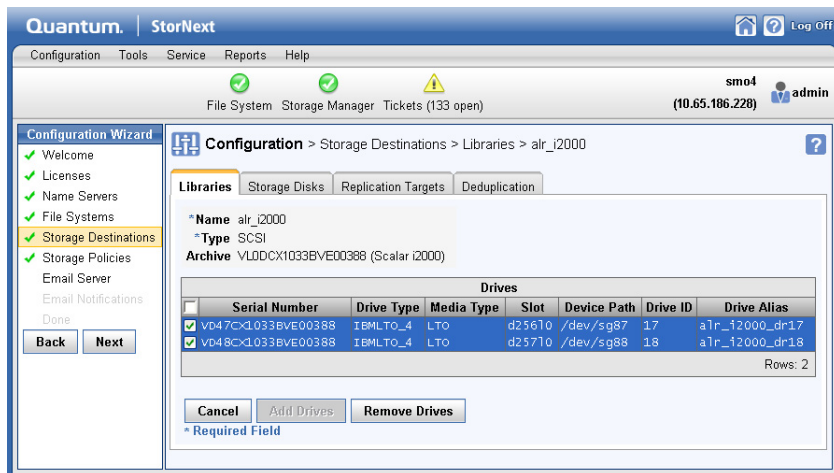
- 5 When you are finished viewing library information, click **Done**.

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 **Libraries** tab.
- 2 Select the library you want to edit.
- 3 Click **Edit**. After you select this option StorNext scans all SCSI devices connected to the library, which could take some time to complete depending on your configuration.

Figure 23 Edit Library Screen



Drives shown with a blue background and a green check in the adjacent checkbox are currently included in the library. Drives shown with a white background are currently not included in the library.

- 4 Select the checkbox next to the drive(s) you want to add or remove. Click the checkbox to the left of the Serial Number heading to select or deselect all available drives.
- 5 Click **Add Drives** to add the selected drives to the library, or **Remove Drives** to remove the drives from the library.
- 6 When the confirmation message appears, click **Yes** to proceed, or **No** to abort.

- 7 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 68.)

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.

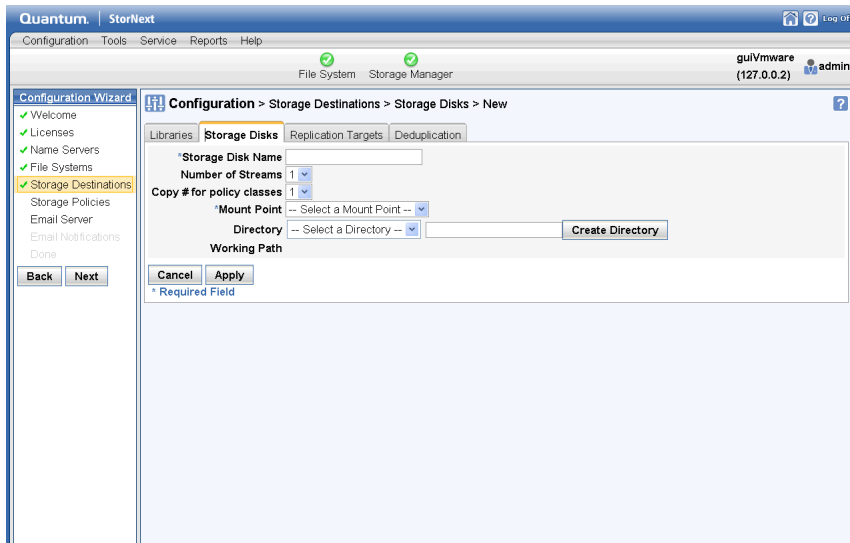
Information for any previously configured storage disks is shown. For each configured storage disk, the screen shows the current state and status (available or unavailable), storage disk name, mount point, directory, maximum number of I/O streams that can concurrently write to the disk, the copy number assigned to the storage disk, usage, and file count.

Figure 24 Configuration > Storage Destinations > Storage Disk Screen



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

Figure 25 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 24](#) on page 68.)
- 3 Select the storage disk whose information you want to view.

4 Click **View**.

Figure 26 View Storage Disk Screen



5 When you are finished viewing Storage Disk 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 68.)
- 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.

Setting up Wide Area Storage Destinations

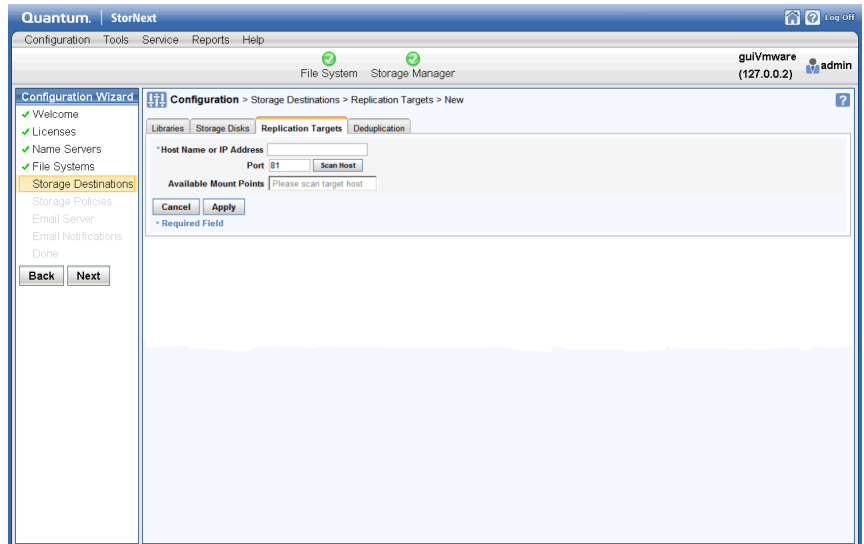
See [Setting Up Lattus Object Storage Destinations](#) on page 424.

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 Targets** screen appears. Click the **New...** button.

Figure 27 Configuration >
Storage Destinations >
Replication Targets / New
Screen



- 2 At the **Host Name or IP Address** field, input a valid host name or IP address for the replication target.
- 3 At the **Port** field, use the default port (81), or input a valid port.
- 4 Click **Scan Host** to identify available mount points on the selected host.
- 5 At the **Available Mount Points** field, Click **Add** to add the new replication target, or **Cancel** to abort without saving.
- 6 Click **Apply** to save your changes.

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 28 Configuration > Storage Destinations > Replication / Deduplication Screen (Blockpool)



Step 7: Storage Policies

There are two kinds of storage policies: Storage Manager storage polices 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:

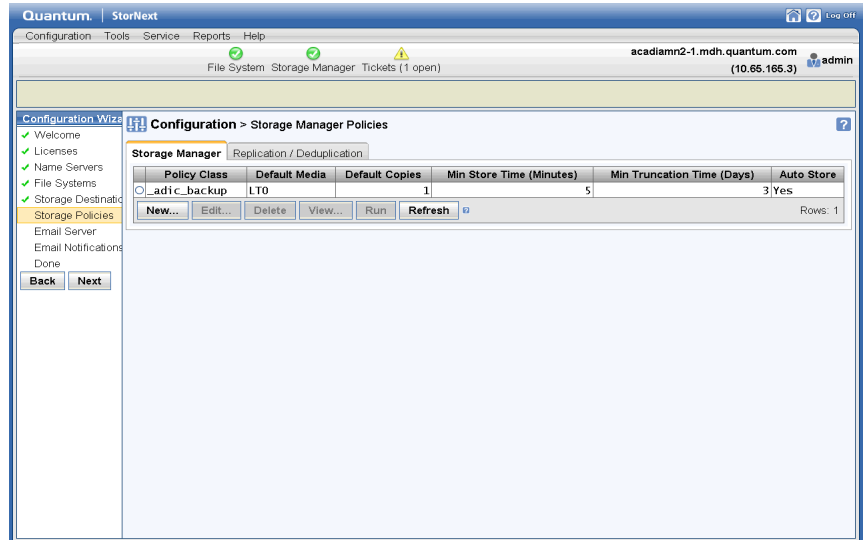
- A directory in which to store backups every night is created. This directory is seldom accessed after the files are copied over. A storage policy could be set up to create two tape copies of the files, store one copy of the files to LTO media after residing on disk for 10 minutes, and then truncate 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:

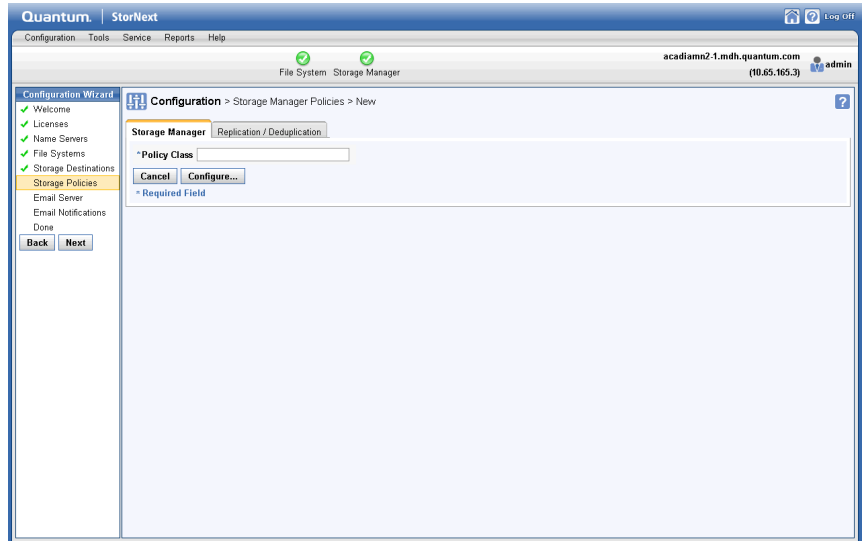
- 1 When the Configuration Wizard is displayed, choose **Storage Policies** on the left side of the screen. (Alternatively, choose **Storage Policies** from the **Configuration** menu.) The **Configuration > Storage Policies** Screen appears.

Figure 29 Configuration > Storage Policies Screen



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

Figure 30 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 31 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 32 Storage Policies >
New > Relocation Tab



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

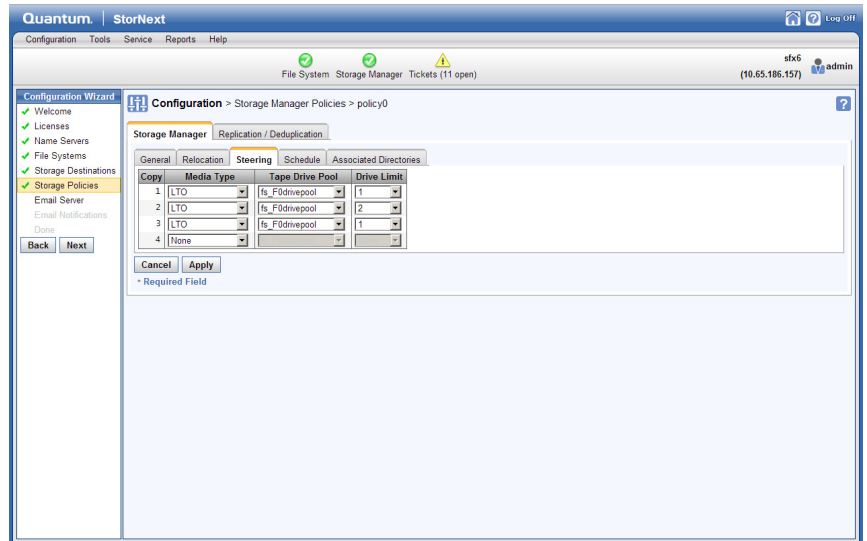
The Steering Tab

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.

Enhanced Control of Tape Drive Allocation

You can manage the number of tape drives to use per store policy by using the **Drive Limit** feature on the Steering Tab (see [Figure 33](#)). To limit the number of tape drives used per policy and copy, new configuration options were added to the policy class commands **fsaddclass**, **fsclassinfo**, and **fsmodclass**. The Storage Manager will use these values to manage the number of drives used per store policy. See the *StorNext Online Help* for additional information and usage. See the *Man Pages Reference Guide* for details on the policy class commands.

Figure 33 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 155.)

Figure 34 Storage Policies >
New > Schedule Tab

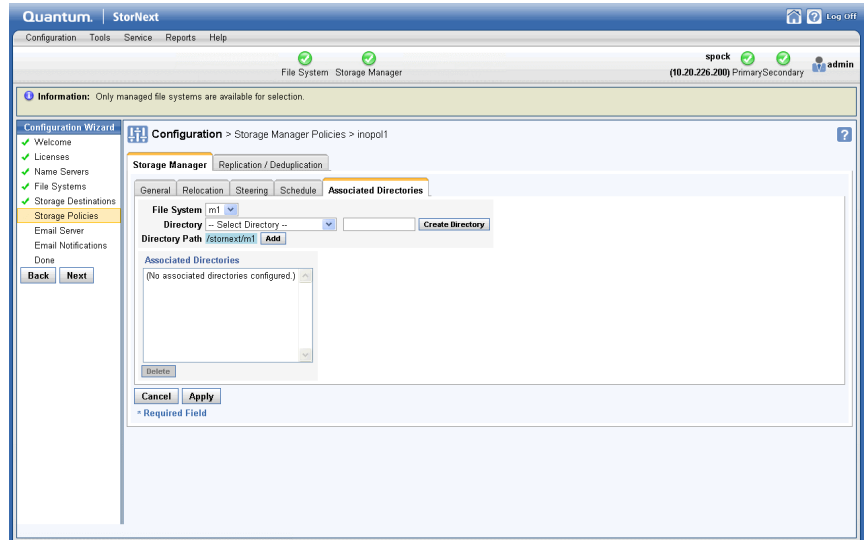


For instructions on what to enter on this screen, see the online help.

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 35 Storage Policies >
New > Associated Directories
Tab



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

Adding a Replication Storage Policy

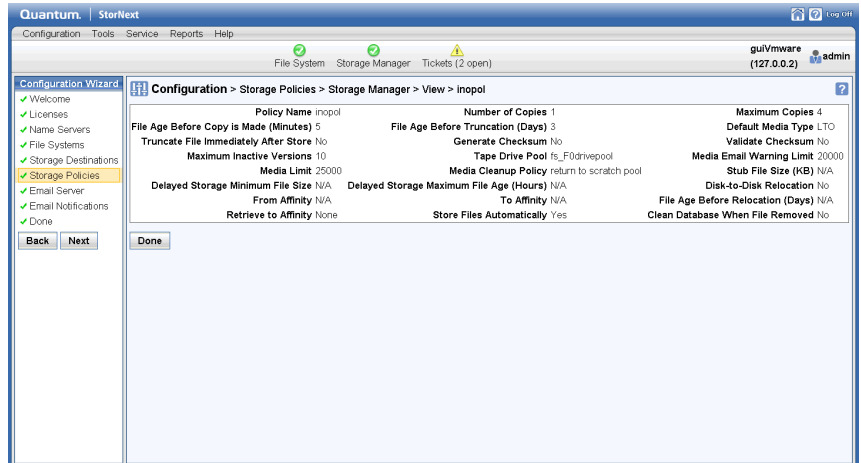
The steps for creating a replication storage policy are described in [Step 4: Create a Replication Storage Policy](#) on page 207.

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 36 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, Outbound Replication, Inbound Replication, Source Directories** and **Blackout** tabs.

If you are editing a Replication target policy, you can edit fields on **Deduplication, Outbound Replication, Inbound Replication,** and **Source Directories** tabs.

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 8: 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.

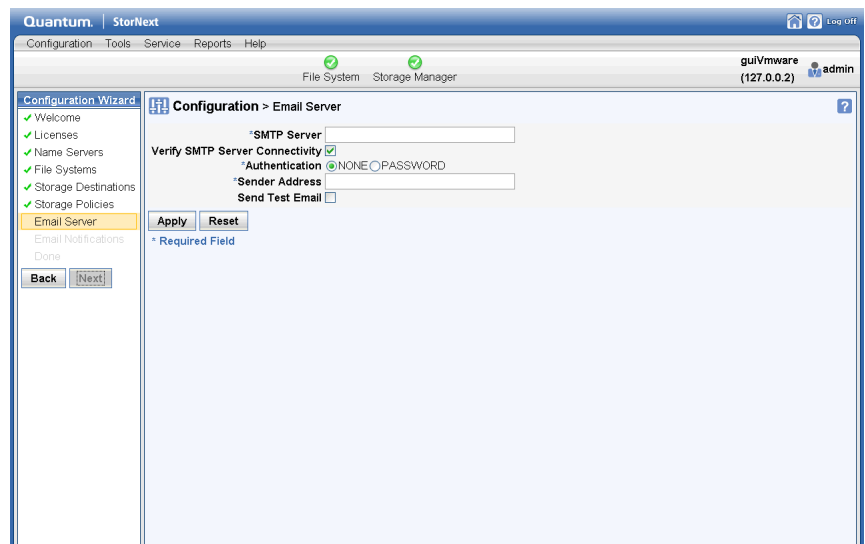
Note: The Email Server option does not configure your email server. Instead, it allows you to specify a previously configured email server so StorNext knows which server is responsible for processing notification messages. Before you use the Email Server option, make sure your email SMTP server is already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

Adding an Email Server

Follow this procedure to add a new email server.

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

Figure 37 Configuration > Email Server Screen



The screenshot shows the Quantum StorNext Configuration Wizard interface. The left sidebar lists various configuration steps, with 'Email Server' highlighted. The main content area is titled 'Configuration > Email Server' and contains the following fields and controls:

- SMTP Server:** A text input field.
- Verify SMTP Server Connectivity:** A checked checkbox.
- Authentication:** Radio buttons for 'NONE' (selected) and 'PASSWORD'.
- Sender Address:** A text input field.
- Send Test Email:** An unchecked checkbox.
- Buttons:** 'Apply' and 'Reset' buttons are located below the fields.
- Legend:** A note indicates that an asterisk (*) denotes a required field.

- 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 9: 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 85.

Adding an Email Recipient

Follow this procedure to add a new email recipient.

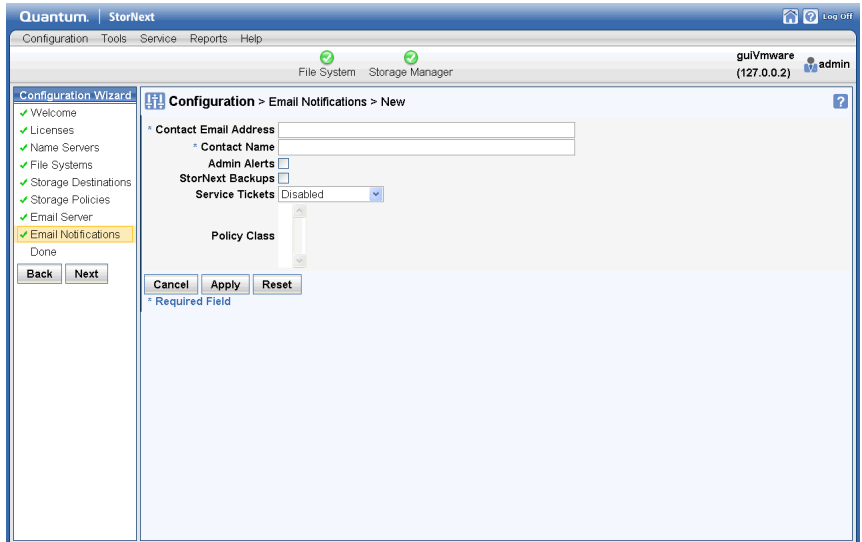
- 1 Choose **Email Notifications** from the **Configuration** menu.
- 2 When the Configuration Wizard is displayed, choose **Email Notifications** on the left side of the screen. (Alternatively, choose **Email Notifications** from the **Configuration** menu.) The **Configuration > Email Notifications**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

Figure 38 Configuration > Email Notifications Screen



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

Figure 39 Configuration > Email Notifications New Screen



- 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.

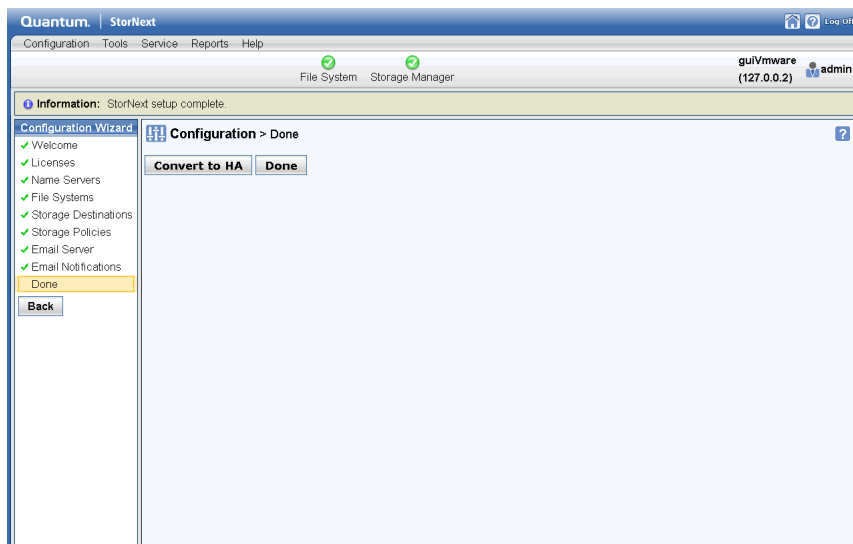
- 1 If you have not already done so, when the Configuration Wizard is displayed, choose **Email Notifications** on the left side of the screen. (Alternatively, choose **Email Notifications** from the **Configuration** menu.)
- 2 On the **Configuration > Email Notifications** screen, review the list of 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 10: 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 351.

Figure 40 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 5: File Systems](#) on page 54, 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
- [Manage Quotas](#): Limit the amount of disk storage consumed on a per user, or per group basis across an entire file system, or within a designated directory hierarchy.

To rename a standalone (unmanaged) StorNext File System, see [Renaming a Standalone \(unmanaged\) StorNext File System](#) on page 122.

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 labels:

- EFI labels are required if you plan to create LUNs that are larger than 2TB. (For Solaris, EFI labels are also required for LUNs with a raw capacity greater than 1TB.)
- 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.

Note: Do not use file system labels of the format `meta_`*any-value* and `shared_`*any-value*. These file system labels are reserved for the HA shared files system on the M-series Metadata Appliances.

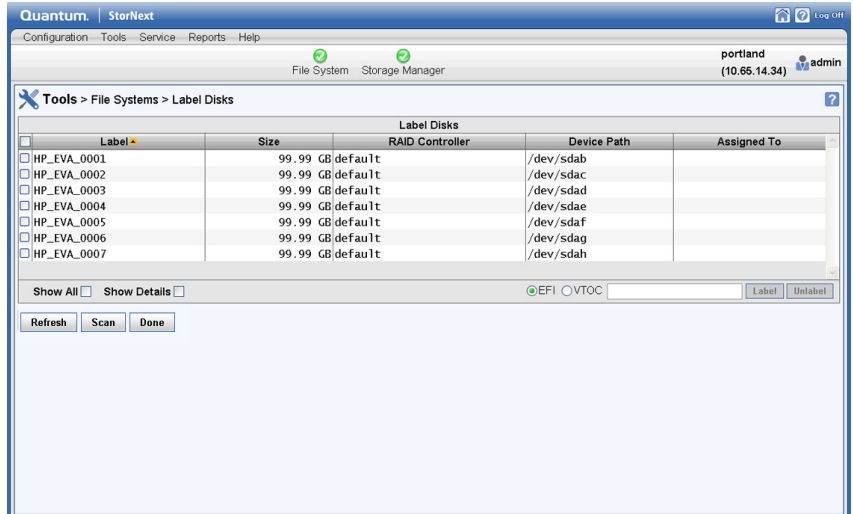
Labeling a Device

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 41 Label Disks Screen



2 If desired, click **Scan** to initiate a scan of the disk devices in your SAN.

Caution: Before you initiate a scan, be aware that in complex SAN environments frequent disk scanning can lead to SAN instability including timeout errors. Before the scan begins you will receive a reminder and be given the opportunity to confirm whether you want to proceed with the scan.

3 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.

4 Specify the label type by choosing **EFI** or **VTOC**.

5 Enter a label name in the text field to the right of the **EFI** and **VTOC** buttons.

6 Click **Label**.

- 7 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 101.)
- 2 Expand the file system. (See [Step 5: File Systems](#) on page 54.)
- 3 Move data stripe groups or metadata/journal stripe groups. (See [Migrate Data](#) on page 109.)
- 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 `sfsdefrag` 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 Unmount all clients to prevent applications that are writing to preallocated files from trying to do IO to the now read-only stripe group.
- 5 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.)
- 6 Stop the File System Manager (FSM) for the desired file system.
- 7 Run `cvupdatefs`.
- 8 Restart the FSM.
- 9 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.
- 10 Remount all clients since all pre-allocated blocks have now been moved to the stripe group.
- 11 Verify that no data remains on the original stripe groups.
- 12 Edit the file system configuration to mark the old stripe groups as "Disabled."
- 13 Stop the FSM.
- 14 Restart the FSM.

Note: The old stripe groups marked "Disabled/ReadOnly" must be left in the file system configuration file.

Moving a Metadata/ Journal Stripe Group

Metadata movement is performed on a LUN level, meaning you must specify the source LUN and the destination LUN. The new `sndiskmove` command that accomplishes metadata movement has two arguments: a source and destination LUN.

After movement is complete, the physical source disk can be removed.

Note: Although a stripe group can consist of multiple disks or LUNs, the `sndiskmove` command moves only a single disk or LUN. Consequently, references to “stripe group” in this section refer to a single disk or LUN when migrating metadata with `sndiskmove`.

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

`Sndiskmove` 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 42 Check File System Screen



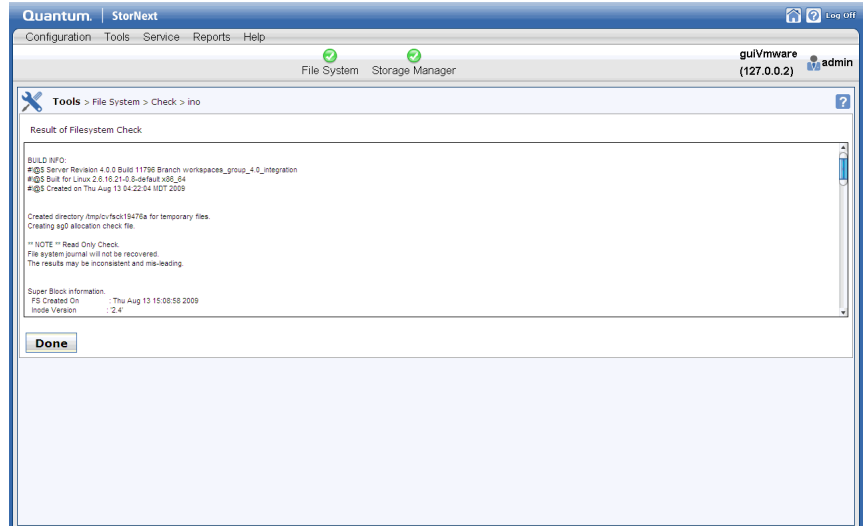
- 2 At the **Temp File Directory** field, enter a new directory if the specified directory does not have enough space to perform the check. (The checking process on large file systems can take hundreds of megabytes or more of local system disk space for working files.)
- 3 Select the file system you 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 43 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. Each stripe group can have zero, one, or more affinities, and a file or directory can have zero or one affinities associated with it. The default behavior is for stripe groups and files to have no affinities.

Exclusivity means a stripe group has one or more affinities and the `exclusive` attribute set to `true`, and can have its space allocated only by files with one of the associated affinities. Files without a matching affinity or with no affinity cannot allocate space from an exclusive stripe group.

Files with an affinity, exclusive or not, cannot be stored on stripe groups without that affinity. If all the stripe groups for an affinity become filled, no more files with said affinity can be stored, even if there are stripe groups with no affinity at all. This is independent of exclusivity.

Files with no affinity can be stored on stripe groups with affinities and available space without the `exclusive=true` attribute or on stripe groups with no affinities at all.

Turning on exclusivity can cause allocation failures for files with no affinity when there is space left on a stripe group. It does not affect allocation failures for files with an affinity, except indirectly by keeping the non-affinity files out of the way and thereby reserving the space just for affinity allocations.

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.

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 `sfnfsdefrag`, 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.

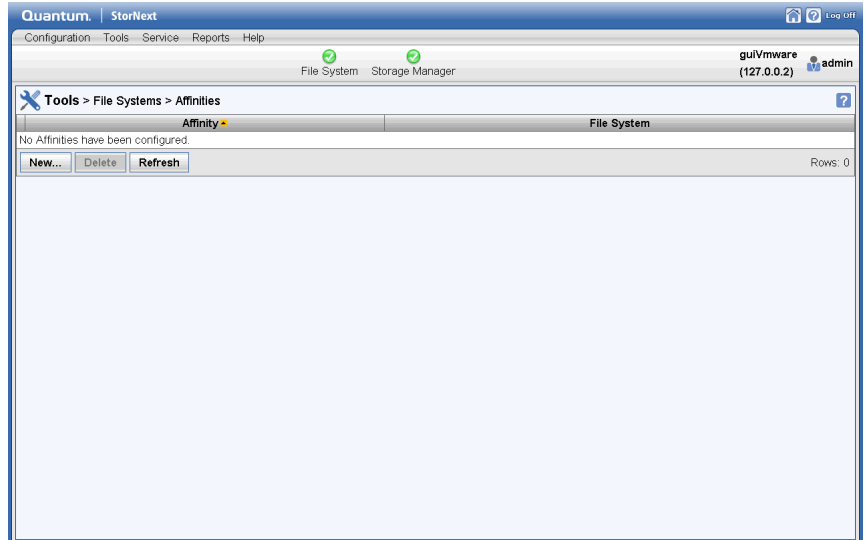
This will not work if critical files may be created later, unless there is a process to move them to the FAST stripe group, or an affinity is set on the critical files by inheritance but removed after their first allocation (to allow them to grow onto non-FAST groups).

Adding a New Affinity

Follow this procedure to add affinities:

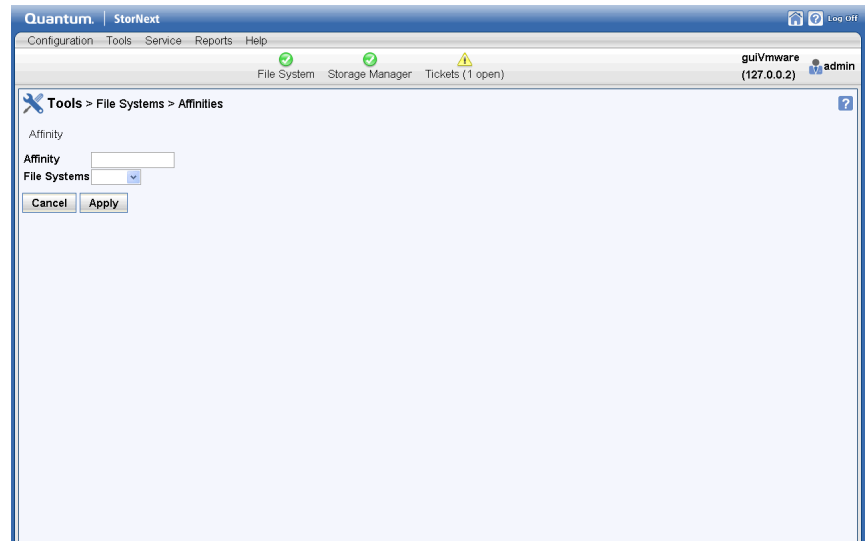
- 1 Choose **Affinities** from the **Tools > File Systems** menu. The **Tools > File Systems > Affinities** screen appears.

Figure 44 Affinities Screen



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

Figure 45 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 other stripe groups on the same file system, and then freeing the source stripe group so it can be removed from the file system.

To access the **Migrate Data** page on the GUI, on the **Tools** menu, click **File Systems**, and then click **Migrate Data** (refer to [Figure 46](#)).

Figure 46 Migrate Data page



You can migrate **User Data** stripe groups or **Metadata/Journal Data** stripe groups.

- To migrate **Metadata/Journal Data** stripe groups, see [How to Migrate Metadata and Journal Data](#) on page 111.
- To migrate **User Data** stripe groups, see [How to Migrate User Data](#) on page 113.

When migrating **User Data** stripe groups, select the source stripe groups only, not the destination stripe groups. 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.

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**.

When migrating **Metadata/Journal Data** stripe groups, select both source stripe group and destination stripe group.

Note: On large disks, the migration process may take some time since all of the blocks on the disks are copied, regardless of the amount of data on the disks.

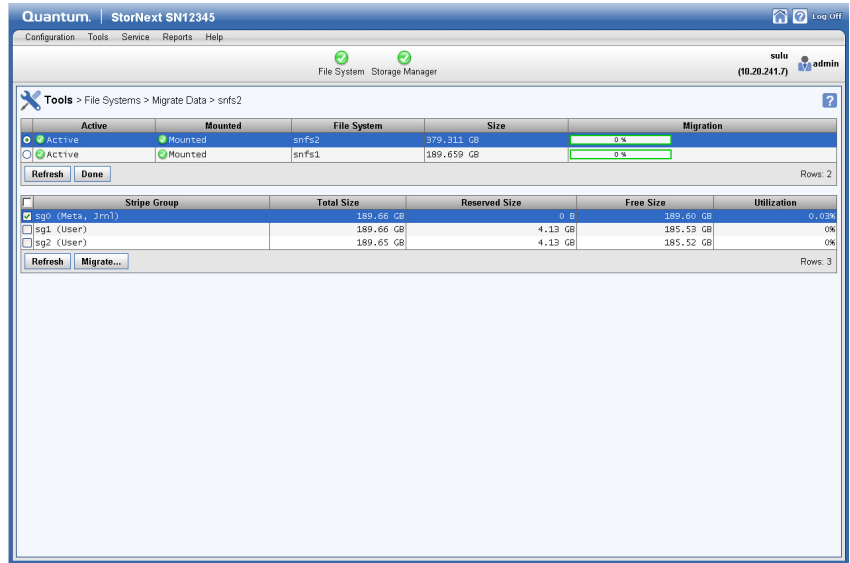
Note: When migrating a metadata stripe group to new disks, the overall metadata size will remain the same despite the size of the target disks. In order to add more metadata capacity, add new metadata stripe groups (refer to [Adding a Stripe Group Without Moving](#) on page 97).

How to Migrate Metadata and Journal Data

Use the following procedure to migrate **Metadata** and **Journal Data** using the GUI.

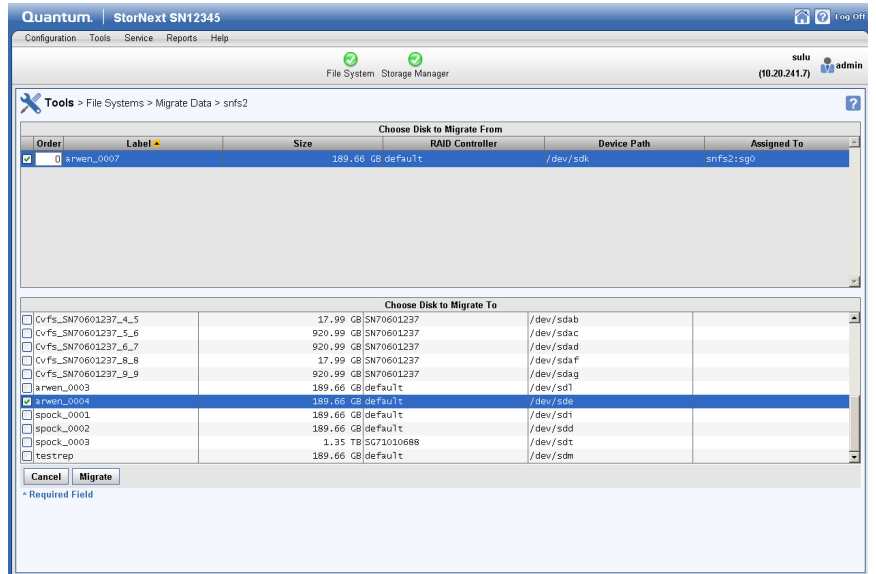
- 1 On the **Tools** menu, click **File Systems**, and then click **Migrate Data**. The **Tools > File Systems > Migrate Data** page appears (refer to [Figure 46](#)).
- 2 Click a **File System**.
- 3 Click a **Metadata/Journal Stripe Group** in the **File System** to migrate (refer to [Figure 47](#)).

Figure 47 Metadata/Journal Data migration page



4 Click **Migrate...** A new page appears (refer to [Figure 48](#)).

Figure 48 Metadata/Journal Data migration (new page)



- 5 Click a **Source Stripe Group**.
- 6 Click a **Destination Stripe Group**.
- 7 Click **Migrate**. The following occurs:
 - The file system stops and is unmounted.
 - The GUI runs the command **sndiskmove**.
 - The source stripe group is re-labeled to **\$LABEL.old**.
 - The destination stripe group is re-labeled to **\$LABEL**.
 - The progress is reported as **percent complete**.
- 8 Click **Refresh** to manually update the status.
- 9 When completed, start and mount the file system using the GUI.

How to Migrate User Data


Use the following procedure to migrate **User Data** using the GUI.

User Data migration runs faster as only areas containing files are processed. Since the file system is still running, it may take several iterations to complete if any clients have open files while the **snfsdefrag** command is executing.

You only specify the source stripe group(s) for **User Data** migration. The migration process will move the data from the source stripe groups to an available user data disk. If any files are not moved due to open file handles, repeat the migration procedure.

- 1 On the **Tools** menu, click **File Systems**, and then click **Migrate Data**. The **Tools > File Systems > Migrate Data** page appears (refer to [Figure 46](#)).
- 2 Click a **File System**.
- 3 Click one or more **User Data Stripe Group(s)** in the **File System** to migrate (refer to [Figure 49](#)).

Figure 49 User Data migration page



4 Click **Migrate...** The following occurs:

- The source stripe groups are set to **read-only**.
- The file system is unmounted on the MDC.
- The GUI runs the command **snfsdefrag**.
- Progress is reported as **percent complete**.

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**.

5 Click **Refresh** to manually update the status.

6 Repeat **Step 1** through **Step 4** until all the files have been migrated off of the source disk.

7 When completed, mount the file system using the GUI.

8 **(Optional)** If you want to re-use the empty source stripe groups, edit the file system and mark the source stripe group as **read-write**.

Note: Migrating stripe groups containing both Metadata/Journal Data and User Data is not supported and cannot be done from within the GUI. Contact Quantum Professional Services to discuss possible workaround procedures.

Note: Quantum recommends you keep User Data and Metadata on separate stripe groups for performance reasons and to allow for stripe group migration using the GUI.

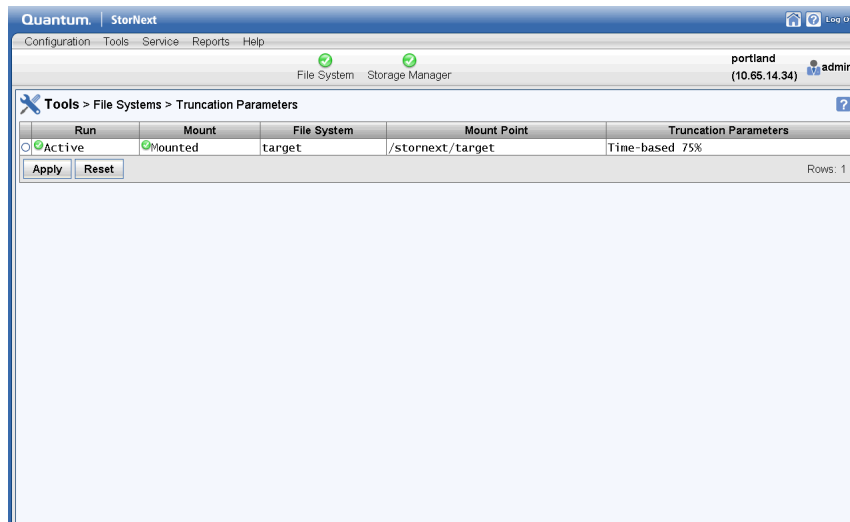
Truncation Parameters

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 50 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.

Manage Quotas

The quota system provides a means for limiting the amount of disk storage consumed on a per user or per group basis across an entire file system or within a designated directory hierarchy. Quota limits apply to the space consumed by disk-block allocations for a user or group, which is not equal to the sum of their file sizes. Disk-block allocations can be less than the file size if the file is sparse, or more if the file system has allocated extra sequential blocks for the efficiency of anticipated future writes.

The Manage Quotas screen enables you to **view, edit, delete** or configure **new** quota values as it pertains to StorNext file system. See Figure 41 and Table 1.

Figure 51 Manage Quotas Screen



The following table provides a brief description of the quota properties, and values you can **view** on the Manage Quotas screen.

Table 1 Manage Quotas Table


Quota Property	Description
Directory Namespace	Specifies the file system and directory name.

Quota Property	Description
Soft Limit	Specifies the soft limit quota value configured on the system. The Soft Limit is the maximum amount of available usage. You are warned upon reaching the Soft Limit quota value.
Usage (Relative to Soft Limit)	Specifies the percent (%) usage relative to the Soft Limit quota value on the system.
Hard Limit	Specifies the hard limit quota value configured on the system. The Hard Limit quota value is the absolute amount of available usage. You cannot go beyond the Hard Limit quota value.
Total Grace Period	Specifies the total grace period configured on the system. The Total Grace Period is used when you have exceeded the Soft Limit quota value, but are still under the Hard Limit quota value. As soon as the Soft Limit quota value has been exceeded, you have the configured Total Grace Period amount of time to free up space to return your usage under the Soft Limit quota value.
Grace Period Remaining	Specifies the grace period remaining on the system. The Grace Period Remaining is the amount of time remaining to free up space to return your usage under the Soft Limit quota value.
Number of Files	Specifies the number of files on the system.

To **edit** the current quota values for a specified file system:

- 1 Under the **Directory Namespace** column, select the file system whose quota values you want to edit.
- 2 Click **Edit...** (see Figure 42).

Figure 52 Edit Quotas Screen



3 Configure the following quota value properties:


- a In the **Hard Limit** field, input a numeric value (decimal values are allowed), and then select the unit of measure from the drop-down. Refer to [Quota Limits](#) on page 120 for additional information on limits.
- b In the **Soft Limit** field, input a numeric value (decimal values are allowed), and then select the unit of measure from the drop-down. Refer to [Quota Limits](#) on page 120 for additional information on limits.
- c In the **Time Limit** field, input a numeric value (decimal values are allowed), and then select the unit of measure from the drop-down.

4 Click **Update** to confirm and save your selection, or click **Cancel** to abort and return to the **Tools > File Systems > Manage Quotas** screen.

To configure **new** quota values for a specified file system:

- 1 Click **New...** (see Figure 43).

Figure 53 New Quotas Screen



- 2 Configure the following quota value properties:
 - a In the **Hard Limit** field, input a numeric value (decimal values are allowed), and then select the unit of measure from the drop-down. Refer to [Quota Limits](#) on page 120 for additional information on limits.
 - b In the **Soft Limit** field, input a numeric value (decimal values are allowed), and then select the unit of measure from the drop-down. Refer to [Quota Limits](#) on page 120 for additional information on limits.
 - c In the **Time Limit** field, input a numeric value (decimal values are allowed), and then select the unit of measure from the drop-down.
 - d For the **Namespace**, select the file system from the drop-down, and then select a directory underneath the **Directory Name** heading.
- 3 Click **Create** to create a new directory namespace with the specified quota values, or click **Cancel** to abort and return to the **Tools > File Systems > Manage Quotas** screen.

To **delete** a configured quota values (see Figure 41):

- 1 Select the **Directory Namespace** whose quota values you want to delete.
- 2 Click **Delete....**
- 3 When asked to confirm the deletion of the configured quota value, click **Yes** to proceed, or **No** to abort.
- 4 When a message notifies you that the quota was successfully deleted, click **OK** to continue.

To **refresh** the configured quota values (see Figure 41):

- Click **Refresh**.

Quota Limits

Each quota entity has two limits associated with it. These are the **Hard Limit**, and the **Soft Limit**.

The **Hard Limit** is the absolute limit which file system space usage should not exceed. Any time the total allocated space is at or over the **Hard Limit**, all further allocations or write requests by the offending user or group will be denied.

The **Soft Limit** is a lesser limit. When you exceed this limit (but not the **Hard Limit**), allocations are still permitted, but a warning will be written to your console. When the **Soft Limit** has been overrun for longer than the **Total Grace Period**, the **Soft Limit** becomes a **Hard Limit** and any further allocations or write requests are denied. When the usage again falls below the **Soft Limit**, allocation requests will again be serviced.

For performance reasons related to the distributed nature of StorNext, quota overruns are not only possible but likely. The overrun size depends upon a number of factors including the size of the allocation request(s) at the time of the quota overrun.

Note: Limits are not enforced against super user accounts.

Note: For all quota types, limits and usage values only apply to regular files, not directories, symlinks, or special device files.

Quota Types

There are **three** types of quotas:

- User Quotas
- Group Quotas
- Directory Quotas

User and **Group Quotas** limit the number of file system blocks that can be allocated by the user or group on which the limit is placed. When quotas are on, the total allocated file system space of all users and groups that own files in the file system are automatically kept.

Directory Quotas are a little different. The system does not automatically keep track of the usage for each directory. The **snquota** command allows directories to be turned into the root of a Directory Quota Name Space (DQNS). Then, the number and size of all files in the directory and all its subdirectories are tracked and (optionally) limited.

Note: When working with **Directory Quotas**, the specified file system must be mounted on the node running **snquota**.

Note: For all quota types, limits and usage values only apply to regular files, not directories, symlinks, or special device files.

Refer to the **snquota(1)** MAN Pages Reference Guide for additional information on the Manage Quotas feature, command syntax usage, and examples.

Renaming a Standalone (unmanaged) StorNext File System

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 using **cvupdatefs**.
 - a **Non-HA mode:**

```
cvupdatefs -R new_file_system_name  
old_file_system_name
```
 - b **HA mode:**

In order to rename the data directory on the secondary you need to manually do that before using **cvupdatefs** on the primary. By default, these directories reside in the `/usr/cvfs/data` directory on UNIX systems and in the `C:\SNFS\data` folder on Windows systems. If you do not rename the data directory on the secondary, it will be left as is, and the HA sync process will create a new data directory with the new file system name.

After renaming the data directory on the secondary, on the primary use the **cvupdatefs** command:

```
cvupdatefs -R new_file_system_name  
old_file_system_name
```

The HA sync process will propagate the change to the secondary.

```
cvupdatefs -R new_file_system_name  
old_file_system_name
```

c Manual HA mode:

In manual HA mode you need to run the same **cvupdatefs** command first on the primary, and then on the secondary:


```
cvupdatefs -R new_file_system_name  
old_file_system_name
```

```
cvupdatefs -R new_file_system_name  
old_file_system_name
```

- 6 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.
- 7 Start the file system, and make it active (**cvadmin**).
- 8 Mount the file system.

For more information, see the **cvupdatefs** man page.

Chapter 4: File System Tasks
Renaming a Standalone (unmanaged) StorNext File System



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.

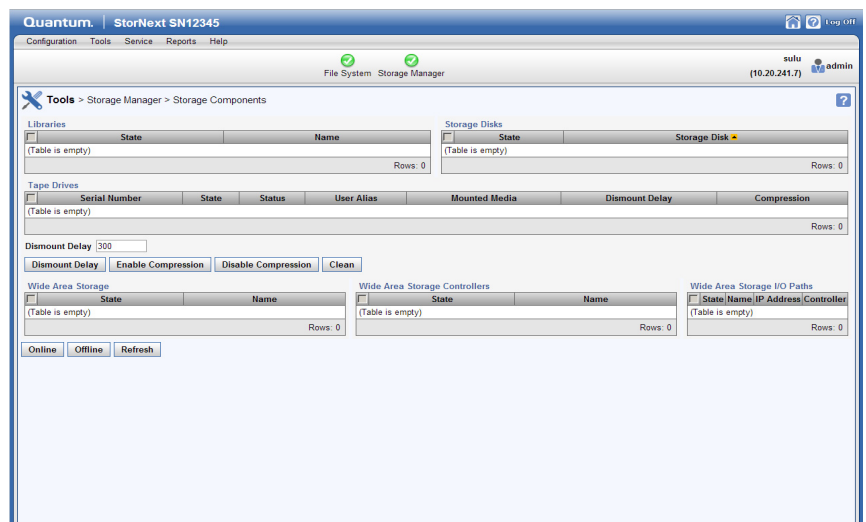
StorNext also features the **Active Vault Policy** feature. See [Active Vault Policy](#) on page 172.

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 54 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 55 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**.

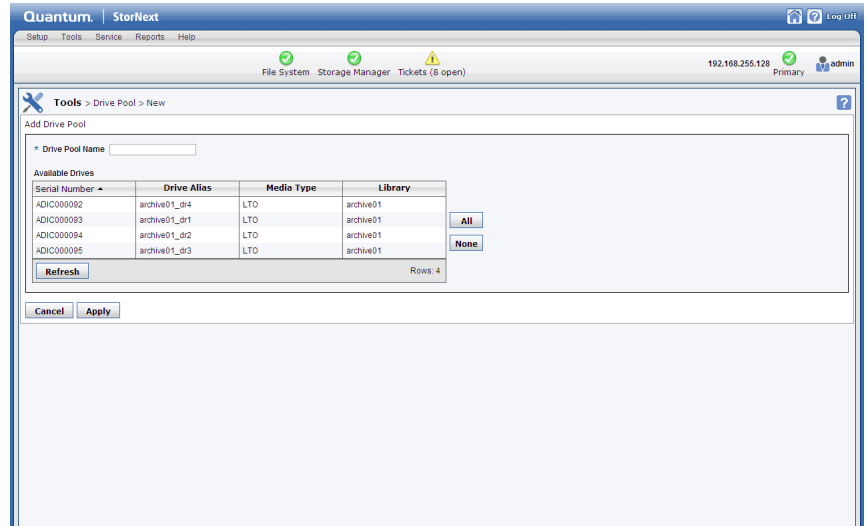
Adding a Drive Pool

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 56 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 57 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 At the **Media Class** field, select the desired media class or choose **Show All Media Classes**.
- 3 If desired, specify a search filter for media IDs at the **Media ID Filter** field. When you specify a filter, only media IDs containing the filter string will be displayed. After you enter the filter string, click **Set Filter** to apply your entry. If necessary, click **Refresh** to update the display.

To remove the filter at any time, click **Clear Filter**.
- 4 Select one or more media to move, or check the box to the left of the **Media ID** heading to select all media.

- 5 At the **Move Media Parameters > Destination Library** field, select the destination library to which you want to move the selected media.
- 6 Click **Apply**.
- 7 When the confirmation message appears, click **Yes** to move the selected media, or **No** to abort.
- 8 Use the Library Operator Interface feature to complete the actual physical move of the media. (See [Library Operator Interface](#) on page 152 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 **fsmedi**n 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 begins 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.

Mixed-Level Tape Drive Compatibility Within the Same Device Family

LTO-3 media in a library containing LTO-5 or LTO-6 drives are considered for store requests unless they are logically marked as write protected. When LTO-3 media is mounted in an LTO-5 or LTO-6 drive, StorNext marks the media as write protected. Quantum recommends circumventing LTO-3 media for store requests by following this procedure:

- 1 In the **Media** menu, click **Attributes**.
- 2 On the **Change Media Attributes** page, select the **LTO-3** media from the list.
- 3 Click **Write Protect**.
- 4 Click **Apply** to make the change.
- 5 Repeat the process for each piece of LTO-3 media.

Notes:

- A similar issue exists for LTO-4 media in a library containing LTO-6 tape drives.
- LTO-3 drives can read but not write LTO-3 tapes.
- A similar issue exists for LTO-5 media in a library containing LTO-7 drives.
- LTO-6 drives can read but not write LTO-5 tapes, and also cannot read LTO-3 tapes at all.

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.

Set WAS Media Availability

This option allows you to set the availability of **namespaces** that have been added to existing WAS destinations. The options are **Available** and **Unavailable**.

- 1 In the **Available Media Actions** list, click **Set WAS Media Availability**. The GUI page displays a list of **namespaces** that exist under different WAS destinations.
- 2 In the **Media** table, click a **namespace**.

- 3 In the **New WAS Media State** list, click **Available** or **Unavailable**.
- 4 To accept your selection, click **Apply**.

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

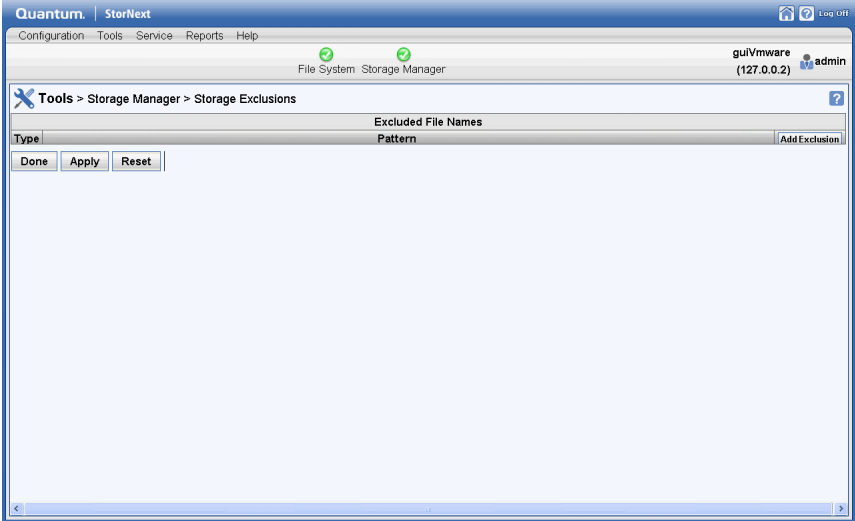
OR

- 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 58 Storage Exclusions Screen




Adding an Exclusion Pattern

To add a new exclusion:

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

Figure 59 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

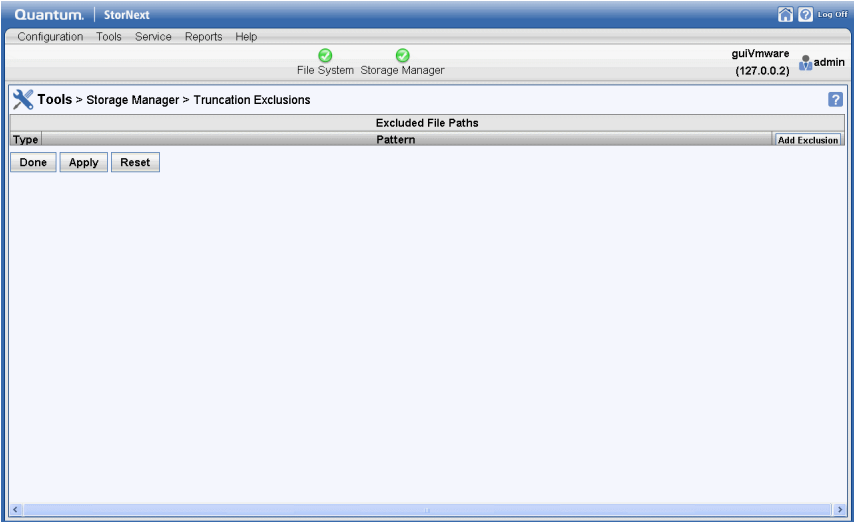
OR

- 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 60 Truncation Exclusions Screen

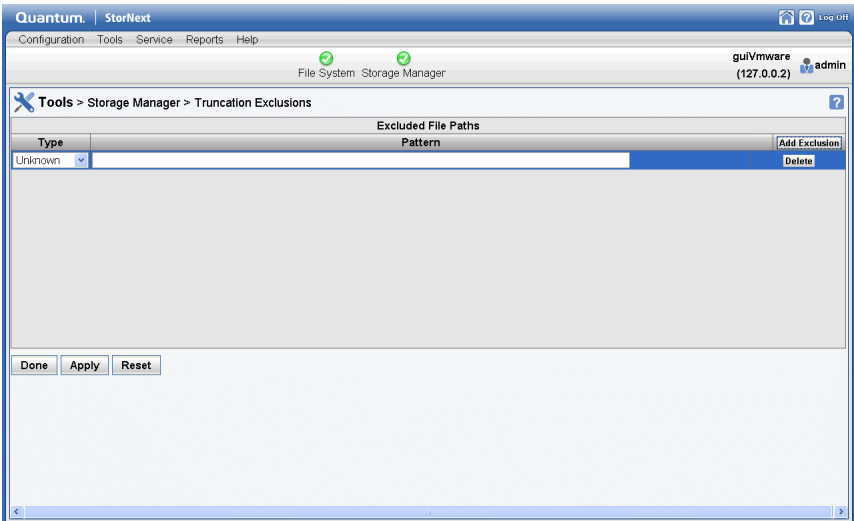


Adding an Exclusion Pattern

To add a new exclusion:

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

Figure 61 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:

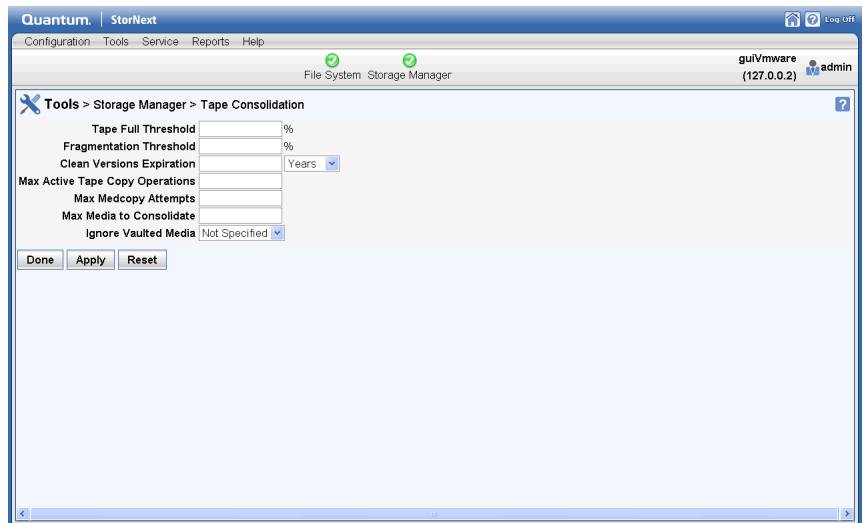
- 1 Setting configuration parameters by specifying criteria for determining which tapes are fragmented
- 2 Creating a schedule to clean old inactive versions of files on tapes
- 3 Creating a schedule for tape defragmentation

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 62 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.

Note: When using the StorNext GUI to defragment or copy tapes on systems with multiple libraries, the GUI does not allow you to specify a library in which the target tape resides. This issue applies only to systems with multiple libraries, archives or partitions attached. To defragment tapes or to copy media to tapes in a desired library, run **fsmedcopy**.

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 155.

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
 - **Details:** Displays information or errors from the back-end system. If information is not available, the field is blank.
- 1 Choose **Library Operator Interface** from the **Tools > Storage Manager** menu. The **Library Operator Interface** Screen appears.

Figure 63 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 In the **Port** list, select 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**.

Note: The **Library Operator Interface (LOI)** page and the **Media Action** page may not display immediately if you have thousands of media requiring attention in the LOI page (for example, media moves, media ejects, etc.) at the same time. A general estimate is that the delay could be approximately 10 seconds for every 500 media requiring attention on the LOI page.
This situation is extremely unlikely to occur, and you may never experience this delay unless you need to export thousands of media at one time. In that situation, you can avoid this issue by breaking up media action requests into smaller chunks.

Software Requests

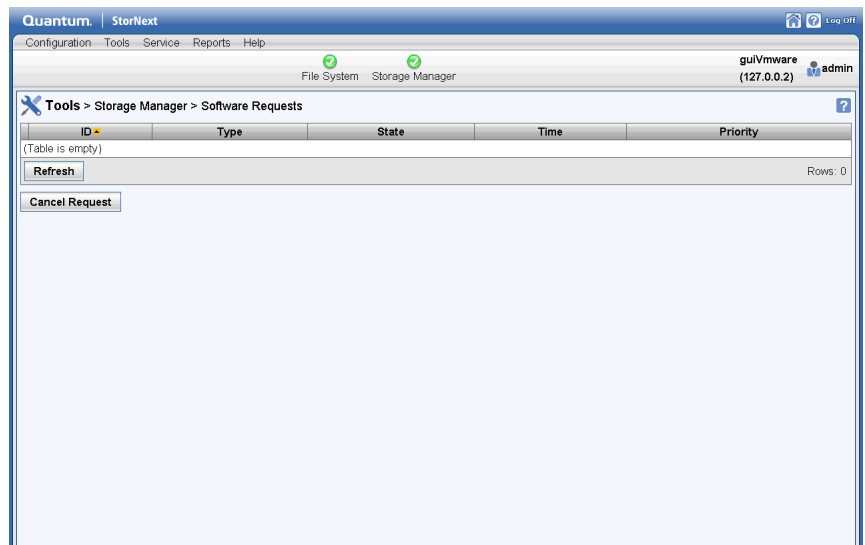
The Software Requests menu option enables you to view software requests currently in process, and to cancel requests.

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 64 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 configuration files, and file system journal files.

Note: MySQL saves an image of the portions of the database journals which have changed.

- **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 149.)

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](#).

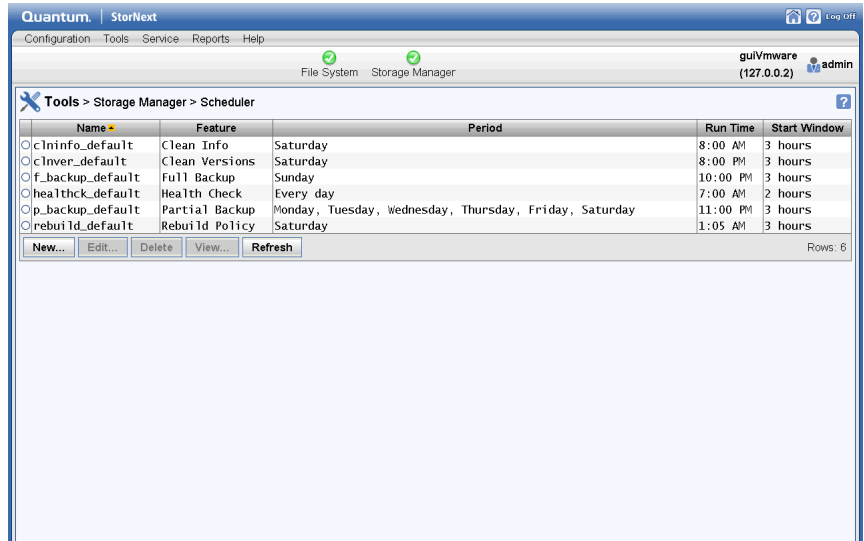
Note: To ensure that scheduled StorNext tasks are started at the correct time, StorNext servers should be rebooted whenever changes are made to the system time.

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 65 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 66 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 67 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.


The data mover process, `fs_fmover`, runs on the metadata controller (MDC) and clients, allowing up to one `fs_fmover` process per tape drive or storage disk (SDisk) stream to run at one time on the MDC and each client.

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

The DDM 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-3 reservations. 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.

The StorNext Distributed Data Mover feature requires that SCSI-3 persistent reservations are enabled. Refer to parameter `FS_SCISI_RESERVE` in `/usr/adic/TSM/config/fs_sysparm.README` to direct the StorNext Manger to use SCSI-3 persistent reservations.

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 -C -K  
0x123456  
[root@sfx14]# sg_persist -n -d /dev/sg78 -o -C -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.

Redhat and SUSE mover clients requires installing these additional quantum-supplied DDM packages:

- quantum_curl
- quantum_openssl
- snfs-mover
- quantum_unixODBC
- snlufs

Note: these packages might require using the **--force** option.

Follow these installation steps for each client:

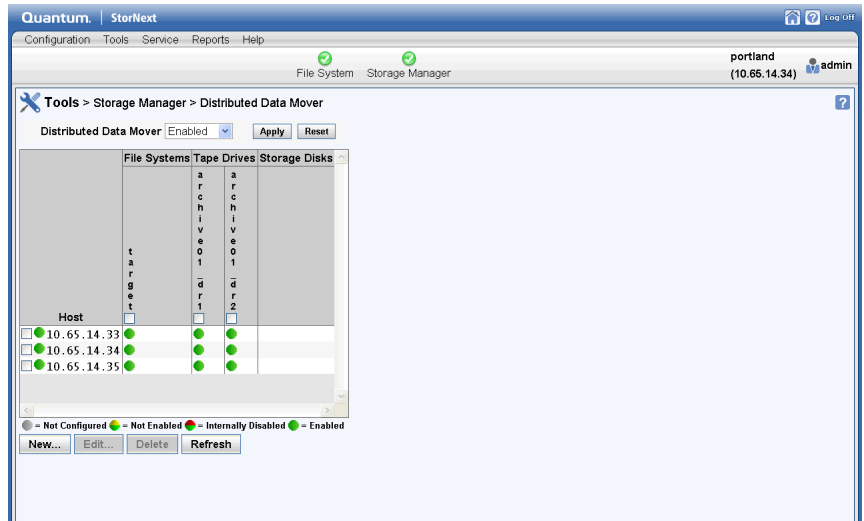
- 1 Log in as root.
- 2 Download the client with DDM package from the MDC.
- 3 Install the rpms in the DDM package .tar archive.

- For a new client installation, run either the command `rpm -ivh *.rpm` or `rpm -Uvh *.rpm`
- For a client upgrade, execute the following:
 - 1 `/etc/init.d/cvfs fullstop`
 - 2 `rpm -Uvh *.rpm`
 - 3 `rpm -Uvh snfs*`
 - 4 `service cvfs start`

Accessing Distributed Data Mover

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

Figure 68 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 page’s **Distributed Data Mover** 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 168.

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 select **Disabled**, **Enabled**, or **Threshold**, click **Apply** 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.

Note: When configuring Distributed Data Movers (DDM), the mount path must exactly match the mount path on the MDCs because the GUI assumes DDM clients have the same directory structure as MDCs. DDM does not do this automatically on DDM client systems; it is a manual process. If a drive, for example, is replaced on a DDM client, the customer must create the directory structure on the new drive on the DDM client to match the MDC mount paths.

Adding a Host

- 1 If you have not already done so, navigate to **Tools > Storage Manager > Distributed Data Mover**.
- 2 Click **New**. Fields appear where you can enter host information

Figure 69 DDM Screen New Host



- 3 At the **Host** field, enter the host name you are adding.
- 4 Enter the remaining fields in the upper portion section of the screen:
 - **Enabled**
 - **Mover Threshold**

- **Maximum Movers (Active)**
- **Maximum Movers (Client or Standby)**

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, navigate to **Tools > Storage Manager > DDM**.
- 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, navigate to **Tools > Storage Manager > DDM**.
- 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 364.

Active Vault Policy

Introduction

Vaulting of media is a process that moves media from a library to a vault that frees up library slots for additional media. Traditionally, vaulting is used for data archival purposes, and for enabling a library to manage more media than it can physically hold. On a StorNext managed system, vaulting can also lower the used capacity of the Storage Manager license.

Active Vault Policy is a StorNext feature that enables you to configure, and schedule automatic execution of custom vault policies for library media. The feature also helps you manage your Storage Manager license more proactively.

This section contains an overview of how StorNext Active Vault Policy works, the steps required to configure a vault, configure a vault policy, and usage tips for the feature.

Overview

Active Vault Policy offers a flexible and fully configurable vault policy that makes it easy for you to customize to your vault needs. You can schedule an Active Vault Policy to run automatically at periodic intervals, or manually on demand.

Policy Options

Active Vault Policy uses the following two system parameters to control the start and stop of the vault process:

- **ACTIVEVAULT_HIGH_USE**: Specifies the percentage of used license capacity at which to begin the Active Vault Policy process (the **default** value is **95.0%**).
- **ACTIVEVAULT_LOW_USE**: Specifies the percentage of used license capacity at which to stop the Active Vault Policy process (the **default** value is **90.0%**).

These two system parameters tie Active Vault Policy directly to Storage Manager used license capacity, allowing you to proactively manage your licensed capacity.

Active Vault Policy uses the following system parameter to select media for vaulting consideration:

- **ACTIVEVAULT_FULL_PERCENT**: Specifies the percentage value used to check against, to determine if a medium is full enough for vaulting consideration by Active Vault Policy (the **default** value is **90.0%**).

You can change the default values of the above system parameters by using the `fs_sysparm_override` file, or override these values for any individual policy to fit your vault needs, as described below.

Active Vault Policy also offers a wide range of criteria for selection of media to vault. Any combination of the following vault policy options (which come from StorNext command **fsactivevault**) can be used to select media for vaulting consideration:

- The list of archives or libraries that media could be selected from.
- The list of copy numbers that media could be selected from.
- Whether or not media belongs to the policies listed.
- Whether or not media is in MIGRATE class.
- The minimal time media has not been accessed.
- The minimal used percentage on a media (to override **ACTIVEVAULT_FULL_PERCENT**).
- The minimal used space size on a media.
- The remaining space on a media if it falls to certain size.
- The override value for **ACTIVEVAULT_HIGH_USE**.
- The override value for **ACTIVEVAULT_LOW_USE**.

Refer to the **fsactivevault** man page for additional information on how to use these policy options.

Active Vault Policy works with multiple vault policies with different vaulting start and stop controls, media selection criteria, and schedules. Each time an Active Vault Policy is run, StorNext invokes the **fsactivevault** command to control the vaulting process and selects media for vaulting based on how the policy was configured. If system used capacity meets or exceeds the **ACTIVEVAULT_HIGH_USE** (or its overridden) value, qualified media which meets the configured vault criteria will be vaulted until used capacity is below the

ACTIVEVAULT_LOW_USE (or its overridden) value or qualified media is exhausted.

Active Vault Policy also offers useful tools for Active Vault policy planning, validation, and troubleshooting. Refer to section [Tools](#) on page 180 for additional information on these tools.

Configuring Active Vault

To use Active Vault Policy, perform the following procedure:

- 1 Configure the library vault.

Active Vault Policy works with multiple vaults, external or internal to the library. An example of the latter is Quantum's Scalar i6000 Tape Storage Library that provides the ability to store vaulted media inside the library using its unique Active Vault partition capability. Additional configuration steps are required to use the i6000 Active Vault partition. Refer to section [Configuring External and Internal Vaults](#) on page 175 for details on how to configure external and internal vaults.

- 2 Create and schedule the vault policy.

Run the command **fsschedule** to create, and schedule an Active Vault Policy.

Note: You can also use the command **fsschedule** to modify, delete, and report a scheduled vault policy.

Currently, the StorNext GUI can only be used to view, edit, and delete the running schedule portion of a vault policy. You cannot create, edit, or schedule an entire vault policy via the GUI.

Refer to the **fsschedule** man page for details on how to create and schedule a vault policy, and other available options.

Refer to the **fsactivevault** man page for additional information on the media selection criteria available for a vault policy.

The StorNext command **fsschedlock** can be used to lock, unlock, and report locks on scheduled Active Vault policies. Refer to the **fsschedlock** man page for additional information.

Refer to section [Command Syntax](#) on page 179 for examples of some of the commands mentioned in this section.

- 3 Wait for scheduled vault policy to run, and vault media.

The scheduled execution of a vault policy may result in one or more media selected for vaulting. If media was selected for vaulting, the StorNext GUI will issue an **Action Required** notification to inform you of the required action to complete the physical move of vaulting media from the library to the designated vault.

On the StorNext GUI, click the **Action Required** icon to display the **Library Operator Interface** screen (alternatively, navigate to **Tools > Storage Manager > Library Operator Interface**) and complete the media move to the vault. If the vault is external, the vaulted media will be moved to the library mailbox, for manual media removal. If the vault is internal as provided by the Scalar i6000 Active Vault partition, the vaulted media will be moved directly to the i6000 Active Vault partition, bypassing the library mailbox. Refer to section [Library Operator Interface](#) on page 152 for additional information about StorNext's GUI Library Operator Interface.

Configuring External and Internal Vaults

Perform the following procedure to configure an external vault:

- 1 On the StorNext GUI **Configuration** menu, click **Storage Destinations**. The **Configuration > Storage Destinations > Libraries** screen appears.
- 2 Click **New...** The **Configuration > Storage Destinations > Libraries > New** screen is displayed.
- 3 In the **Type** field, click **Vault**.
- 4 In the **Name** field, enter a name for the vault.
- 5 Click **Apply** to confirm, or **Cancel** to abort and return to the **Configuration > Storage Destinations > Libraries** screen.

Vaulting processes for this type of vault will occur at the mailbox.

Perform the following procedure to configure the Scalar i6000 Internal Vault:

Note: For proper functional use with the StorNext Active Vault Policy, it is required that the following general attributes and settings are properly configured within the Scalar i6000 Tape Storage Library and the StorNext system.

Part 1 of 2 - On the Scalar i6000 Tape Storage Library:

- Verify the required licenses for i6000 Active Vault features are installed.
- Verify the required licenses for i6000 EDLM features, if preferred for this configuration, are installed.
- Verify the Active Vault partition(s) are correctly established in the i6000 library for use with StorNext Active Vault Policy.
- Verify the EDLM partition is correctly established, if preferred for use in this system environment.
- Verify the SNAPI Client Plug-in for Scalar i6000 is installed and configured for use and communication to the StorNext SNAPI.

Note: You can download the SNAPI Client Plug-in for Scalar i6000, and access official documentation here:

<http://www.quantum.com/ServiceandSupport/SoftwareandDocumentationDownloads/S6K/Index.aspx>

To download the SNAPI Client Plug-in for Scalar i6000:

On the web page, click the **DRIVERS** tab, and then click **DOWNLOAD** under the SNAPI Client Plug-in for Scalar i6000 heading.

For additional information, refer to the *Scalar i6000 SNAPI Plug-in Release Notes*, and the *Scalar i6000 i10 User's Guide* which can be found on the **DOCUMENTATION** tab of the web page.

- Verify the policies for Active Vault and (or) EDLM, are properly configured for use within the i6000 library, and prepared for interrelated use with StorNext, as preferred.
- Verify the Active Vault partitions within the i6000 library are identified for their given name and spelling, so that the creation of the StorNext vault(s) intended for use with i6000 Active Vault partition(s) follow the same naming standards.

Part 2 of 2 - On the StorNext system:

- Verify the StorNext SNAPI is installed and configured for use and communication to the SNAPI Client Plug-in for Scalar i6000.

- Verify any external vault(s) configured in StorNext for use with the mailbox must be configured with a unique name to any vault names intended for use with the i6000 Active Vault partitions.

Perform the remaining procedure to configure the Scalar i6000 internal vault:

- 1 On the StorNext GUI **Configuration** menu, click **Storage Destinations**. The **Configuration > Storage Destinations > Libraries** screen appears.
- 2 Click **New...** The **Configuration > Storage Destinations > Libraries > New** screen is displayed.
- 3 In the **Type** field, click **Vault**.
- 4 In the **Name** field, enter a name for the internal vault partition previously created in the i6000 library.

Note: The name of the newly created StorNext vault must be identical to the named vault partition in the i6000 library.

Vaulting processes for this type of vault will occur within the i6000 internal library Active Vault partitions.

Access and Retrieval of Media from Vaults

The following procedure allows you to access media contained in the i6000 Library – “Library Managed” Active Vault partition(s) and control the media’s subsequent movement to the host (StorNext) managed “Standard” partition for use.

Access to media located in the i6000 Active Vault partition may be needed when files stored on media need retrieval.

The procedure for access and movement of media are managed through a combination of user actions involving the i6000 Library Management Console – Move Media wizard and use of the StorNext GUI Library Operator Interface. These steps will include internal use of the I/E element - the mailbox.

Note: It is crucial that the operator should not physically touch (open or close) the mailbox during this process.

Perform the steps below for media access and movement from the Active Vault partition of the i6000 Library to the StorNext partition:

Note: Normal operations used for management of external “shelf” type media vaulting with StorNext do not require use of the procedures below.

- 1 Upon the need to access media from the Active Vault partition and usage of StorNext to implement that process has occurred, a **Library Operator Interface (LOI) – Action Required** icon in the GUI is normally posted. You should disregard any action to the LOI at this time, leaving it for activity later in this process.
- 2 Open the **Operations – Move Media** wizard from the drop-down menu on the LMC screen. Follow the directions to move through the wizard pages.
- 3 On the next page, select the **Source Partition and Destination Partition** as required to gain access to the list of media.
- 4 On the **Select Media to Move** page, ensure that the **Source Slot Type: Storage Slots** is shown (default).
 - a Locate and select the media to be retrieved from the list.
 - b Choose and select the **Destination Slot Type: I/E Slots** from the selectable menu.

Note: The choice of I/E slots for the **Destination Slot Type** is crucial to the movement of media back into the StorNext partition. This I/E selection ensures that the media will be moved to the mailbox where the robot, StorNext and the LOI can access it and update the database correctly.

- 5 Upon completion of the LMC - Move Media processes, the media will have been moved by the robot to the I/E mailbox.
- 6 At this time, return to the StorNext GUI – LOI Action Required process. Initiate the sequence of actions required to complete the operations to move the media from the mailbox and return it to storage within the StorNext “Standard” partition of the i6000 library.
- 7 When the media is returned successfully to the StorNext partition, normal activity to access the media, files and tape devices will occur.

This completes the procedures to access media in the i6000 Active Vault partition and return it to the StorNext partition.

Command Syntax

Below are some usage examples of the Active Vault Policy commands.

To list all scheduled Active Vault policies or a specific policy by name:

```
fsschedule [-f activevault | -n name] [-l]
```

To schedule an Active Vault Policy:

```
fsschedule -a -n name -f activevault -p period [-e weekday | -y monthday] -t runtime [-w window] -- activevault_options
```

To modify a scheduled Active Vault Policy:

```
fsschedule -m -n name -p period [-e weekday | -y monthday] -t runtime [-w window] -- activevault_options
```

The “-- **activevault_options**” are options specified for the Active Vault Policy. Anything after “--” is unique to, and only used for Active Vault Policy.

Refer to the **fsactivevault** man page for a complete list of valid options and descriptions:

```
fsactivevault [ -archive a1,... ]
[ -vault dest ]
[ -copy c1,... ]
[ -used size ]
[ -remaining size ]
[ -age age ]
[ -sort column ]
[ -migrate | -nomigrate ]
[ -pending | -nopending ]
[ -highmark pct ]
[ -lowmark pct ]
```

```
[ -fullpct pct ]  
[ -report ]  
[ -include-policy p1,... ]  
[ -exclude-policy p1,... ]  
[ -capacity ]  
[ -dryrun ]  
[ -limit num ]  
[ -notify level ]  
[ -noheader ]  
[ -debug ]  
[ -help ]  
[ -policy name ]
```

For example, to schedule an Active Vault Policy named `av1` to run daily at 1:00 AM, that vaults from an archive named `i6k` to a vault named `vault01`, at most 10 qualified media which have been used for copy number 1, and have not been accessed for at least 1 month:

```
fsschedule -a -n av1 -f activevault -p daily -t 0100  
-- -archive i6k -vault vault01 -limit 10 -copy 1 -age  
1month
```

Tools

Active Vault Policy also offers the following features:

- The **`fsactivevault -report`** command option is integrated with other vault options to generate a media report tailored to your needs.
- The **`fsactivevault -dryrun`** command option reports (but does not carry-out) the media to be vaulted based on the current policy. It is used internally by SNMS to validate a vault policy prior to scheduling or modifying.
- The **`fsactivevault -notify`** command option sets level (none, error, warn, info) for email notification.
- The SNMS Active Vault Policy high-level operations are logged in:

/usr/adic/TSM/logs/history/hist_01

- The SNMS Active Vault Policy low-level operations are logged in:

/usr/adic/TSM/logs/fsactivevault

Usage Tips

- Prior to deployment of StorNext configurations that will utilize Active Vault Policy functions in the Scalar i6000 library, it is recommended that all required firmware, licenses, software and configuration attributes on the i6000 library are met.
- When use of the i6000 Active Vault partition is desired, the i6000 library must have the SNAPI Client Plug-in software installed and configured. Additionally, the StorNext SNAPI software must be installed and configured as well.
- When creating an internal vault in StorNext for media vaulting to the i6000 Active Vault partition, the vault name must be identical to the Active Vault partitions name. For example, if the i6000 Active Vault partition name is **VaultA**, the StorNext vault name must also be **VaultA**.
- When managing vaulting processes for media storage within the i6000 library, you are required to perform StorNext GUI LOI (Library Operator Interface) activity to complete movement of the selected media to the Active Vault partition.



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](#)
- [Replication Target Relocating Procedures](#)
- [Troubleshooting Replication](#)
- [Data Deduplication Overview](#)
- [Setting Up Deduplication](#)
- [Data Deduplication Functions](#)
- [Replication / Deduplication Removal Procedures](#)

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 198.

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 70 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 389.

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 Manger 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 71 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:

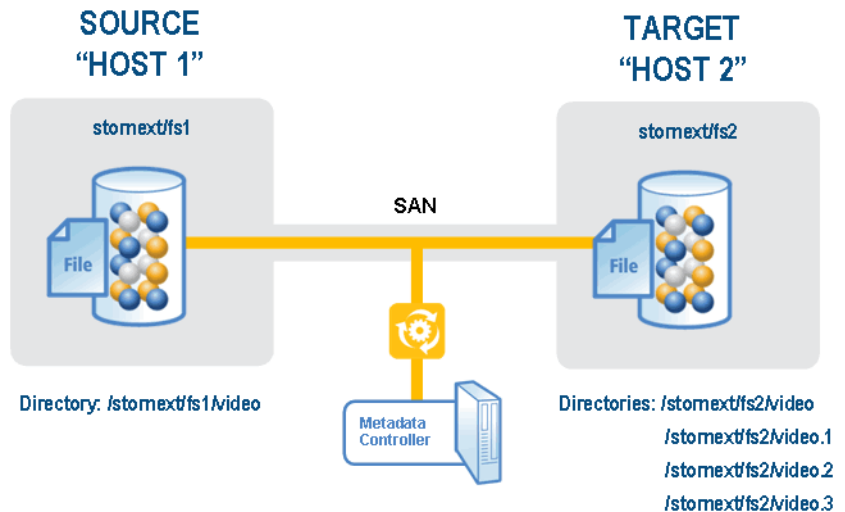
- `/stornext/fs2/video` (Contains the most recent realization)
- `/stornext/fs2/video.1` (Contains the second-most recent realization)

- /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 72 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 73 Replication Scenario
3

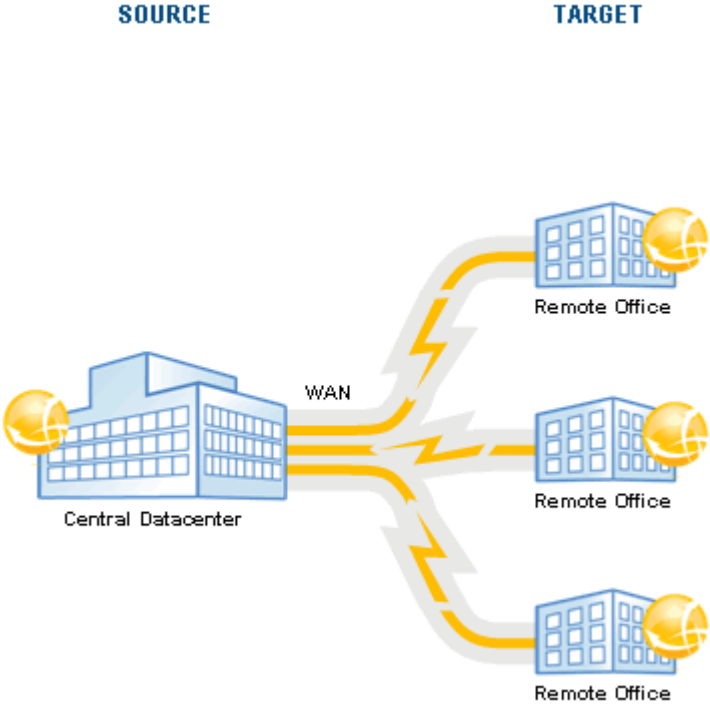


Additional Replication Possibilities

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 74 Replicating From
One Source to Multiple Targets



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

Figure 75 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 76 Non-Supported Replication From Source to Target



"Chained" Replication

"Chained replication" is replicating from source to a 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 194.)

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.

Caution: You cannot move a file from one directory to another directory at the same level, if both directories have the same policy.

Note: 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 10 TB blockpool.

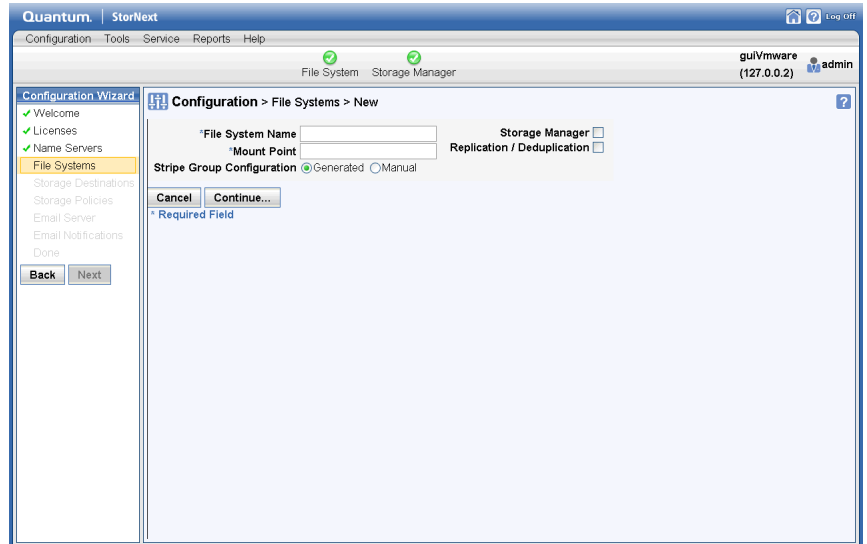
Step 1: Create Source and Target File Systems

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.

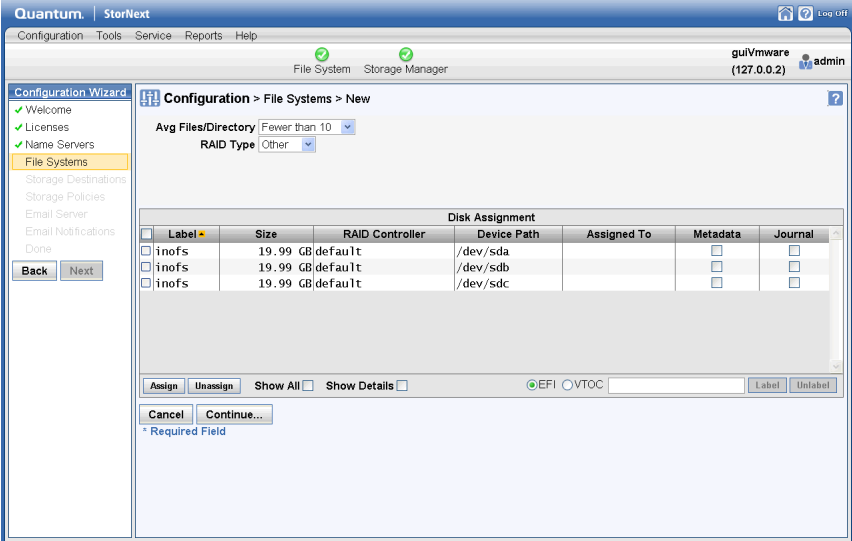
- 1 If you have not already done so, launch the StorNext Configuration Wizard and proceed through **Welcome**, **License** and **Name Servers** steps to the **File System** step.
- 2 The **Configuration > File System** screen appears.
- 3 On the Configuration > File System screen, click **New**. The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

Figure 77 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 Select the **Generate** option, and then click **Continue** to proceed.

Figure 78 Configuration > File System > New Screen 2



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

Figure 79 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 5: File Systems](#) on page 54.)

Note: If you use either the StorNext GUI or the **snpolicy** command to create or modify a replication/deduplication policy, a policy text file is written to the file system. For example, suppose that `/stornext/photos/` is the mount point for file system named `photos`. If a policy named `pol_replicate_1` is created in that file system, a text copy of the policy information called `/stornext/photos/.rep_private/config/pol_replicate_1` is created. 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 `/usr/cvfs/data/fsname/policy_history/policyname.date_time`) In the previous example, the file system name (`fsname`) is `photos` and the policy name is `pol_replicate_1`. So, the backup copy would have a name such as `/usr/cvfs/data/photos/policy_history/pol_replicate_1.2010-10-29_14-07-13`. The backup copy directory is not in the same file system as `photos`. If Storage Manager is used on the machine, all the policy backup files will be backed up along with the regular Storage Manager backups. Quantum suggests that after upgrading to the latest version of StorNext, you run the command `snpolicy_gather -b > some_file`. This will save a copy of your current configuration. The `-b` option will also create a copy of policy information in the `usr/cvfs/data/fsname/policy_history` directory.

Creating a Target File System and Blockpool File System

- 1 Repeat the above file system creation process (sub-Steps 1 - 8) and create the file system you intend to use as a target for replication on this same server.

Note: The Replication/Deduplication option (in sub-Step 5) must be enabled for both source file system and target file system. If a file system (source and/or target) is also to be used for Storage Manager, Data/Migration option (in sub-Step 5) must be enabled.

- 2 Configure another file system for the Blockpool that has neither Data Migration nor Replication/Deduplication enabled.

Note: The step above assumes that the source file system and target file system reside on the same machine. If the target file system resides on a different machine (target machine), follow the above file system creation process (sub-Steps 1 - 8) on the target machine where the target file system for replication resides. In addition, configure another file system for the Blockpool on the target machine that has neither Data Migration nor Replication/Deduplication enabled.

Step 2: Setting up the Blockpool

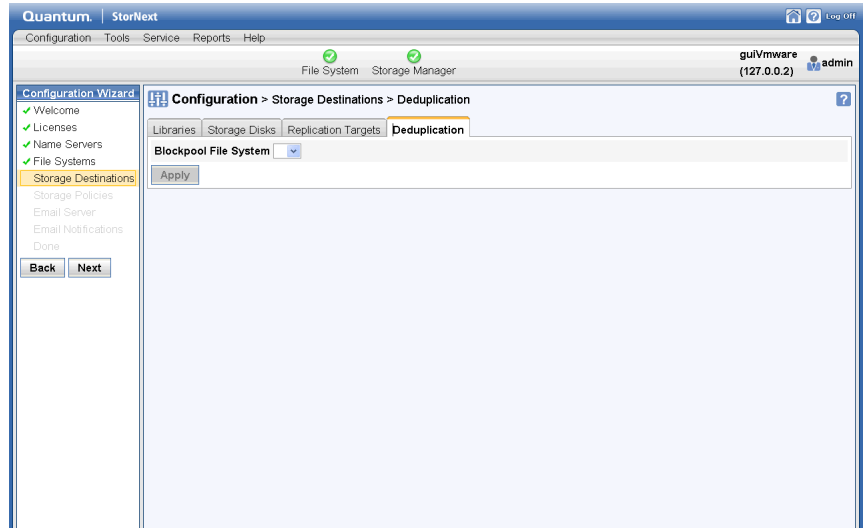
In this step, set up the blockpool on the blockpool file system you just created in the previous step. If the source file system and target file system reside on different machines, perform the following sub steps on both source machine and target machine.

- 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 80 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.

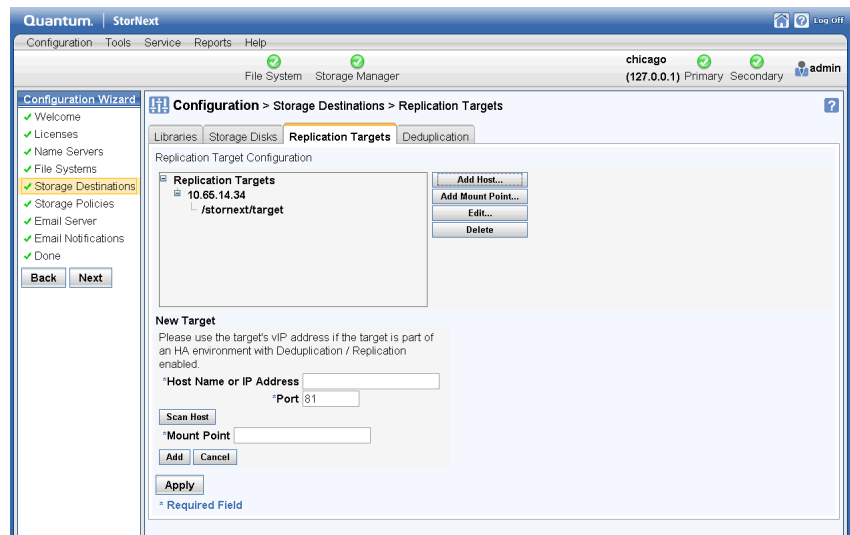
- 4 After you select the blockpool file system, click **Apply**. A background job is started to create the blockpool.

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 81 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. If multiple vIPs are configured for the target HA cluster, select one vIP address that is accessible from the source host.
- 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 216, 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

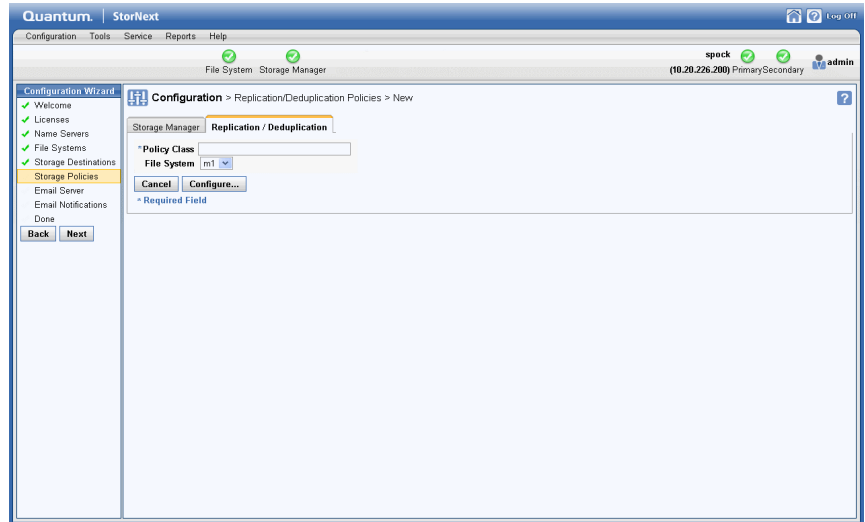
The next step in configuring replication is to create a replication storage policy. This policy contains the replication "rules" specific to your replication source and target file systems. You must create a replication policy for the source directory and enable inbound replication for the target file 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 82 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.

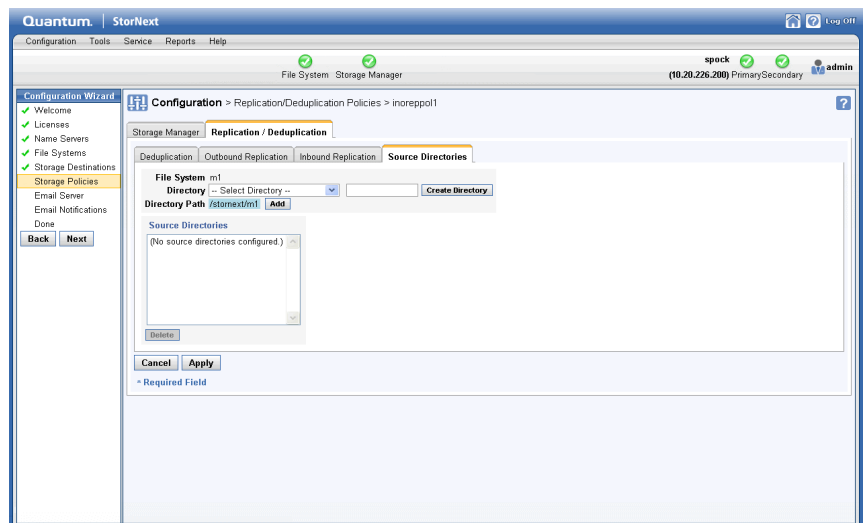
-
- **File System:** Choose the source file system from the dropdown list.

- 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 83 Configuration > Storage Policies > New / Source Directories Screen

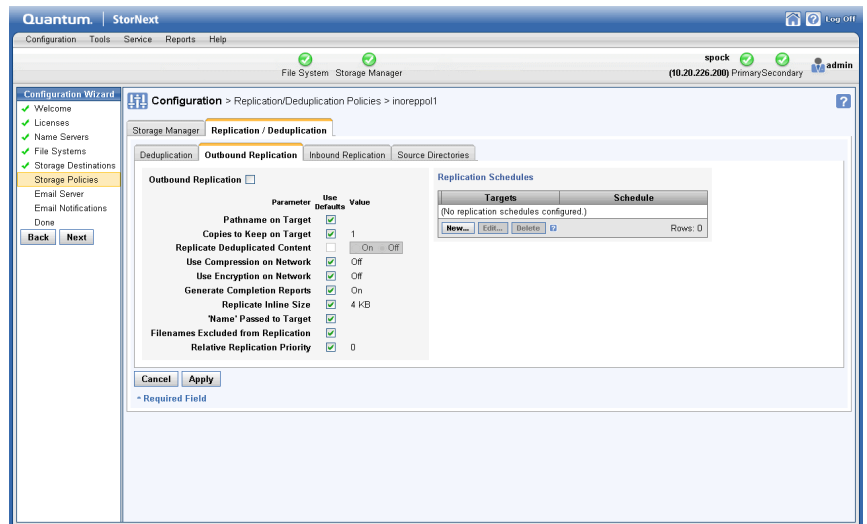


- 3 At the **Directory** field, select from the pulldown list an existing directory you want to use for the policy.
- 4 To create a new directory, enter a directory location at the field to the left of the **Create Directory** button, and then click **Create Directory** to create the specified directory.
- 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 84 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.)

Note: The **Filenames Excluded from Replication** option allows you to exclude specific files from the replication process. This option works the same way as a UNIX shell which lets you pattern match names. For example, entering `*.o` core would exclude all `.o` files and also files named core. You could also skip all core files by entering `rep_skip=core*`.

- 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 85 Outbound
Replication Tab > Replication
Schedule

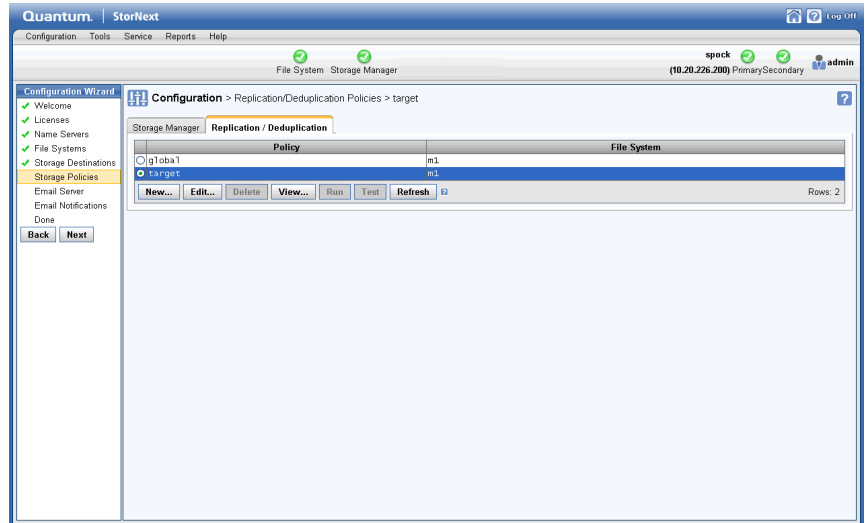


- 5 Under the heading **Select Available Targets**, select the target file system on the target server.
- 6 Create a schedule by making a selection in every column. If you select none of the schedule columns, this creates an unscheduled policy that must be run manually. The schedule shown in [Figure 85](#) 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 86 Configuration > Storage Policies Screen (Select "target")



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

Figure 87 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

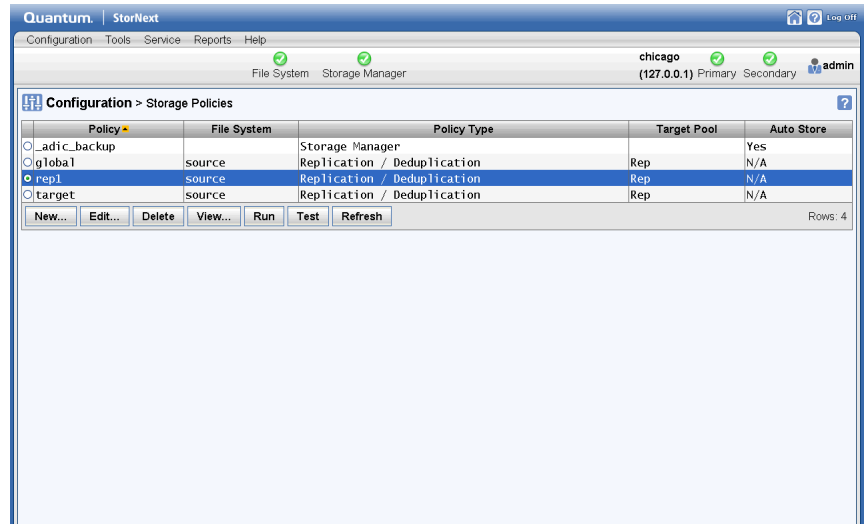
The preceding four configuration steps accomplished the following:

- 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 88 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.

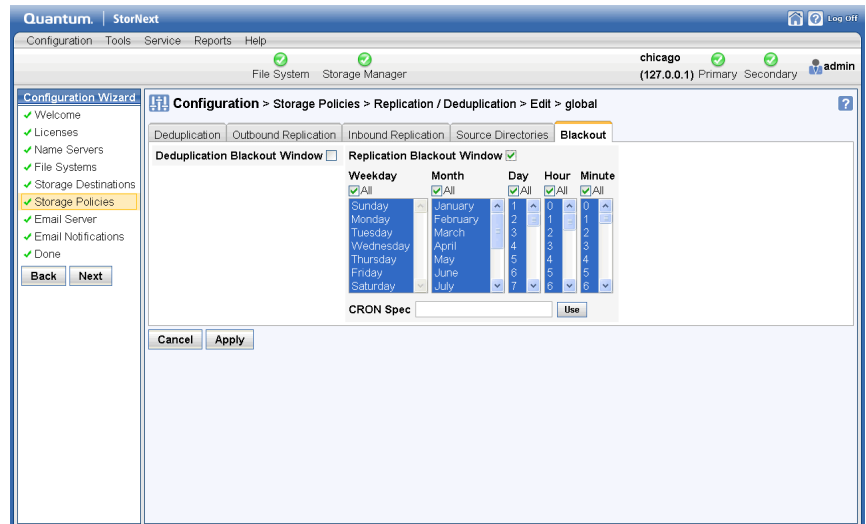
However, note the following caveats about blackouts:

- 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 89 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 190 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 351.

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

Figure 90 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**). If multiple vIPs are configured, select one vIP address for the blockpool server.
- 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.

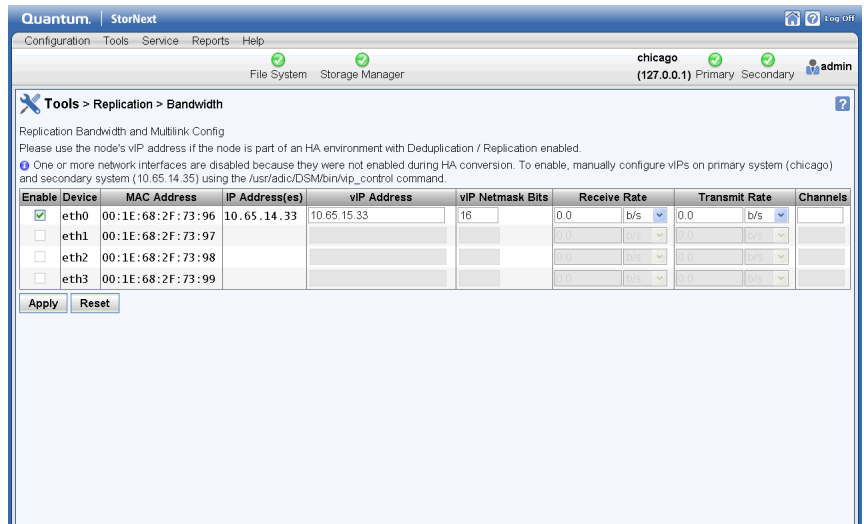
Note: The IP address of the Blockpool Server in HA Clusters can also be configured from the tab **Outbound Replication** and **Inbound Replication** for the "global" policy. There is no difference in configuring from different tabs and the setting of the IP address for the Blockpool from one tab is automatically reflected in the corresponding field of the other two tabs.

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 91 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 88](#).)
- 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.

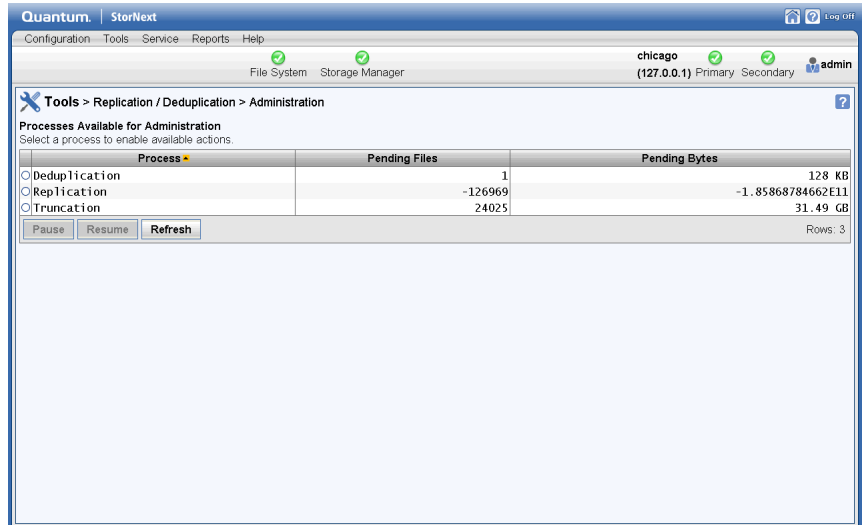
For more information about replication reports, see [Replication Deduplication Reports](#) on page 389.

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 92 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 370.

Replication Target Relocating Procedures

In the normal case, replication occurs between the source file system and the designated target file system(s), either on a schedule or on a demand basis. However, in certain cases, the target file systems could be relocated to meet certain needs. This section provides examples of cases when a target file system could be relocated, and the procedures to relocate a target file system.

This section is aimed at audiences who have advanced knowledge and expertise of StorNext replication/deduplication and are responsible for the configuration, administration of StorNext file systems including replication/deduplication.

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

Conventions	Examples
Command line examples are shown in a bold monospace font, with a hash mark (#) at the start	# /usr/adic/bin/adic_control start
Computer output is shown in a non-bold monospace font	found 0 entries
File and directory names, policy parameters, menu commands and button names are shown in bold font.	/data/upload

Replication Target Relocating Basics

This section discusses several replication terms that are related to Replication Target Relocating. It then describes the cases where Replication Target Relocating is needed. Finally, assumptions are made for the Replication Target Relocating.

Replication Policy

In StorNext, user-defined policy is used to control the replication/deduplication behavior.

A policy is comprised of a set of named and typed attributes. For example, policy attribute "dedup" has type Boolean and can be configured to "on" or "off" to indicate whether deduplication is involved; policy attribute "rep_target" is a parameter used to specify one or multiple replication targets.

Each target has the location of target host and the mount point, also an optional cron-spec can be attached to specify the times when replication is scheduled.

There are several pre-defined named policies, for example policy "global" defines certain policy attributes for the whole file system. Named policy "target" inherits its attributes from policy "global". When its policy attribute "rep_input" is set to on, policy "target" is used as the default policy for a replication target.

The top directory to be replicated in a source file system is called the *replication source directory*. When a replication policy is configured on the source file system, it is assigned to the replication source directory. All files and sub-directories under the replication source directory will inherit the attached policy. When the source directory is replicated to the target file system, the same directory structure is regenerated, in a process called *namespace realization*. The corresponding top directory on the target file system is called the *realized target directory* or *target namespace*. On the target side, normally you turn on the policy attribute "rep_input" for the predefined policy "target" so the replication target directory inherits its policy from policy "target".

Caution: You cannot move a file from one directory to another directory at the same level, if both directories have the same policy.

Deduplicated vs. Non-Deduplicated Replication

The content of files replicated can be either raw file data (non-deduplicated) or deduplicated file data. When non-deduplicated replication is configured, any changes to a file under the replication source directory will cause the whole file to be copied to the target.

For deduplication configuration, a file's unique data content is stored in the local blockpool repository. Each file just stores reference pointers called *blockpool tags* that point to its unique data in the blockpool. When deduplicated replication is configured, the blockpool tags are copied to the target file. If the corresponding unique data segment is not in the target blockpool, the data segment is also copied to target blockpool from the source blockpool.

To configure deduplicated replication, set policy attributes "dedup=on", "rep_dedup=on". If both attributes are set to "off", then the replication is non-deduplicated.

Cross-Mount vs. Network-Mount Replication

A replication source file system and its target file system can be mounted on the same MDC (Metadata Controller) that runs the snpolicyd (the snpolicy daemon that enables replication/deduplication).

This type of configuration is called *Cross-mount of Replication Source Target File System*. If they're mounted on different MDCs, then it is called *Network-mount of Replication Source Target File System*. With cross-mount configuration, the data content is copied locally from source file system to target file system without going through a network. This can potentially lead to higher performance.

Where Target Relocating Is Potentially Needed

Normally, after replication is configured, replication is performed between the replication source and the designated replication target(s). In certain scenarios, it is better to relocate the target for certain needs.

One scenario is the initial replication seeding. Before the first replication starts, the source site could already have millions of files. It could take a very long time if the source and actual target sites are far apart and network bandwidth is limited. It may significantly reduce the time to finish the initial replication if the storage for the target file system can be shipped to the source site and installed in the SAN fabric with the storage for source file system. Then a replication cross-mount can be configured to implement replication. Once the initial replication is complete, the storage for the target file system can be shipped back and installed in the original target site so regular replication can proceed.

Another scenario is data recovery. If the source site suffers a disaster that makes all source data unavailable, the replicated data on the target

site can be used for data recovery. During recovery, the target site data can be replicated back to the new source product host. If the amount of data to be replicated back is prohibitively large, it could be much faster to just ship the storage for the target file system from target site to source site, install it with the storage for the source file system, and perform cross-mount replication. Refer to document “StorNext Replication and Data Recovery” for more information on replication data recovery.

Assumptions about Replication Target Relocating

Since the replication target relocating involves relocating the storage for the target file system, we make the following assumptions:

- The storage components (disk array, disks, etc.) for the target file system can be separated from other storage on the target site, and can be shipped to a remote site. These storage components should be dedicated to the target file system, not shared with other file systems.
- The shipping method for storage devices should offer very high reliability and security. Since storage for target file system is shipped between source and target sites, it should ensure no data damage or compromise occurs during the shipping and handling. Data backup to tape, for example, is needed in case some data should be damaged in shipping and handling.
- The storage components for the target file system can be installed on the source site. The SAN fabric on the source site can offer enough space, power, I/O bandwidth for both source and target file systems.
- The cvlabels for the disks used to construct target file system must be unique. There is no disk label conflict when storage disks are relocated to another location.
- The performance in cross-mount replication is much higher than Network-mount replication, especially when the target is far away from the source site
- The total amount of data to be replicated is prohibitively large, and could take a very long time to replicate.
- The non-deduplicated replication is configured. Since deduplicated replication stores unique data content in a blockpool repository, this does not work with replication target relocating unless the target

blockpool file system is also relocated to the source site. Currently, deduplicated replication configuration is not supported for target relocating. No TSM relation point is configured for the target namespace. Currently, retargeting a target file system if TSM is involved is not supported.

Target Relocating Procedures

This section describes the detailed procedures to relocate a target file system between source and target site.

Collect and Understand Target Replication Configuration

The first step for the replication target relocating is to understand your current target replication configuration. The configuration information includes:

- Which file system is the target file system and where is it mounted?
- What stripe groups and disk devices are used for the target file system and what are the disks' labels?

You can find this information using the GUI (**Configuration>File Systems>Edit>Stripe Group/Disk Management**), or by examining the file system configuration file (fsname.cfgx).

The xml element `<snfs:stripeGroups>` defines the stripe groups and the disks used in the stripe group. Attribute `diskLabel` specifies the cvlabel of the disk. For example, the following stripe group configuration from a file system configuration file shows that one stripe group "sg0" is used to construct the file system. One disk with label "repdisk3" is used in the stripe group.

```
<snfs:stripeGroups>
<snfs:stripeGroup index="0" name="sg0" status="up" stripeBreadth="2097152"
read="true" write="true" metadata="true" journal="true" userdata="true" realTimeIOs="0"
realTimeIOsReserve="0" realTimeMB="0" realTimeMBReserve="0"
realTimeTokenTimeout="0" multipathMethod="rotate">
<snfs:disk index="0" diskLabel="repdisk3" diskType="GENERIC_4294948831" ordinal="0"/>
</snfs:stripeGroup>
</snfs:stripeGroups>
```

- Identify the corresponding storage components (disks, disk array, etc.) for the target file system in the storage SAN fabric.

Relocate Storage Hardware

In Replication Target Relocating, the affected storage hardware needs to be transported from one location to another location.

Follow the procedures below:

- Perform all necessary decommissioning procedures needed for the designated storage components. For example, for an active file system that is to be relocated, make sure there are no active operations on the file system, then unmount the file system, stop the file system using `cvadmin`, etc.
- Follow proper procedures to power off the storage components. Detach them from the current SAN fabric.
- Use reliable and secure transportation to move the storage hardware to the intended location.
- Follow proper procedures to install the storage hardware in the new location's SAN fabric.
- Make necessary SAN configurations to bring back the storage hardware in the new location's SAN fabric.

Procedures for Relocating Target from Cross-Mount to Network-Mount

This is mainly used for initial replication seeding. There are two steps for initial replication seeding using cross-mount:

- 1 [Perform Cross-Mount Initial Replication](#)
- 2 [Relocating the Target File System to the Target Site](#)

Note: If the replication target directory is intended to sit under a TSM relation point, the relation point should be added after the initial replication is complete, and the target file system has been relocated to the actual final target site.

Perform Cross-Mount Initial Replication

Since this is intended for first replication, replication has not been conducted to the target before for the designated replication source directory. We assume either the replication policy has not been configured on either source or target site, or the outbound replication

on the replication source directory is turned off, or the inbound replication on the target file system is turned off.

- 1 If the storage components for target file system have been installed on target site, follow the procedures in section [Relocate Storage Hardware](#) on page 228 to transport the storage components to the source site and install and configure them in the SAN fabric. Otherwise, if the storage components already in the source site, install and configure them in the SAN fabric in the source sites.

- 2 Configure target file system on the source MDC

If the target file system has been configured, the file system configuration file should already exist. Follow the steps below reconfigure target file system on the source MDC.

- a Copy the config file, for example, `tgt1.cfgx`, of the target file system from target MDC to `/usr/cvfs/config` on source MDC.
- b Add the target file system to file `/usr/cvfs/config/fsmlist` so the file system can be started automatically.
- c Add the target file system to mount table `/etc/fstab`. For example, if the target file system is `tgt1` and the mount point is `/stornext/tgt1`, add entry `"tgt1 /stornext/tgt1 cvfs rw 0 0"` to file `/etc/fstab`.
- d Start and mount the target file system:

```
# /usr/cvfs/bin/cvadmin -e "start tgt1"
```

```
# mount tgt1
```

If target file system has not been configured before, create it from GUI.

- e Open the StorNext GUI for the source MDC, navigate to **"Configuration>File Systems"**, then click **New** to create a new file system for the intended target file system on the disks relocated from target site.

- 3 Enable `"rep_input"` for policy `"target"` on the target file system

From the StorNext GUI, navigate to **"Configuration>ReplicationDeduplication Policies"**, select the predefined policy `"target"` for the target file system, click **Edit**, and then change **Inbound Replication** to **On**.

- 4 Add a new target to replication targets.

From StorNext GUI, navigate to “**Configuration>Storage Destinations>Replication Targets**”, click **New**, type the IP address of the source MDC, select the mount point for the target file system, and then add to replication targets.

- 5 Configure source replication policy, setting the target to the newly configured target.

If the source replication policy has not been created already, create one:

- a From StorNext GUI, navigate to “**Configuration>Replication/Deduplication**”, click **New** to create the replication policy. In addition, set other policy parameters, check **Outbound Replication**, and set the target to be the target configured in step 4.
- b Add the designated replication source directory to **Source Directories** for this policy.
- c Click **Apply** when all other parameters are configured.

Note: Since the source replication directory may already have millions of files, the completion of the policy creation could take a long time after you click **Apply**.

If the source replication policy has been configured before, change the target to the new target configured in step 4.

- From StorNext GUI, navigate to “**Configuration>Replication/Deduplication**”, select the source replication policy, click **Edit**, click **Outbound Replication**, delete the original target, add the new target configured in step 4. Click **Apply**.

Note: Since the source replication directory may already have millions of files, the completion of the policy creation could take a long time after you click **Apply**.

- 6 Perform the initial replication.

From StorNext GUI, navigate to “**Configuration>Replication/Deduplication**”, select the source replication policy, and then click **Run**. Alternatively, run the following command:

```
# /usr/cvfs/bin/snpolicy -replicate=mnt_path/source_dir -wait
```

Note: This can take a very long time depending on the total amount of data to be copied from source replication directory to the target file system.

Relocating the Target File System to the Target Site

Once the initial replication is complete, future replication is incremental for the files that change during normal operation. The network should have enough bandwidth to replicate these incremental changes to the target. Now it is time to relocate the target file system to its intended location, the target site.

Follow the procedures below:

- 1 Turn off the policy attribute **Outbound Replication** of the source replication policy after the initial replication is complete so that no further replication occurs to the temporary target.
 - From StorNext GUI, navigate to “**Configuration>Replication/Deduplication**”, select the source replication policy, click **Edit**, click **Outbound Replication**, clear **Outbound Replication** to turn off replication.
- 2 Follow the procedures in section [Relocate Storage Hardware](#) on page 228 to detach the storage components for the target file system, transport to the target site, install and configure them properly.
- 3 If the target file system was not created on the target site, copy the file system configuration file and configure it:
 - a Copy the target file system config file from the source MDC to /usr/cvfs/config on the target MDC.
 - b Add the target file system to /usr/cvfs/config/fsmlist
 - c Use **mkdir(1)** to create a mount point for the file system
 - d Add a mount entry to file /etc/fstab so the target file system can be mounted automatically when it is started.
- 4 Start and mount the target file system by running **cvadmin**, for example, run:

```
# /usr/cvfs/bin/cvadmin -e “start tgt1”
```

```
# /bin/mount tgt1
```

- 5 Redirect the replication target of the source replication policy to the new target.

- a On the source MDC, create a new replication target based on the target file system on target site:

From StorNext GUI, navigate to “**Configuration>Storage Destinations>Replication Targets**”, click **New**, type the IP address of the target MDC, select the mount point for the target file system, add to replication targets, then click **Apply**.

- b Change the source replication policy to the new target:

From StorNext GUI, navigate to “**Configuration>Replication/Deduplication**”, select the source replication policy, click **Edit**, click **Outbound Replication**, mark **Outbound Replication** to turn on replication, delete the previous temporary target and add the new target, and then click **Apply**.

- 6 Perform a Replication Test

Now perform a replication test on source to make sure all replication setup works properly for regular daily replications.

- a Create a temporary file in the replication source directory, for example, run “**touch foo**”.
- b Run `snpolicy -replicate` command on the source MDC to perform namespace replication.

```
# /usr/cvfs/bin/snpolicy -replicate=mnt_path/sourc_dir -wait
```

Procedures to Relocate Target from Network-Mount to Cross-Mount

This is mainly used for replication data recovery. This involves a pre-detach configuration, detaching and transportation, installation and configuration in the source site. After the data recovery is complete, relocate the target back to the original target site.

- 1 Turn off **Inbound Replication** for the pre-defined policy “target” on the target file system.
 - From StorNext GUI, navigate to “**Configuration>Replication/Deduplication**”, select the predefined replication policy “target”, click **Edit**, click **Inbound Replication**, set **Inbound Replication** to **Off**, and then click **Apply**.
- 2 Collect the target file system configuration.

- Follow the procedures in section [Collect and Understand Target Replication Configuration](#) on page 227.

3 Unmount the target file system.

- a Ensure no replication activities are enabled on the target file system
- b Run the following command where `mnt_path` is the mount path of the target file system.

```
# /bin/umount mnt_path
```

4 Stop the target file system.

- Run the `cvadmin` command to stop the target file system:

```
# /usr/cvfs/bin/cvadmin -e "stop tgt1"
```

5 Follow the procedures in section [Relocate Storage Hardware](#) on page 228 to relocate the storage components for the target system to the source MDC.

6 Configure the target file system on the source MDC.

- a Copy the config file, for example, `tgt1.cfgx`, of the target file system from target MDC to `/usr/cvfs/config` on source MDC.
- b Add the target file system to file `/usr/cvfs/config/fsmlist` so the file system can be started automatically.
- c Add the target file system to mount table `/etc/fstab`. For example, if the target file system is `tgt1` and the mount point is `/stornext/tgt1`, add entry `"tgt1 /stornext/tgt1 cvfs rw 0 0"` to file `/etc/fstab`.
- d Start and mount the target file system by running `cvadmin`, for example, run:

```
# /usr/cvfs/bin/cvadmin -e "start tgt1"
```

```
# /bin/mount tgt1
```

So far, the target file system is cross-mounted with the source file system.

7 Follow the proper procedures in document "StorNext Replication and Data Recovery" to perform further operations on this cross-mount target file system.

- 8 After data recovery is complete, if the target storage and file system needs to move back to the original target site, follow the procedures in section [Relocating the Target File System to the Target Site](#) on page 231 to relocate the target storage to the original target site.

Troubleshooting Replication

The Troubleshooting appendix in this guide contains simple troubleshooting procedures related to replication. For more information, see [Troubleshooting Replication](#) on page 647.

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 93 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 5: File Systems](#) on page 54.

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 204 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 82](#).)

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 94 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 92](#).)

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 222.)

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 389.

Replication / Deduplication Removal Procedures

StorNext replication/deduplication provides a tightly coupled set of services for StorNext file systems. Data deduplication removes

duplicated data by identifying duplicated data segments in files and only storing the unique data segments in a blockpool repository. Reference pointers to the unique data segments are stored in files for data retrieval. Replication copies source directories to one or more target directories through scheduled policy or run on-demand. The content of files replicated can be either raw file data (non-deduplicated) or deduplicated file data. Deduplication provides data reduction while replication provides data protection. User-defined policy is used to control the replication/deduplication behavior.

Although the GUI provides support to set up replication/deduplication, it does not support the decommissioning or removal of replication/deduplication. This section aims to provide the necessary procedures to remove the replication/deduplication configuration and return the system to the original configuration without replication/deduplication.

This section is aimed at audiences who have advanced knowledge and expertise of StorNext replication/deduplication and are responsible for the configuration, administration of StorNext file systems including replication/deduplication.

Assumptions

- It is assumed StorNext software 4.1.x or later is installed. The information in this document only applies to those releases.
- It is assumed that full removal of replication/deduplication configuration will be performed. It is not intended for partial removal that leaves some file system's replication/deduplication configuration intact.

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

Conventions	Examples
Command line examples are shown in a bold monospace font, with a hash mark (#) at the start	# /usr/adic/bin/adic_control start
Computer output is shown in a non-bold monospace font	found 0 entries

Conventions	Examples
File and directory names, policy parameters, menu commands and button names are shown in bold font.	/data/upload

Removal Procedures

This section describes the detailed procedures to remove a replication/deduplication configuration and restore StorNext file system to the original configuration. Examples are used to demonstrate how to perform each step.

Collect and Understand Replication/Deduplication Configurations

The first step for the removal of replication/deduplication is to understand your current replication/deduplication configurations. The configuration information includes:

- Which file systems are snpolicy-managed file systems and where are they mounted?
- What replication/deduplication policies have been defined? Is a policy defined for deduplication only, replication only, or replication with deduplication?
- Does the snpolicy-managed file system work as a replication source site, or target site or both? Where are the target sites which are defined in the replication policies on source site?
- Which directories are source directories for the source site file systems, which directories are realized namespaces for the target site file systems? Is there a TSM relation point associated with the source directory or target namespace?

Note: On a target, the realized namespace must land under a TSM relation point.

- Where is the blockpool repository located? Is the file system where the blockpool repository resides only used for the blockpool?

Such information can be acquired either through the GUI or command line. Start from the replication source side host; obtain the list of target

site host(s); then collect the configuration information on all targets. For deduplication-only configuration, there are no target hosts involved.

Obtain Information from StorNext GUI

- From the GUI **“Configuration->File Systems”**, the StorNext File systems are listed with mounting point if mounted. Select a file system and click **“Edit”**, if **“Replication/Deduplication”** is checked, then the file system is an snpolicy-managed file system.
- From the GUI **“Configuration->Storage Policies->Replication/Deduplication”**. Policies for all snpolicy-managed file systems are displayed. Select a policy and click **“View”**. Determine whether deduplication is **“on”** or replication is **“on”**. If replication is **“on”**, and **“Outbound replication”** is also **“on”**, this is a policy defined for source site replication. You can also find the associated directories, the source replication directories. You’ll also find the target location and directories it populates.
- From the GUI **“Configuration->Destinations->Replication Targets”**, you’ll find all defined replication targets (the host and the directory to be replicated into).
- From the GUI **“Configuration->Storage Policies->Storage Manager”**, you’ll find all TSM policy classes. Select a class and click **“View”**, you’ll find the directories associated with the class. If the directory is also a snpolicy-managed directory or is the parent of a snpolicy-managed directory, then the directory has both an snpolicy policy and a TSM relation point associated.
- From the GUI **“Configuration->Destinations->Deduplication”**, the blockpool file system is displayed. Normally the blockpool repository is in a sub-directory **“blockpool”** of the mount point of the blockpool file system.

Obtain Information from Command Line

- 1 To obtain snpolicy-managed file systems, run:

```
# /usr/cvfs/bin/snpolicy -listfilesystems=localhost
```

```
fsname:      snfs1 [replication dedup] up 110:49:07  
mount:      /stornext/snfs1
```

blockpool: Running up 110:49:07

2 To obtain policy information, run:

```
# /usr/cvfs/bin/snpolicy -listpolicies=mnt_path
```

```
# /usr/cvfs/bin/snpolicy -listpolicies=/stornext/snfs1
NAME: default
  NAME: global           inherits from: default
  NAME: target          inherits from: global
  NAME: rep_pol1        inherits from: global
    DIR: /stornext/snfs1/test (key: 371660016)
      active: dedup      inherits from: rep_pol1
    DIR: /stornext/snfs1/test1 (key: 371660016)
      active: dedup rep  inherits from: rep_pol1
```

The above output indicates that there are two snpolicy-managed directories: directory **/stornext/snfs1/test** has a replication policy associated, while directory **/stornext/snfs1/test1** has a policy with both replication and deduplication configured.

3 To view the policy configuration, run:

```
# /usr/cvfs/bin/snpolicy -dumppolicy=mnt_path -name=policy_name
```

```
# /usr/cvfs/bin/snpolicy -dumppolicy=/stornext/snfs1 -name=rep_pol1
name=rep1
inherit=global
dedup=on
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_min_size=4K
trunc_low_water=0
trunc_high_water=0
rep_output=true
rep_dedup=true
rep_report=true
rep_target="target://stornext/tgt1@10.65.189.39:"
rep_inline_size=4K
```

From the output, it can be seen that this is a replication source policy. It has `rep_dedup = true`, so deduplication is enabled. It also has `rep_output = true` so this is a replication source policy, the replication target is host 10.65.89.39 (**rep_target**), the intended namespace will be realized under `/stornext/tg1` on the target. As a result, the associated directory (`/stornext/snfs1/test1`) is a source replication directory.

If `rep_input = true`, the policy is a target policy. Normally this is configured on policy "**target**". A host that has a policy (typically policy "**target**") configured with **rep_input** turned on is a target host.

- 4 To check whether a directory is associated with a TSM relation point, run:

```
# /usr/adic/TSM/bin/fsdirclass path
```

This will show the TSM policy class if the directory is associated with a TSM relation point.

5 The blockpool repository can be found in file `/usr/cvfs/config/blockpool_root`

```
# cat /usr/cvfs/config/blockpool_root
```

```
BFST_ROOT=/stornext/snfs1/blockpool/
```

```
CURRENT_SETTINGS=_stornext1TB
```

BFST_ROOT points to the blockpool repository, in this case, the blockpool is located at `/stornext/snfs1/blockpool`.

Replication Removal on a Target Host

From the previous section, [Obtain Information from Command Line](#) on page 242, you obtained the replication/deduplication configuration on source host and target host(s). Now you start the removal on the target hosts. There are 10 steps described below. Follow these steps to remove replication/deduplication on target host(s).

Note: If there are only deduplication policies, and no replication policy is configured, you should skip this section and jump to section [Replication Removal on a Source Host](#) on page 250.

1 Backup Replication/Deduplication Configurations

In case you need to reuse the current replication/deduplication configuration in the future, it is recommended to back up the current replication/deduplication configuration first. Run:

```
# /usr/cvfs/bin/snpolicy_gather &>snpolicy_dump
```

The configuration information is saved to file `snpolicy_dump`.

2 Suspend Replication/Deduplication Activities

The next step is to suspend replication/deduplication activities so that the `snpolicy` daemon becomes idle. Run the following commands to suspend potential ingest processing, replication processing, truncate processing, blockpool delete and compact processing. Run the following commands where the `mnt_path` is the mount path of the `snpolicy`-managed file system.

```
# /usr/cvfs/bin/snpolicy -runingest=mnt_path -suspend  
# /usr/cvfs/bin/snpolicy -runreplicate=mnt_path -suspend  
# /usr/cvfs/bin/snpolicy -runtruncate=mnt_path -suspend  
# /usr/cvfs/bin/snpolicy -rundelate=mnt_path -suspend  
# /usr/cvfs/bin/snpolicy -compact=mnt_path -suspend
```

In addition, disable `rep_input` by setting `rep_input` to "off" of policy "target".

```
# /usr/cvfs/bin/snpolicy -updatepolicy=mnt_path -  
name=target -policy='rep_input=false'
```

3 Namespace and Private Directory Removal

In this step, we'll find the target keys of all realized and unrealized (due to some error) namespaces, identify whether a realized namespace lands under a TSM relation point, then remove all realized namespaces and clean up their content in the snpolicyd-managed private directory.

- a Obtain the realized namespace directories and target keys. Run:

```
# /usr/cvfs/bin/snpolicy -listrepcopies=mnt_path
```

```
# /usr/cvfs/bin/snpolicy -listrepcopies=/stornext/tgt1  
source://stornext/snfs1@10.65.189.57:14500/test?key=2231941 ->  
target://stornext/tgt1@10.65.189.39:?key=350  
  0 -> (Not Realized)  
source://stornext/snfs1@10.65.189.57:14500/test?key=9465955003 ->  
target://stornext/tgt1@10.65.189.39:?key=430  
  0 -> /stornext/tgt1/repository/test
```

In this example, there were two replication streams. The first one had target key 350, but it was not realized due to some error. The second one had target key 430. It had 1 copy and was realized under `/stornext/tgt1/repository/test`.

- b Determine whether the realized namespace is under a TSM relation point. Run:

```
# /usr/adic/TSM/bin/fsdirclass /stornext/tgt1/repository/test
FS0070 27 1002737568 fsdirclass completed: /stornext/
tgt1/repository/test located in class smpoll1.
```

The output of the command indicates that the realized namespace was associated with a TSM relation point.

c Remove realized namespaces

Note: This is for StorNext 4.1 **only**. Skip this step if your StorNext version is 4.2.0 or greater.

For each realized replication stream, run:

```
# /usr/cvfs/bin/snpolicy -rmrepcopy=mnt_path -
key=tgt_key -allcopies
```

```
# /usr/cvfs/bin/snpolicy -rmrepcopy=/stornext/tgt1 -
key=430 -allcopies
```

```
I [0127 10:49:48.191082 24776] Removed namespace 0 at /
stornext/tgt1/test
```

d Remove content in private directory

For each replication stream (whether realized or not), remove the content in the private replication directory. Run:

For StorNext 4.1.x, run:

```
# /usr/cvfs/bin/snpolicy -repcleanup=mnt_path -
key=tgt_key
```

For StorNext 4.2.0 and greater, run:

```
# /usr/cvfs/bin/snpolicy -repcleanup=mnt_path -
key=tgt_key -allcopies
```

```
# /usr/cvfs/bin/snpolicy -repcleanup=/stornext/tgt1 -
key=2162719
```

```
I [0127 11:04:26.790013 25224] Removed 0 old files
```

e Check whether all replication target streams have been processed

Run snpolicy command **listrepcopies** for each snpolicy-managed file system, ensuring all replication streams have been processed. If the output of the command still shows replication streams, go back to **Step c** and **Step d** to remove the remaining replication streams.

f Check for any replication source policies and directories

Run the snpolicy command **listpolicies** for each snpolicy-managed file system on this target host to see whether there are any policies defined for source side replication directories or dedup-only directories. If there are directories defined for replication source or dedup-only, go to Section 2.3 for the replication removal on the source host.

```
# /usr/cvfs/bin/snpolicy -listpolicies=/stornext/tgt1
```

```
NAME: default
```

```
NAME: global                inherits from: default
```

```
NAME: target                inherits from: global
```

The above output indicates that there is no policy or replication directory that has not been cleaned up.

4 Stop Snpolicy Daemon and Blockpool Server

In this step, the snpolicy daemon and blockpool server need to be stopped. Run the following commands:

```
# /usr/cvfs/bin/cvadmin -e "stopd snpolicyd"
```

```
# /usr/cvfs/bin/bp_stop
```

Check whether the process "snpolicyd" and "blockpool" have been stopped. You may run command "**ps -ef**" to check active processes. If the above commands still cannot stop them, try to kill them by running command "**kill -9 pid**".

5 Remove Snpolicy-managed Private Directories

For each snpolicy-managed file system, a private directory is created to store replication/deduplication related information. For a realized namespace directory that lands under a TSM relation point, a private directory is also created under the relation point. Run the following commands to remove them:

```
# /bin/rm -rf mnt_path/.rep_private
```

For each relation point that has realized namespace landed under it, run

```
# /bin/rm -rf relation_point_path/.rep_private
```

6 Remove Replication History Logs

Note: Skip this step if you want to retain the Replication History log files.

Run the following command where **fsname** is the file system name of the snpolicy-managed file system.

```
# /bin/rm -rf /usr/cvfs/data/fsname/rep_reports/*  
# /bin/rm -rf /usr/cvfs/data/fsname/policy_history
```

7 Remove Related Event Files

The snpolicy managed event files are located under `/usr/adic/TSM/internal/event_dir`. Run the following commands to remove any existing snpolicy-managed event files.

```
# /bin/rm -f /usr/adic/TSM/internal/event_dir/*.blocklet  
# /bin/rm -f /usr/adic/TSM/internal/event_dir/*.blocklet_delete  
# /bin/rm -f /usr/adic/TSM/internal/event_dir/*.blocklet_truncate  
# /bin/rm -f /usr/adic/TSM/internal/event_dir/*.replicate  
# /bin/rm -f /usr/adic/TSM/internal/event_dir/*.replicate_src
```

8 Remove Blockpool and its Configurations

This step removes a blockpool repository and its configuration files. As mentioned in section [Obtain Information from Command Line](#) on page 242, the blockpool repository path can be found from file `/usr/cvfs/config/blockpool_root`. Run the following commands:

```
# /bin/rm -rf blockpool_repository_path  
# /bin/rm -f /usr/cvfs/config/blockpool_root  
# /bin/rm -f /usr/cvfs/config/blockpool_config.txt
```

If the file system where the blockpool repository resides is used only for blockpool, you may use it for other purpose or unmount it and remove the file system to reuse the disks in other file system.

9 Turn off the “Snpolicy-managed” Attribute in File System Configuration File

This step turns off the “snpolicy-managed” attribute of the currently snpolicy-managed file systems. Perform the following for each snpolicy-managed file system:

- a Unmount the snpolicy-managed file system

```
# /bin/umount /stornext/tgt1
    b Stop file system
# /usr/cvfs/bin/cvadmin -e "stop tgt1"
    c Update the file system configuration file /usr/cvfs/config/
    tgt1.cfgx
# /bin/sed -e 's/<snfs:snPolicy>true/
<snfs:snPolicy>false/g' /usr/cvfs/config/tgt1.cfgx
>tgt1.tmp
# /bin/mv tgt1.tmp /usr/cvfs/config/tgt1.cfgx
    d Start file system
# /usr/cvfs/bin/cvadmin -e "start tgt1"
    e Mount file system
# /bin/mount -t cvfs tgt1 /stornext/tgt1
```

10 Restart StorNext GUI

The StorNext GUI service needs to be restarted in order to view the changed configurations. Run:

```
# /sbin/service stornext_web restart
```

Replication Removal on a Source Host

The replication removal on the source host is very similar to that on a target host. For simplicity, refer to the corresponding step in a previous section if a step is the same as mentioned before.

On a replication source host, snpolicy-managed directories need a policy assigned directly to them. The policy can be either deduplication only, replication only or replication with deduplication. For directories that have deduplication enabled, the content may have been truncated by the snpolicy daemon. Snpolicy command `removepolicy` will retrieve the truncated content back before removing the policy from a file if no TSM relation point is associated with it.

The removal of replication/deduplication on host side has 11 steps. It is assumed you have already collected the replication/deduplication configuration described in section [Collect and Understand Replication/Deduplication Configurations](#) on page 241.

1 Backup Replication/Deduplication Configurations

This saves the replication/deduplication configurations on the source host. See [Backup Replication/Deduplication Configurations](#) on page 245 on how to back up.

2 Suspend Replication/Deduplication Activities

This is similar to [Suspend Replication/Deduplication Activities](#) on page 245 except that there is no need to change policy “target”.

Note: If you have multiple snpolicy-managed file systems, you must stop replication/deduplication activities for each file system.

3 Remove Policy from Snpolicy-managed Directories

For each snpolicy-managed directory obtained from section [Obtain Information from Command Line](#) on page 242 (from the output of snpolicy command **listpolicies**), remove the policy key from the directory. Note, this can take a long time if the directory has millions of files. Run the following command where **dir_path** is the path of the snpolicy-managed directory:

```
# /usr/cvfs/bin/snpolicy -removepolicy=dir_path
```

```
# /usr/cvfs/bin/snpolicy -removepolicy=/stornext/snfs1/test
I [0126 09:48:06.955781 28768] Removed policy from /
stornext/snfs1/test.
```

Note: If you have multiple snpolicy-managed file systems, run this command for every snpolicy-managed directory on each snpolicy-managed file system.

After all snpolicy-managed directories have been removed the associated policies, run snpolicy command **listpolicies** to ensure no directory is left.

```
[root@ylu-rep-src1 ylu]$ snpolicy -listpolicies=/
stornext/snfs1
```

```
NAME: default
```

```
NAME: global                inherits from: default
```

```
NAME: target                inherits from: global
```

```
NAME: rep_pol1              inherits from: global
```

4 Remove Replication Targets from StorNext GUI

In the StorNext GUI, click **Configuration > Storage Destinations > Replication Targets**, and delete all replication targets defined there.

5 Stop Snpolicy Daemon and Blockpool Server

Stop the snpolicy daemon and blockpool server on the source host. See **Stop Snpolicy Daemon and Blockpool Server** on page 248 for details.

6 Remove Snpolicy-managed Private Directories

This step removes the snpolicy-managed private directories. See [Remove Snpolicy-managed Private Directories](#) on page 248 for details.

Note: On the source host, no private directory is created under a TSM relation point, so it is not necessary to remove a private directory under a TSM relation point as shown in [Remove Snpolicy-managed Private Directories](#) on page 248. Also, the private directories for all snpolicy-managed file systems must be removed.

7 Remove Replication History Logs

See [Remove Replication History Logs](#) on page 249 for details.

8 Remove Related Event Files

See [Remove Related Event Files](#) on page 249 for details.

9 Remove Blockpool and its Configurations

See [Remove Blockpool and its Configurations](#) on page 249 for details.

10 Turn off “Snpolicy-managed” Attribute in File System Configuration File

See [Turn off the “Snpolicy-managed” Attribute in File System Configuration File](#) on page 249 for details.

11 Restart StorNext GUI

See [Restart StorNext GUI](#) on page 250 for details.



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.
- [Lattus Certificates](#): View, create, import, convert, download, and delete Wide Area Storage (WAS) certificates.
- [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

- [Upgrade Firmware](#)

Note: This menu option is only available on StorNext M660, M440 and M330 Metadata Appliances.

[HTTPS Configuration](#) on page 281

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.

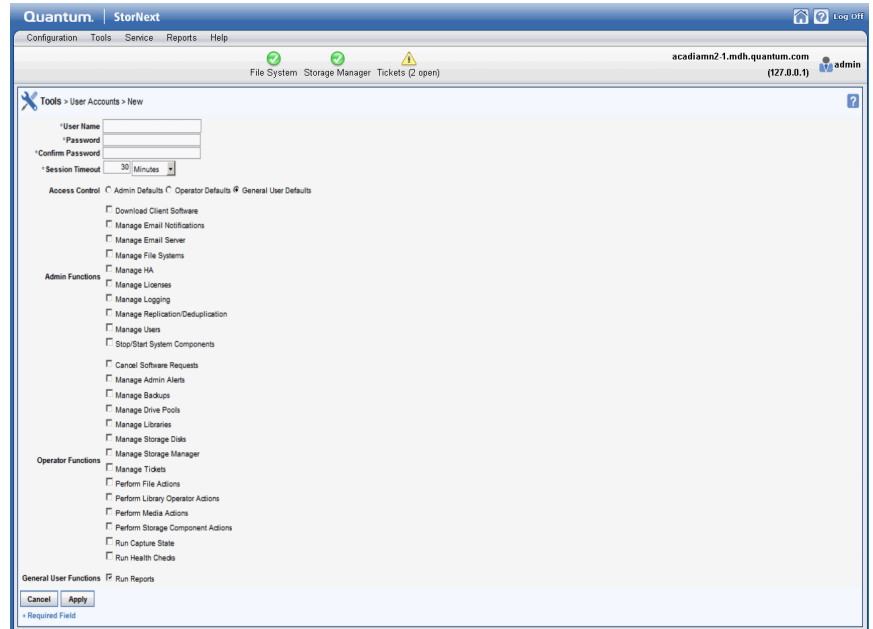
- 1 Choose **User Accounts** from the **Tools** menu. The **User Accounts** page appears. All existing users and the admin are shown.

Figure 95 User Accounts Page



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

Figure 96 New User Page



- 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 In the **Session Timeout** field, type a number and select **Minutes** or **Hours** from the drop-down list. The **Session Timeout** specifies the number of minutes or hours that a session can remain idle before the server terminates it automatically. The **default** is 30 minutes, and the valid range is from 10 minutes to 12 hours.

Note: Access to the **Session Timeout** feature is available when a user has the **Manage Users** privilege checked within the **Admin Functions** heading.

- 6 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.

- 7 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


- 8 When you are satisfied with the permissions you have assigned, click **Apply** to save your changes. (To exit without saving, click **Cancel**.)
- 9 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** page, select the user whose information you want to view, and then click **View**. A page shows the parameters for the selected user.

Figure 97 View User Page



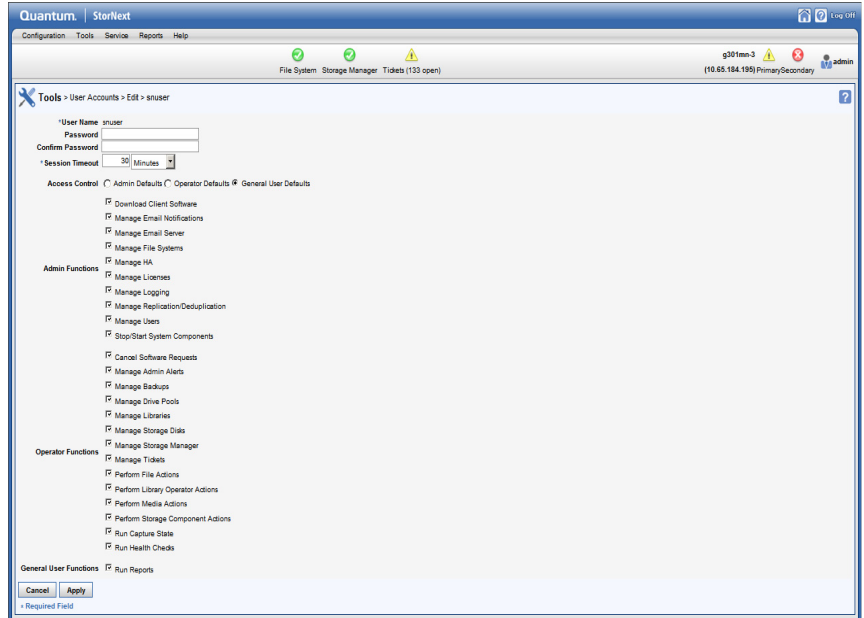
- 2 When you are finished viewing user profile information, click **Back** to return to the User Accounts page.

Modifying an Existing User

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

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

Figure 98 Edit User Page



- 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** page, select the user you want to delete, and then click **Delete**. (See [Figure 95](#).)

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

Note: You cannot delete the **admin** user account.

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

Enable or Disable User Accounts

This feature allows you to re-enable or disable any GUI accounts on the system, including the default **admin** and **service** user accounts. If you disable the **admin** user account, a warning message appears to inform you that if you disable the **admin** user account, you may not be able to fully administer the system unless another user with all privileges has been created and is enabled.

The **Enable** and **Disable** features requires the **Manage Users** privilege to be checked within the **Admin Functions** section. If all users with **Manage Users** privilege are disabled, changes cannot be made.

- 1 From the **Tools > User Accounts** page, select the user account, and then click **Enable** or **Disable**.
- 2 When the confirmation message appears, click **Yes** to proceed, or **No** to return to the **User Accounts > [admin name]** page without applying the changes.

If a user account is disabled, or an incorrect password is entered, the following text is displayed:

Could not login due to

- Incorrect Username and/or Password
- Account has been disabled

Note: To restore GUI access, contact Quantum Technical Support.

Client Download

The StorNext client software lets you mount and work with StorNext file systems. Note that LAN client and LAN client gateways are included with the standard client software packages. For more information about LAN Clients, see [About StorNext LAN Clients](#) on page 3 and [Gateway Server and Client Network Tuning](#) on page 449.

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 162.

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** page appears.

Figure 99 Client Download
Page



- 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 100 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**.

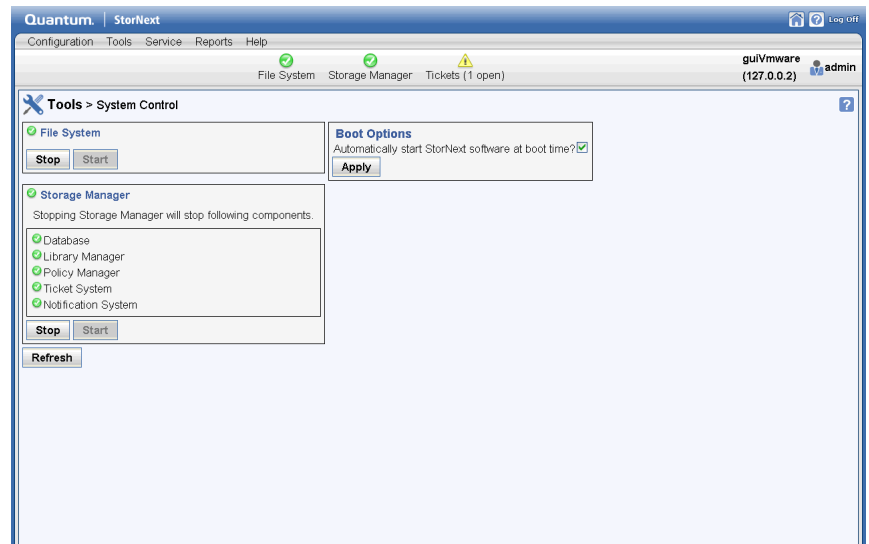
- 9 Continue with the installation procedure for your operating system as described in the *StorNext Installation Guide* or the online help.

System Control

The **System Control** page 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 page 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** page, choose **System Control** from the **Tools** menu. The **Tools > System Control** page appears.

Figure 101 System Control Page



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**.

Lattus Certificates

The Tools menu's **Lattus Certificates** option enables you to manage, and perform various actions to the public and private certificates that various applications requiring SSL authentication use. To access the **Tools > Lattus Certificates** page, on the **Tools** menu, click **Lattus Certificates** (see [Figure 102](#)).

For configuration details, see [HTTPS Configuration](#) on page 281.

If you are working on a Lattus system which does not have an existing SSL certificate, see [Creating a Single Public/Private SSL Certificate for use on Both a StorNext MDC and a Lattus System](#) on page 282.

HTTPS Default CA ROOT Certificate File or Path

Starting with StorNext 5 release 5.2, `/usr/cvfs/config/ssl` is no longer the default repository referenced by Storage Manager for SSL certificates when using HTTPS. The default certificate file or repository will depend on the OS vendor:

- **Debian:** `/etc/ssl/certs/ca-certificates.crt`
- **Red Hat:** `/etc/pki/tls/certs/ca-bundle.crt` or `/etc/ssl/certs/ca-bundle.crt`
- **SUSE:** `/etc/ssl/certs/`

Considerations for SUSE Platforms

If you are using `/usr/cvfs/config/ssl` as your certificate repository, you will have a conflict with the default root certificate repository `/etc/ssl/certs`. You have two options (below):

Option 1

Use `/usr/cvfs/config/ssl` as your default certificate repository; set the `FS_OBJSTORAGE_CAPATH=/usr/cvfs/config/ssl` in the `/usr/adic/TSM/config/fs_sysparm_override` file, then copy all the root certificates from `/etc/ssl/certs` to `/usr/cvfs/config/ssl`.

Note: Be sure to execute **c_rehash** on the directory that will be used as your default certificate repository afterward.

Option 2

Do not use `/usr/cvfs/config/ssl`; instead use the default root certificate repository by copying your certificates from `/usr/cvfs/config/ssl` to `/etc/ssl/certs`.

Note: Be sure to execute **c_rehash** on the directory that will be used as your default certificate repository afterward.

Considerations for Red Hat Platforms

If you are using `/usr/cvfs/config/ssl` as your certificate repository, you will not have a conflict with the default root certificate file. You will have to set `FS_OBJSTORAGE_CAPATH=/usr/cvfs/config/ssl` in the `/usr/adic/TSM/config/fs_sysparm_override` file.

How to Update Expired CA Root Certificates

Root Certificates may expire. When they do, you can update all your Root Certificates to the latest available from <http://rpmfind.net/linux/rpm2html/search.php?query=ca-certificates>. Select the one that fits your system.

- 1 Determine the default configured CA Root Certificate configured for StorNext using **libcurl**:

```
# curl-config --ca  
/etc/pki/tls/certs/ca-bundle.crt
```

- 2 Download the RPM that matches your system. In this example, we downloaded `ca-certificates-2014.1.98-65.1.el6.noarch.rpm`.
- 3 View the contents of the RPM.

```
# rpm -q -filesbypkg -p ca-certificates-2014.1.98-65.1.el6.noarch.rpm
ca-certificates          /etc/pki/ca-trust
ca-certificates          /etc/pki/ca-trust/README
.. snip ..
ca-certificates          /etc/pki/tls
ca-certificates          /etc/pki/tls/cert.pem
ca-certificates          /etc/pki/tls/certs
ca-certificates          /etc/pki/tls/certs/ca-
bundle.crt
ca-certificates          /etc/pki/tls/certs/ca-
bundle.trust.crt
.. snip ..
ca-certificates
/usr/share/pki/ca-trust-source/ca-
bundle.supplement.p11-kit
ca-certificates          /usr/share/pki/ca-trust-
source/ca-bundle.trust.crt
```

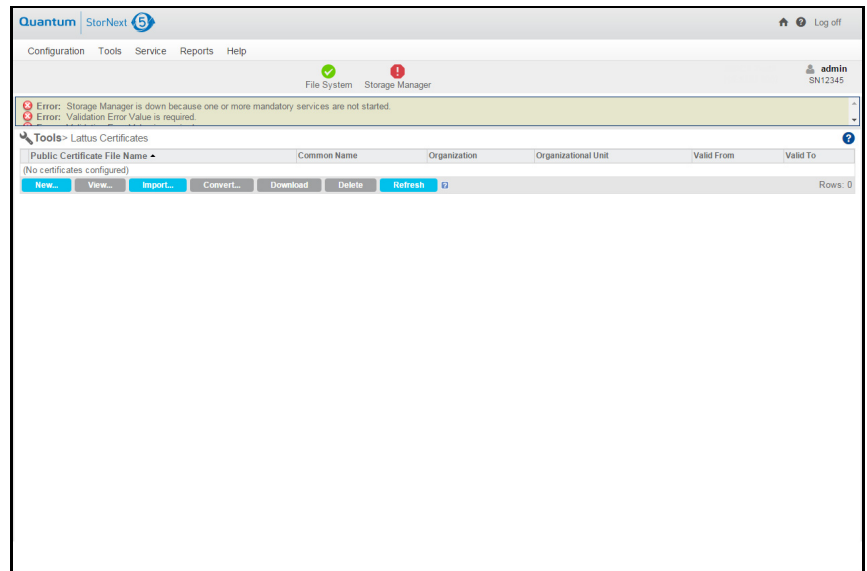
4 Install `/etc/pki/tls/certs/ca-bundle.crt`.

```
# mv /etc/pki/tls/certs/ca-bundle.crt etc/pki/tls/
certs/ca-bundle.crt.bak
# rpm2cpio ca-certificates-2014.1.98-65.1.el6.noarch.rpm | cpio -ivd
/etc/pki/tls/certs/ca-bundle.crt
```

- 5 Install the complete latest RPM.
- 6 Backup any files that you do not want replaced. This step may require you to install required dependencies.


```
# rpm -hiv ca-certificates-2014.1.98-65.1.el6.noarch.rpm  
  
p11-kit >= 0.18.4-2 is needed by ca-certificates-2014.1.98-65.1.el6.noarch  
  
p11-kit-trust >= 0.18.4-2 is needed by ca-certificates-2014.1.98-65.1.el6.noarch
```

Figure 102 Lattus Certificates Page



The table below provides the information displayed for each certificate, on the **Tools > Lattus Certificates** page:

Heading	Description	Examples
Public Certificate File Name	<p>The Privacy Enhanced Mail (PEM) filename and its respective filename extension (for example, .pem, or .der), of the Lattus certificate.</p> <ul style="list-style-type: none"> • There will always be a public certificate, but in some instances where you create a private certificate using this feature, the name will be both the name of the public and private certificate. • There will always be a public file name, and a private if you use this feature to generate your certificates. 	<p>accounts.mycompany.pem accounts.mycompany.der</p>
Common Name	<p>The fully qualified domain name (FQDN) of your server. This must match exactly what you type in your web browser or you will receive a name mismatch error.</p>	<p>*.mycompany.com controller.mycompany.com</p>
Organization	<p>The legal name of your organization. This should not be abbreviated and should include suffixes such as Inc, Corp, or LLC.</p>	<p>Mycompany Corp</p>
Organizational Unit	<p>The division of your organization handling the certificate.</p>	<p>Information Technology IT Department</p>
Valid From	<p>The date the certificate is valid from, in the form of yyyy-mm-dd hh:mm:ss time zone.</p>	<p>1998-08-22 11:41:51 MST</p>
Valid To	<p>The date the certificate is valid to, in the form of yyyy-mm-dd hh:mm:ss time zone.</p>	<p>1998-08-22 11:41:51 MST</p>

To create a **new** certificate, or **import** a certificate, click one of the following buttons on the **Tools > Lattus Certificates** page:

- [New...](#)

- [Import...](#)

Otherwise, select an existing certificate for which you want to perform the action, and then click one of the following buttons:

- [View...](#)
- [Convert...](#)
- [Download](#)
- [Delete](#)

To **refresh** the Lattus Certificates page, click the following button:

- [Refresh](#)

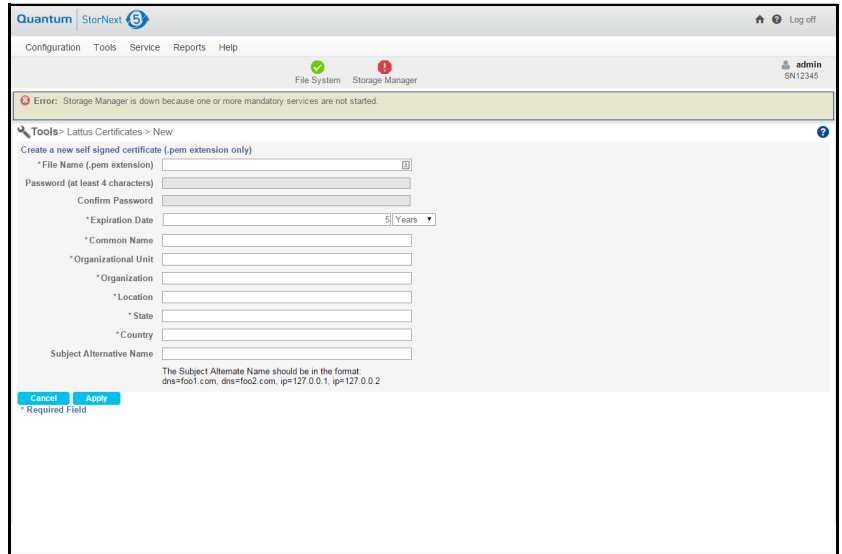
New...

Click **New...** to create a new self-signed certificate.

Note: The filename extension / format of the self-signed Lattus certificate must be `.pem`. You cannot create a self-signed Lattus certificate with filename extension or format other than `.pem`.

- 1 On the **Tools** menu, click **Lattus Certificates**. The **Tools > Lattus Certificates** page appears (see [Figure 102](#)).
- 2 On the **Tools > Lattus Certificates** page, click **New...** The **Tools > Lattus Certificates > New** page appears (see [Figure 103](#)).

Figure 103 Creating a New self-signed Lattus Certificate



3 In the various text boxes, input the appropriate certificate data. The table below describes the various text boxes on the **Tools > Lattus Certificates > New** page:

Note: Text box fields on the **Tools > Lattus Certificates > New** page, designated with an asterisk (*) are required.

Text Box	Description	Examples
File Name (.pem extension)	The Privacy Enhanced Mail (PEM) filename and its respective filename extension of the self-signed Lattus certificate. Note: Adding a certificate with the same name generates an error, instructing you to delete the certificate with that name first.	accounts.mycompany.pem

Text Box	Description	Examples
Password (at least 4 characters)	The Password input is an optional field. If a Password is entered, the input mimics the OpenSSL command password requirements as follows: <ul style="list-style-type: none"> • The Password input, and the Confirm Password input must match. • The Password input must be at least 4 characters, and can be all empty spaces or contain spaces. 	mypassword1234
Confirm Password	See the requirements for the Password input.	mypassword1234
Expiration Date	The Expiration Date low value is at least 1 day in the future. You can input a numeric value, and then select the unit of measurement from the drop-down list. The available unit of measurements are Years, Months, and Days . Note: There is no limit on the high end; however, if you input a value that is out of bounds for OpenSSL, then the OpenSSL command will generate an error.	5 Years
Common Name	The fully qualified domain name (FQDN) of your server. This must match exactly what you type in your web browser or you will receive a name mismatch error.	*.mycompany.com controller.mycompany.com
Organizational Unit	The division of your organization handling the certificate.	Information Technology IT Department
Organization	The legal name of your organization. This should not be abbreviated and should include suffixes such as Inc, Corp, or LLC.	Mycompany Corp

Text Box	Description	Examples
Location	The city where your organization is located.	Englewood
State	The state where your organization is located. This should not be abbreviated.	Colorado
Country	The two-letter ISO code for the country where your organization is located.	US
Subject Alternative Name	The Subject Alternative Name is an optional field. If entered, it should be in the following format (also specified under the text box): dns=foo1.com, dns=foo2.com, ip=127.0.0.1, ip=127.0.0.2	dns=foo1.com, dns=foo2.com, ip=127.0.0.1, ip=127.0.0.2

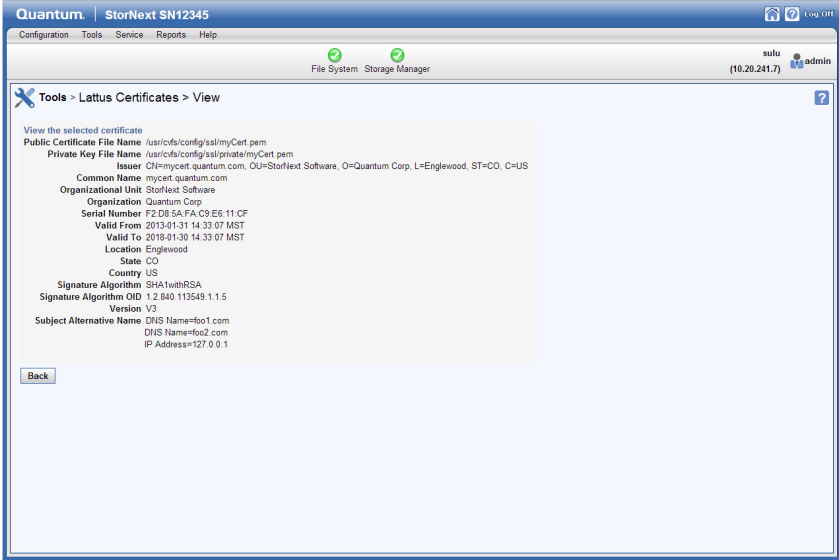
- 4 Click **Apply** to submit your inputs and create a new self-signed Lattus Certificate, or click **Cancel** to reset the form, and return to the **Tools > Lattus Certificates** page. If the submission is successful, your newly created self-signed Lattus Certificate appears on the **Tools > Lattus Certificates** page.

View...

Click **View...** to display the details of a specified Lattus certificate.

- 1 On the **Tools** menu, click **Lattus Certificates**. The **Tools > Lattus Certificates** page appears (see [Figure 102](#)).
- 2 On the **Tools > Lattus Certificates** page, click the option button to the left of a Lattus certificate to select it, and then click **View...** The **Tools > Lattus Certificates > View** page appears (see [Figure 104](#)).

Figure 104 Viewing a Lattus Certificate



The table below describes the various fields on the **Tools > Lattus Certificates > View** page:

Name	Description	Examples
Public Certificate File Name	The Privacy Enhanced Mail (PEM) filename and its respective filename extension (for example, .pem, or .der), of the Lattus certificate.	/usr/cvfs/config/ssl/myCert.pem
Private Key File Name	The filename of the private key in the Lattus certificate. Note: If the certificate was not created through this feature, you will receive following text (in red/bold): Certificates that were imported do not have Private Keys associated to them.	/usr/adic/gui/.ssl/myCert.pem

Name	Description	Examples
Issuer	This property contains the name of the certificate authority (CA) that issued the certificate. The distinguished name for the certificate is a textual representation of the certificate subject or issuer.	CN=mycert.mycompany.com, OU=StorNext Software, O=Mycompany Corp, L=Englewood, ST=CO, C=US
Common Name	The fully qualified domain name (FQDN) of your server. This must match exactly what you type in your web browser or you will receive a name mismatch error.	mycert.mycompany.com
Organizational Unit	The division of your organization handling the certificate.	Information Technology IT Department
Organization	The legal name of your organization. This should not be abbreviated and should include suffixes such as Inc, Corp, or LLC.	Mycompany Corp
Serial Number	The serial number of the selected certificate.	F2:D8:5A:FA:C9:E6:11:CF
Valid From	The date the certificate is valid from , in the form of yyyy-mm-dd hh:mm:ss time zone.	2013-01-31 14:33:07 MST
Valid To	The date the certificate is valid to , in the form of yyyy-mm-dd hh:mm:ss time zone.	2018-01-30 14:33:07 MST
Location	The city where your organization is located.	Englewood
State	The state where your organization is located. This should not be abbreviated.	Colorado
Country	The two-letter ISO code for the country where your organization is located.	US

Name	Description	Examples
Signature Algorithm	The algorithm used to create the signature of the certificate.	SHA1withRSA
Signature Algorithm OID	The object identifier (OID) identifies the type of signature algorithm used by the certificate.	1.2.840.113549.1.1.5
Version	The version number of the certificate.	V3
Subject Alternative Name	The Subject Alternative Name is the name of the user of the certificate. The alternative name for the certificate is a textual representation of the subject or issuer of the certificate.	DNS Name=foo1.com DNS Name=foo2.com IP Address=127.0.0.1

3 Click **Back** to return to the **Tools > Lattus Certificates** page.

Import...

Click **Import...** to import a certificate.

Notes and Considerations

- Files that do not have a .pem extension will need to be converted to .pem for use in SSL communication. See [Convert...](#) on page 278 to convert a file to the .pem format. Quantum only supports the .pem format for certificates.
- You can import one file, which contains multiple public keys. Doing so will create individual rows for each key file with the filename `_multiple.pem`. If any of the multiple keys is deleted, since they comprise the same file, the entire certificate is deleted, and all of the public keys are no longer persisted.
- You can view a certificate on an individual basis by selecting the certificate to view.
- You can import any type of valid public key file, as long as the certificate is not expired. If the certificate is expired, the import will fail, and you will be notified via an **Error** notification. If you import a file with multiple public keys, and any of the public keys in the file are expired, then the entire file is rejected.

- Empty files and files exceeding 10 MB are not permitted. If you want to change the 10 MB limit, you must manually edit the `/usr/adic/gui/config/component.properties` file, and modify the following value: `objectstorage.ssl.maxCertSizeMb=10`
 - You cannot upload a private certificate file; however, you can create a private certificate. If your private / public key is in a `.pem` file, open the file in a text editor and remove the private key.
- 1 On the **Tools** menu, click **Lattus Certificates**. The **Tools > Lattus Certificates** page appears (see [Figure 102](#)).
 - 2 On the **Tools > Lattus Certificates** page, click **Import...** The **Import A Certificate** dialog box appears.
 - 3 In the **Import A Certificate** dialog box, click **Choose File** to select a file to import. The **Open** dialog box appears. Alternatively, click **Close** to cancel the import.
 - 4 In the **Open** dialog box, navigate to the certificate file you want to import, and then click **Open**.

If the import is successful, the **Information** notification at the top of the **Tools > Lattus Certificates** page displays, as an example, "Certificate `certificate_name.com.pem` uploaded successfully."

Convert...

Click **Convert...** to convert a file to the `.pem` format. You **must** convert a file, if you upload a file that is not already in the `.pem` format. Quantum only supports the `.pem` format.

Notes and Considerations

- If a file with the same name exists, you cannot convert the file to the `.pem` format. Delete the existing file first.
- If the file can be converted, the interface will attempt to convert it to the `.pem` format. The standard extension is `.pem`.
- The PEM format can contain private keys (RSA or DSA), public keys (RSA or DSA) and X.509 certificates. It is the default format for OpenSSL, and stores the data in either ASN.1 or DER format, surrounded by ASCII headers. Therefore, it is suitable for sending files as text, between systems.
- A file can contain multiple certificates.

- Below is a complete listing of files that can be converted:
 - a **PKCS7**: This is the Cryptographic Message Syntax Standard. A file can contain multiple certificates. Optionally they can be hashed. Optionally a certificate can be accompanied by a private key. As well as the original PKCS #7, there are three revisions: a, b, and c. The standard extensions for these four versions are .spc, .7m, .p7s, .p7a, .p7c, .p7b, and .p7z respectively.
 - b **DER**: This format can contain private keys (RSA or DSA), public keys (RSA or DSA) and X.509 certificates. It is the default format for most browsers. A file can contain only one certificate. Optionally, the certificate can be encrypted. The standard extension is .cer, but might be .der or .crt in some installations. If any of these file formats are actually ASCII base65 PEM files, the conversion will fail.
- Below are formats that **cannot** be converted to .pem:
 - a **PKCS12**: This format can contain private keys (RSA or DSA), public keys (RSA or DSA) and X.509 certificates. It stores them in a binary format. The standard extension is .pfx or .p12.
- 1 On the **Tools** menu, click **Lattus Certificates**. The **Tools > Lattus Certificates** page appears (see [Figure 102](#)).
- 2 On the **Tools > Lattus Certificates** page, click the option button to the left of a Lattus certificate to select it, and then click **Convert...**. The **Convert Certificate** dialog box appears.
- 3 In the **Convert Certificate** dialog box, click **Yes** to convert the file, or **No** to cancel the conversion process and return to the **Tools > Lattus Certificates** page.

If the conversion is successful, the .pem file appears in the **Lattus Certificates** table.

Download

Click **Download** to download and save a file listed on the **Tools > Lattus Certificates** page. This feature allows you to conveniently backup any certificate listed on the **Tools > Lattus Certificates** page.

Notes and Considerations

- You can download any file listed on the **Tools > Lattus Certificates** page.

- If you download a file created using the [New...](#) procedure, both the public and private certificate files are downloaded as one file.

- 1 On the **Tools** menu, click **Lattus Certificates**. The **Tools > Lattus Certificates** page appears (see [Figure 102](#)).
- 2 On the **Tools > Lattus Certificates** page, click the option button to the left of a Lattus certificate to select it, and then click **Download**. The **Download Private/Public Key Pair** dialog box appears.
- 3 In the **Download Private/Public Key Pair** dialog box, click the file link to begin the download.

If the download is successful, the .pem file appears in your local download directory.

- 4 In the **Download Private/Public Key Pair** dialog box, click **Done** to return to the **Tools > Lattus Certificates** page.

Delete

Click **Delete** to remove a file listed on the **Tools > Lattus Certificates** page.

Notes and Considerations

- You can delete any file listed on the **Tools > Lattus Certificates** page.
 - If an error occurs, you are provided with an error message.
 - After the file is deleted, the file is backed up to `/usr/cvfs/config_history/ssl`, with the same filename as the original, in addition to the standard time stamp `yyyyMMddHHmmss`.
- 1 On the **Tools** menu, click **Lattus Certificates**. The **Tools > Lattus Certificates** page appears (see [Figure 102](#)).
 - 2 On the **Tools > Lattus Certificates** page, click the option button to the left of a Lattus certificate to select it, and then click **Delete**. The **Delete Private/Public Certificate(s)** dialog box appears.
 - 3 In the **Delete Private/Public Certificate(s)** dialog box, click the button next to the appropriate file, and then click **Yes** to delete the file, or click **No** to return to the **Tools > Lattus Certificates** page.

If the file is deleted successfully, the **Information** notification at the top of the **Tools > Lattus Certificates** page displays, as an example, "File

backed up to `{/usr/cvfs/config_history/ssl/accounts.google.der.20130213155919}`.”

Refresh

Click **Refresh** to manually update the file list in the **Lattus Certificates** table.

Notes and Considerations

- The **Refresh** feature scans the `/usr/cvfs/config/ssl` directory, and adds any public certificates found within the directory to the **Lattus Certificates** table.
 - The **Refresh** feature works independently of the user interface. If an administrator using the command line interface, manually creates, updates, or deletes any of the certificates found in `/usr/cvfs/config/ssl`, the certificates are automatically updated on the **Lattus Certificates** table.
 - If an invalid certificate is manually placed in the list using the command line interface, an error message is displayed until the invalid file is removed. Until you remove the invalid file by manually removing the invalid certificate, other certificates are not displayed.
- 1 On the **Tools** menu, click **Lattus Certificates**. The **Tools > Lattus Certificates** page appears (see [Figure 102](#)).
 - 2 On the **Tools > Lattus Certificates** page, click **Refresh**.

HTTPS Configuration

You must have the following binary files installed for proper functionality and use of this feature:

Binary File	Description
<code>objectstorage.openssl.binary</code>	If the <code>objectstorage.openssl.binary</code> file is not installed, the New... button is disabled, and you will receive an error message.

Binary File	Description
objectstorage.c_rehash.binary	If the objectstorage.c_rehash.binary file is not installed, both the New... and Import... buttons are disabled, and you will receive an error message.

For the installation procedure and configuration of the binary files, see the *StorNext Installation Guide*.

Creating a Single Public/Private SSL Certificate for use on Both a StorNext MDC and a Lattus System

If you are working on a Lattus system which does not have an existing SSL certificate, this section outlines what you need to do to use both the private and public portions of the SSL certificate.

This section discusses how to use the PEM (Privacy Enhanced Mail) file that you create using the StorNext GUI. A typical PEM file will look like the server.pem file referenced in [Appendix A, Basic Secure Sockets Layer \(SSL\) Guidelines](#).

See [Appendix A, Basic Secure Sockets Layer \(SSL\) Guidelines](#), as it outlines some standard information about using private and public certificates.

To create a private and public SSL certificate for use on a Lattus system and a StorNext MDC, perform the following procedure:

- 1 On the **Tools** menu, click **Lattus Certificates**. The **Tools > Lattus Certificates** page appears.
- 2 On the **Tools > Lattus Certificates** page, click **New...** The **Tools > Lattus Certificates > New** page appears.
- 3 In the various text boxes, input the appropriate certificate data. The table in the [New...](#) section describes the various text boxes on the **Tools > Lattus Certificates > New** page.

Note: Text box fields on the **Tools > Lattus Certificates > New** page, designated with an asterisk (*) are required.

- a For the purposes of Lattus, do **NOT** enter a password in the **Password** field.

- b** In the **Subject Alternative Name** field, input the DNS and IP entries of all the servers for the certificate to work for.

For example:

```
dns=ibis1-controller1, dns=ibis1-  
controller1.mycompany.com, ip=192.168.166.94,  
ip=192.168.166.97, ip=192.168.10.3,  
ip=192.168.20.3
```

- 4** Click **Apply** to submit your inputs and create a private and public SSL certificate for use on a Lattus system and a StorNext MDC, or click **Cancel** to reset the form, and return to the **Tools > Lattus Certificates** page. If the submission is successful, your newly created private and public SSL certificate for use on a Lattus system and a StorNext MDC appears on the **Tools > Lattus Certificates** page.
- 5** To obtain the private and public SSL certificate to be used on the Lattus system, select the server .pem file and click **Download**. In the **Download Private/Public Key Pair** dialog box, click the file for “Click the Private Self-Signed Certificate file link to begin the download” and save the file where the Lattus CMC can access it.
- 6** Verify the Lattus system is working with your server .pem file.
- 7 (Optional)** Delete the server .pem file from the StorNext MDC, as it is no longer needed by the MDC.
 - a** On the **Tools > Lattus Certificates** page, click the option button to the left of the server .pem certificate to select it, and then click **Delete**. The **Delete Private/Public Certificate(s)** dialog box appears.
 - b** In the **Delete Private/Public Certificate(s)** dialog box, click “Check this to delete the Private Self-Signed Certificate file.”, and then click **Yes** to delete the file, or click **No** to return to the **Tools > Lattus Certificates** page.

Lattus Certificates

The Tools menu's **Lattus Certificates** option enables you to manage, and perform various actions to the public and private certificates that various applications requiring SSL authentication use. To access the **Tools > Lattus Certificates** page, on the **Tools** menu, click **Lattus Certificates** (see [Figure 102](#)).

For configuration details, see [HTTPS Configuration](#) on page 281.

If you are working on a Lattus system which does not have an existing SSL certificate, see [Creating a Single Public/Private SSL Certificate for use on Both a StorNext MDC and a Lattus System](#) on page 282.

HTTPS Default CA ROOT Certificate File or Path

Starting with StorNext 5 release 5.2, `/usr/cvfs/config/ssl` is no longer the default repository referenced by Storage Manager for SSL certificates when using HTTPS. The default certificate file or repository will depend on the OS vendor:

- **Debian:** `/etc/ssl/certs/ca-certificates.crt`
- **Red Hat:** `/etc/pki/tls/certs/ca-bundle.crt` or `/etc/ssl/certs/ca-bundle.crt`
- **SUSE:** `/etc/ssl/certs/`

Considerations for SUSE Platforms

If you are using `/usr/cvfs/config/ssl` as your certificate repository, you will have a conflict with the default root certificate repository `/etc/ssl/certs`. You have two options (below):

Option 1

Use `/usr/cvfs/config/ssl` as your default certificate repository; set the `FS_OBJSTORAGE_CAPATH=/usr/cvfs/config/ssl` in the `/usr/adic/TSM/config/fs_sysparm_override` file, then copy all the root certificates from `/etc/ssl/certs` to `/usr/cvfs/config/ssl`.

Note: Be sure to execute **c_rehash** on the directory that will be used as your default certificate repository afterward.

Option 2

Do not use `/usr/cvfs/config/ssl`; instead use the default root certificate repository by copying your certificates from `/usr/cvfs/config/ssl` to `/etc/ssl/certs`.

Note: Be sure to execute **c_rehash** on the directory that will be used as your default certificate repository afterward.

Considerations for Red Hat Platforms

If you are using `/usr/cvfs/config/ssl` as your certificate repository, you will not have a conflict with the default root certificate file. You will have to set `FS_OBJSTORAGE_CAPATH=/usr/cvfs/config/ssl` in the `/usr/adic/TSM/config/fs_sysparm_override` file.

How to Update Expired CA Root Certificates

Root Certificates may expire. When they do, you can update all your Root Certificates to the latest available from <http://rpmfind.net/linux/rpm2html/search.php?query=ca-certificates>. Select the one that fits your system.

- 1 Determine the default configured CA Root Certificate configured for StorNext using **libcurl**:

```
# curl-config --ca  
/etc/pki/tls/certs/ca-bundle.crt
```

- 2 Download the RPM that matches your system. In this example, we downloaded `ca-certificates-2014.1.98-65.1.el6.noarch.rpm`.
- 3 View the contents of the RPM.

```
# rpm -q -filesbypkg -p ca-certificates-2014.1.98-65.1.el6.noarch.rpm
ca-certificates          /etc/pki/ca-trust
ca-certificates          /etc/pki/ca-trust/README
.. snip ..
ca-certificates          /etc/pki/tls
ca-certificates          /etc/pki/tls/cert.pem
ca-certificates          /etc/pki/tls/certs
ca-certificates          /etc/pki/tls/certs/ca-
bundle.crt
ca-certificates          /etc/pki/tls/certs/ca-
bundle.trust.crt
.. snip ..
ca-certificates
/usr/share/pki/ca-trust-source/ca-
bundle.supplement.p11-kit
ca-certificates          /usr/share/pki/ca-trust-
source/ca-bundle.trust.crt
```

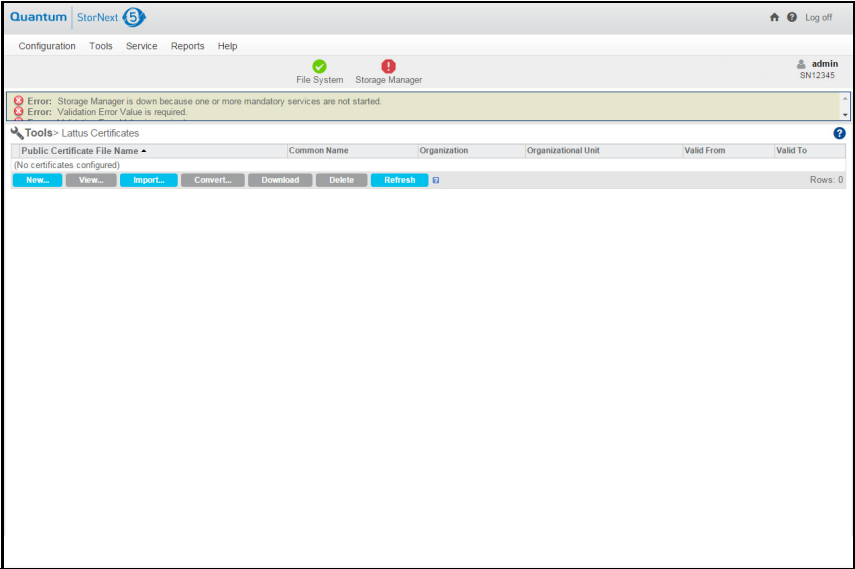
4 Install `/etc/pki/tls/certs/ca-bundle.crt`.

```
# mv /etc/pki/tls/certs/ca-bundle.crt etc/pki/tls/
certs/ca-bundle.crt.bak
# rpm2cpio ca-certificates-2014.1.98-65.1.el6.noarch.rpm | cpio -ivd
/etc/pki/tls/certs/ca-bundle.crt
```

- 5 Install the complete latest RPM.
- 6 Backup any files that you do not want replaced. This step may require you to install required dependencies.

```
# rpm -hiv ca-certificates-2014.1.98-65.1.el6.noarch.rpm  
  
p11-kit >= 0.18.4-2 is needed by ca-certificates-2014.1.98-65.1.el6.noarch  
  
p11-kit-trust >= 0.18.4-2 is needed by ca-certificates-2014.1.98-65.1.el6.noarch
```

Figure 105 Lattus Certificates Page



The table below provides the information displayed for each certificate, on the **Tools > Lattus Certificates** page:

Heading	Description	Examples
Public Certificate File Name	<p>The Privacy Enhanced Mail (PEM) filename and its respective filename extension (for example, .pem, or .der), of the Lattus certificate.</p> <ul style="list-style-type: none"> • There will always be a public certificate, but in some instances where you create a private certificate using this feature, the name will be both the name of the public and private certificate. • There will always be a public file name, and a private if you use this feature to generate your certificates. 	<p>accounts.mycompany.pem accounts.mycompany.der</p>
Common Name	<p>The fully qualified domain name (FQDN) of your server. This must match exactly what you type in your web browser or you will receive a name mismatch error.</p>	<p>*.mycompany.com controller.mycompany.com</p>
Organization	<p>The legal name of your organization. This should not be abbreviated and should include suffixes such as Inc, Corp, or LLC.</p>	<p>Mycompany Corp</p>
Organizational Unit	<p>The division of your organization handling the certificate.</p>	<p>Information Technology IT Department</p>
Valid From	<p>The date the certificate is valid from, in the form of yyyy-mm-dd hh:mm:ss time zone.</p>	<p>1998-08-22 11:41:51 MST</p>
Valid To	<p>The date the certificate is valid to, in the form of yyyy-mm-dd hh:mm:ss time zone.</p>	<p>1998-08-22 11:41:51 MST</p>

To create a **new** certificate, or **import** a certificate, click one of the following buttons on the **Tools > Lattus Certificates** page:

- [New...](#)

- [Import...](#)

Otherwise, select an existing certificate for which you want to perform the action, and then click one of the following buttons:

- [View...](#)
- [Convert...](#)
- [Download](#)
- [Delete](#)

To **refresh** the Lattus Certificates page, click the following button:

- [Refresh](#)

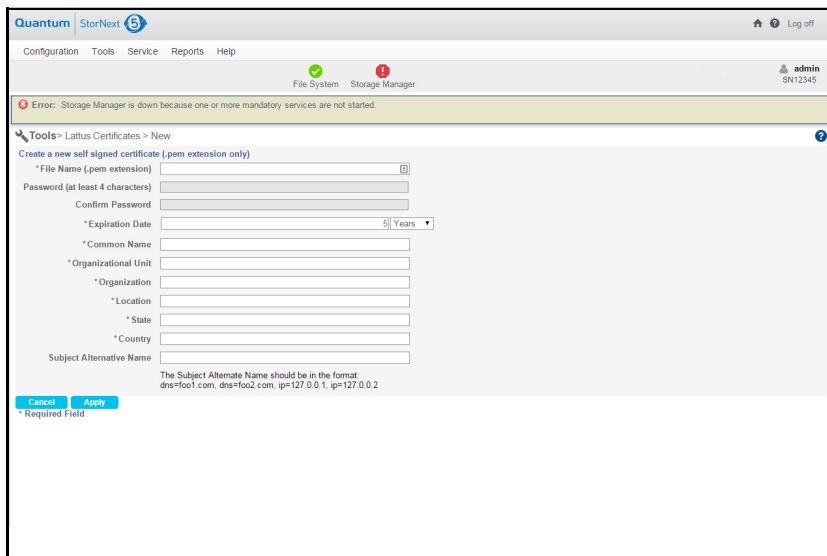
New...

Click **New...** to create a new self-signed certificate.

Note: The filename extension / format of the self-signed Lattus certificate must be `.pem`. You cannot create a self-signed Lattus certificate with filename extension or format other than `.pem`.

- 1 On the **Tools** menu, click **Lattus Certificates**. The **Tools > Lattus Certificates** page appears (see [Figure 102](#)).
- 2 On the **Tools > Lattus Certificates** page, click **New...** The **Tools > Lattus Certificates > New** page appears (see [Figure 103](#)).

Figure 106 Creating a New self-signed Lattus Certificate



3 In the various text boxes, input the appropriate certificate data. The table below describes the various text boxes on the **Tools > Lattus Certificates > New** page:

Note: Text box fields on the **Tools > Lattus Certificates > New** page, designated with an asterisk (*) are required.

Text Box	Description	Examples
File Name (.pem extension)	The Privacy Enhanced Mail (PEM) filename and its respective filename extension of the self-signed Lattus certificate. Note: Adding a certificate with the same name generates an error, instructing you to delete the certificate with that name first.	accounts.mycompany.pem

Text Box	Description	Examples
Password (at least 4 characters)	The Password input is an optional field. If a Password is entered, the input mimics the OpenSSL command password requirements as follows: <ul style="list-style-type: none"> • The Password input, and the Confirm Password input must match. • The Password input must be at least 4 characters, and can be all empty spaces or contain spaces. 	mypassword1234
Confirm Password	See the requirements for the Password input.	mypassword1234
Expiration Date	The Expiration Date low value is at least 1 day in the future. You can input a numeric value, and then select the unit of measurement from the drop-down list. The available unit of measurements are Years, Months, and Days . Note: There is no limit on the high end; however, if you input a value that is out of bounds for OpenSSL, then the OpenSSL command will generate an error.	5 Years
Common Name	The fully qualified domain name (FQDN) of your server. This must match exactly what you type in your web browser or you will receive a name mismatch error.	*.mycompany.com controller.mycompany.com
Organizational Unit	The division of your organization handling the certificate.	Information Technology IT Department
Organization	The legal name of your organization. This should not be abbreviated and should include suffixes such as Inc, Corp, or LLC.	Mycompany Corp

Text Box	Description	Examples
Location	The city where your organization is located.	Englewood
State	The state where your organization is located. This should not be abbreviated.	Colorado
Country	The two-letter ISO code for the country where your organization is located.	US
Subject Alternative Name	The Subject Alternative Name is an optional field. If entered, it should be in the following format (also specified under the text box): dns=foo1.com, dns=foo2.com, ip=127.0.0.1, ip=127.0.0.2	dns=foo1.com, dns=foo2.com, ip=127.0.0.1, ip=127.0.0.2

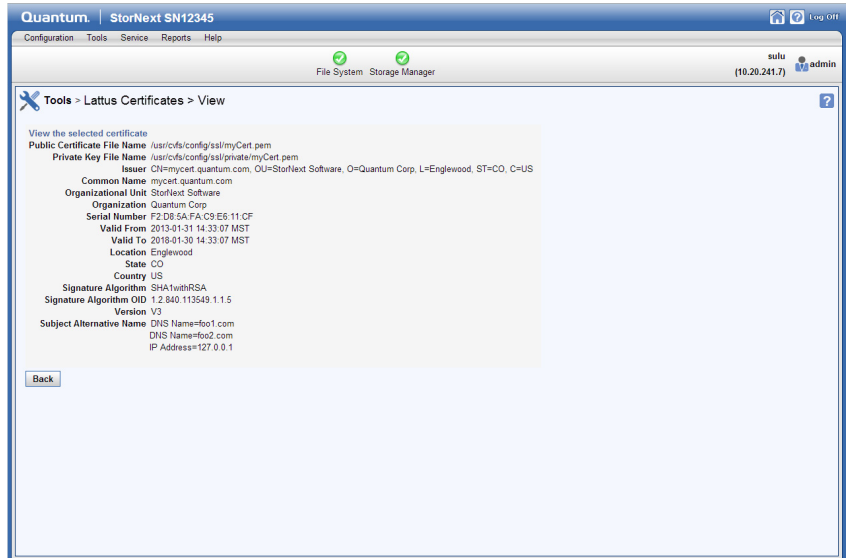
- 4 Click **Apply** to submit your inputs and create a new self-signed Lattus Certificate, or click **Cancel** to reset the form, and return to the **Tools > Lattus Certificates** page. If the submission is successful, your newly created self-signed Lattus Certificate appears on the **Tools > Lattus Certificates** page.

View...

Click **View...** to display the details of a specified Lattus certificate.

- 1 On the **Tools** menu, click **Lattus Certificates**. The **Tools > Lattus Certificates** page appears (see [Figure 102](#)).
- 2 On the **Tools > Lattus Certificates** page, click the option button to the left of a Lattus certificate to select it, and then click **View...** The **Tools > Lattus Certificates > View** page appears (see [Figure 104](#)).

Figure 107 Viewing a Lattus Certificate



The table below describes the various fields on the **Tools > Lattus Certificates > View** page:

Name	Description	Examples
Public Certificate File Name	The Privacy Enhanced Mail (PEM) filename and its respective filename extension (for example, .pem, or .der), of the Lattus certificate.	/usr/cvfs/config/ssl/myCert.pem
Private Key File Name	The filename of the private key in the Lattus certificate. Note: If the certificate was not created through this feature, you will receive following text (in red/bold): Certificates that were imported do not have Private Keys associated to them.	/usr/adic/gui/.ssl/myCert.pem

Name	Description	Examples
Issuer	This property contains the name of the certificate authority (CA) that issued the certificate. The distinguished name for the certificate is a textual representation of the certificate subject or issuer.	CN=mycert.mycompany.com, OU=StorNext Software, O=Mycompany Corp, L=Englewood, ST=CO, C=US
Common Name	The fully qualified domain name (FQDN) of your server. This must match exactly what you type in your web browser or you will receive a name mismatch error.	mycert.mycompany.com
Organizational Unit	The division of your organization handling the certificate.	Information Technology IT Department
Organization	The legal name of your organization. This should not be abbreviated and should include suffixes such as Inc, Corp, or LLC.	Mycompany Corp
Serial Number	The serial number of the selected certificate.	F2:D8:5A:FA:C9:E6:11:CF
Valid From	The date the certificate is valid from , in the form of yyyy-mm-dd hh:mm:ss time zone.	2013-01-31 14:33:07 MST
Valid To	The date the certificate is valid to , in the form of yyyy-mm-dd hh:mm:ss time zone.	2018-01-30 14:33:07 MST
Location	The city where your organization is located.	Englewood
State	The state where your organization is located. This should not be abbreviated.	Colorado
Country	The two-letter ISO code for the country where your organization is located.	US

Name	Description	Examples
Signature Algorithm	The algorithm used to create the signature of the certificate.	SHA1withRSA
Signature Algorithm OID	The object identifier (OID) identifies the type of signature algorithm used by the certificate.	1.2.840.113549.1.1.5
Version	The version number of the certificate.	V3
Subject Alternative Name	The Subject Alternative Name is the name of the user of the certificate. The alternative name for the certificate is a textual representation of the subject or issuer of the certificate.	DNS Name=foo1.com DNS Name=foo2.com IP Address=127.0.0.1

3 Click **Back** to return to the **Tools > Lattus Certificates** page.

Import...

Click **Import...** to import a certificate.

Notes and Considerations

- Files that do not have a .pem extension will need to be converted to .pem for use in SSL communication. See [Convert...](#) on page 278 to convert a file to the .pem format. Quantum only supports the .pem format for certificates.
- You can import one file, which contains multiple public keys. Doing so will create individual rows for each key file with the filename `_multiple.pem`. If any of the multiple keys is deleted, since they comprise the same file, the entire certificate is deleted, and all of the public keys are no longer persisted.
- You can view a certificate on an individual basis by selecting the certificate to view.
- You can import any type of valid public key file, as long as the certificate is not expired. If the certificate is expired, the import will fail, and you will be notified via an **Error** notification. If you import a file with multiple public keys, and any of the public keys in the file are expired, then the entire file is rejected.

- Empty files and files exceeding 10 MB are not permitted. If you want to change the 10 MB limit, you must manually edit the `/usr/adic/gui/config/component.properties` file, and modify the following value: `objectstorage.ssl.maxCertSizeMb=10`
 - You cannot upload a private certificate file; however, you can create a private certificate. If your private / public key is in a `.pem` file, open the file in a text editor and remove the private key.
- 1 On the **Tools** menu, click **Lattus Certificates**. The **Tools > Lattus Certificates** page appears (see [Figure 102](#)).
 - 2 On the **Tools > Lattus Certificates** page, click **Import...** The **Import A Certificate** dialog box appears.
 - 3 In the **Import A Certificate** dialog box, click **Choose File** to select a file to import. The **Open** dialog box appears. Alternatively, click **Close** to cancel the import.
 - 4 In the **Open** dialog box, navigate to the certificate file you want to import, and then click **Open**.

If the import is successful, the **Information** notification at the top of the **Tools > Lattus Certificates** page displays, as an example, "Certificate `certificate_name.com.pem` uploaded successfully."

Convert...

Click **Convert...** to convert a file to the `.pem` format. You **must** convert a file, if you upload a file that is not already in the `.pem` format. Quantum only supports the `.pem` format.

Notes and Considerations

- If a file with the same name exists, you cannot convert the file to the `.pem` format. Delete the existing file first.
- If the file can be converted, the interface will attempt to convert it to the `.pem` format. The standard extension is `.pem`.
- The PEM format can contain private keys (RSA or DSA), public keys (RSA or DSA) and X.509 certificates. It is the default format for OpenSSL, and stores the data in either ASN.1 or DER format, surrounded by ASCII headers. Therefore, it is suitable for sending files as text, between systems.
- A file can contain multiple certificates.

- Below is a complete listing of files that can be converted:
 - a **PKCS7**: This is the Cryptographic Message Syntax Standard. A file can contain multiple certificates. Optionally they can be hashed. Optionally a certificate can be accompanied by a private key. As well as the original PKCS #7, there are three revisions: a, b, and c. The standard extensions for these four versions are .spc, .7m, .p7s, .p7a, .p7c, .p7b, and .p7z respectively.
 - b **DER**: This format can contain private keys (RSA or DSA), public keys (RSA or DSA) and X.509 certificates. It is the default format for most browsers. A file can contain only one certificate. Optionally, the certificate can be encrypted. The standard extension is .cer, but might be .der or .crt in some installations. If any of these file formats are actually ASCII base65 PEM files, the conversion will fail.
- Below are formats that **cannot** be converted to .pem:
 - a **PKCS12**: This format can contain private keys (RSA or DSA), public keys (RSA or DSA) and X.509 certificates. It stores them in a binary format. The standard extension is .pfx or .p12.
- 1 On the **Tools** menu, click **Lattus Certificates**. The **Tools > Lattus Certificates** page appears (see [Figure 102](#)).
- 2 On the **Tools > Lattus Certificates** page, click the option button to the left of a Lattus certificate to select it, and then click **Convert...**. The **Convert Certificate** dialog box appears.
- 3 In the **Convert Certificate** dialog box, click **Yes** to convert the file, or **No** to cancel the conversion process and return to the **Tools > Lattus Certificates** page.

If the conversion is successful, the .pem file appears in the **Lattus Certificates** table.

Download

Click **Download** to download and save a file listed on the **Tools > Lattus Certificates** page. This feature allows you to conveniently backup any certificate listed on the **Tools > Lattus Certificates** page.

Notes and Considerations

- You can download any file listed on the **Tools > Lattus Certificates** page.

- If you download a file created using the [New...](#) procedure, both the public and private certificate files are downloaded as one file.

- 1 On the **Tools** menu, click **Lattus Certificates**. The **Tools > Lattus Certificates** page appears (see [Figure 102](#)).
- 2 On the **Tools > Lattus Certificates** page, click the option button to the left of a Lattus certificate to select it, and then click **Download**. The **Download Private/Public Key Pair** dialog box appears.
- 3 In the **Download Private/Public Key Pair** dialog box, click the file link to begin the download.

If the download is successful, the .pem file appears in your local download directory.

- 4 In the **Download Private/Public Key Pair** dialog box, click **Done** to return to the **Tools > Lattus Certificates** page.

Delete

Click **Delete** to remove a file listed on the **Tools > Lattus Certificates** page.

Notes and Considerations

- You can delete any file listed on the **Tools > Lattus Certificates** page.
 - If an error occurs, you are provided with an error message.
 - After the file is deleted, the file is backed up to `/usr/cvfs/config_history/ssl`, with the same filename as the original, in addition to the standard time stamp `yyyyMMddHHmss`.
- 1 On the **Tools** menu, click **Lattus Certificates**. The **Tools > Lattus Certificates** page appears (see [Figure 102](#)).
 - 2 On the **Tools > Lattus Certificates** page, click the option button to the left of a Lattus certificate to select it, and then click **Delete**. The **Delete Private/Public Certificate(s)** dialog box appears.
 - 3 In the **Delete Private/Public Certificate(s)** dialog box, click the button next to the appropriate file, and then click **Yes** to delete the file, or click **No** to return to the **Tools > Lattus Certificates** page.

If the file is deleted successfully, the **Information** notification at the top of the **Tools > Lattus Certificates** page displays, as an example, "File

backed up to `{/usr/cvfs/config_history/ssl/accounts.google.der.20130213155919}`.”

Refresh

Click **Refresh** to manually update the file list in the **Lattus Certificates** table.

Notes and Considerations

- The **Refresh** feature scans the `/usr/cvfs/config/ssl` directory, and adds any public certificates found within the directory to the **Lattus Certificates** table.
 - The **Refresh** feature works independently of the user interface. If an administrator using the command line interface, manually creates, updates, or deletes any of the certificates found in `/usr/cvfs/config/ssl`, the certificates are automatically updated on the **Lattus Certificates** table.
 - If an invalid certificate is manually placed in the list using the command line interface, an error message is displayed until the invalid file is removed. Until you remove the invalid file by manually removing the invalid certificate, other certificates are not displayed.
- 1 On the **Tools** menu, click **Lattus Certificates**. The **Tools > Lattus Certificates** page appears (see [Figure 102](#)).
 - 2 On the **Tools > Lattus Certificates** page, click **Refresh**.

HTTPS Configuration

You must have the following binary files installed for proper functionality and use of this feature:

Binary File	Description
<code>objectstorage.openssl.binary</code>	If the <code>objectstorage.openssl.binary</code> file is not installed, the New... button is disabled, and you will receive an error message.

Binary File	Description
objectstorage.c_rehash.binary	If the objectstorage.c_rehash.binary file is not installed, both the New... and Import... buttons are disabled, and you will receive an error message.

For the installation procedure and configuration of the binary files, see the *StorNext Installation Guide*.

Creating a Single Public/Private SSL Certificate for use on Both a StorNext MDC and a Lattus System

If you are working on a Lattus system which does not have an existing SSL certificate, this section outlines what you need to do to use both the private and public portions of the SSL certificate.

This section discusses how to use the PEM (Privacy Enhanced Mail) file that you create using the StorNext GUI. A typical PEM file will look like the server.pem file referenced in [Appendix A, Basic Secure Sockets Layer \(SSL\) Guidelines](#).

See [Appendix A, Basic Secure Sockets Layer \(SSL\) Guidelines](#), as it outlines some standard information about using private and public certificates.

To create a private and public SSL certificate for use on a Lattus system and a StorNext MDC, perform the following procedure:

- 1 On the **Tools** menu, click **Lattus Certificates**. The **Tools > Lattus Certificates** page appears.
- 2 On the **Tools > Lattus Certificates** page, click **New...** The **Tools > Lattus Certificates > New** page appears.
- 3 In the various text boxes, input the appropriate certificate data. The table in the [New...](#) section describes the various text boxes on the **Tools > Lattus Certificates > New** page.

Note: Text box fields on the **Tools > Lattus Certificates > New** page, designated with an asterisk (*) are required.

- a For the purposes of Lattus, do **NOT** enter a password in the **Password** field.

- b In the **Subject Alternative Name** field, input the DNS and IP entries of all the servers for the certificate to work for.

For example:

```
dns=ibis1-controller1, dns=ibis1-  
controller1.mycompany.com, ip=192.168.166.94,  
ip=192.168.166.97, ip=192.168.10.3,  
ip=192.168.20.3
```

- 4 Click **Apply** to submit your inputs and create a private and public SSL certificate for use on a Lattus system and a StorNext MDC, or click **Cancel** to reset the form, and return to the **Tools > Lattus Certificates** page. If the submission is successful, your newly created private and public SSL certificate for use on a Lattus system and a StorNext MDC appears on the **Tools > Lattus Certificates** page.
- 5 To obtain the private and public SSL certificate to be used on the Lattus system, select the server .pem file and click **Download**. In the **Download Private/Public Key Pair** dialog box, click the file for "Click the Private Self-Signed Certificate file link to begin the download" and save the file where the Lattus CMC can access it.
- 6 Verify the Lattus system is working with your server .pem file.
- 7 **(Optional)** Delete the server .pem file from the StorNext MDC, as it is no longer needed by the MDC.
 - a On the **Tools > Lattus Certificates** page, click the option button to the left of the server .pem certificate to select it, and then click **Delete**. The **Delete Private/Public Certificate(s)** dialog box appears.
 - b In the **Delete Private/Public Certificate(s)** dialog box, click "Check this to delete the Private Self-Signed Certificate file.", and then click **Yes** to delete the file, or click **No** to return to the **Tools > Lattus Certificates** page.

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** page, choose **File and Directory Actions** from the **Tools** menu. (See [Figure 108](#).)

The following information is displayed for the available files:

- **Name:** The name of the file
- **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 page is a drop-down 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](#)
- [Assign Affinities](#)

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** page appears.

Figure 108 File and Directory
Action Page



- 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.

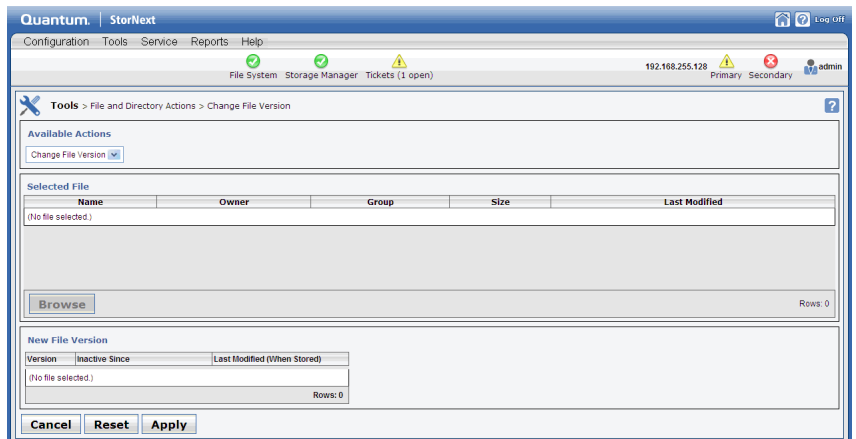
When the confirmation message appears, click **Yes** to mount the store the file, or **No** to abort.

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** page appears. (See [Figure 108.](#))
- 2 Choose **Change File Version** from the **Available Actions** drop-down list.

Figure 109 Change File Version Page



- 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** page appears. (See [Figure 108.](#))
- 2 Choose **Recover Files** from the **Available Actions** drop-down list.
- 3 Click **Browse**. The **Tools > File and Directory Actions > Recover Files > Browse** page appears.

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

Figure 110 Recover Files
Browse Page



- 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** page, 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.

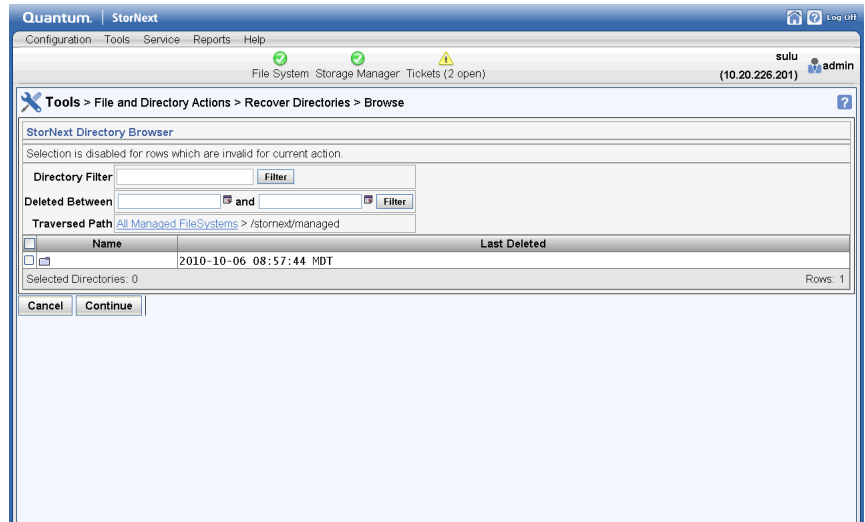
Recover Directories

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** page appears. (See [Figure 108](#).)
- 2 Choose **Recover Directories** from the **Available Actions** drop-down list.
- 3 Click **Browse**. The **Tools > File and Directory Actions > Recover Directories > Browse** page appears.

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

Figure 111 Recover Directories
Browse Page



- 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** page, 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** page appears. (See [Figure 108](#).)
- 2 Choose **Retrieve Files** from the **Available Actions** drop-down list.

Figure 112 Retrieve Files



- 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** page appears. (See [Figure 108.](#))
- 2 Choose **Retrieve Directory** from the **Available Actions** drop-down list.

Figure 113 Retrieve Directory Page




- 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** page appears. (See [Figure 108.](#))
- 2 Choose **Truncate Files** from the **Available Actions** drop-down list.

Figure 114 Truncate Files Page



- 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** page appears. (See [Figure 108.](#))

2 Choose **Move Files** from the **Available Actions** drop-down list.

Figure 115 Move Files Page



- 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.

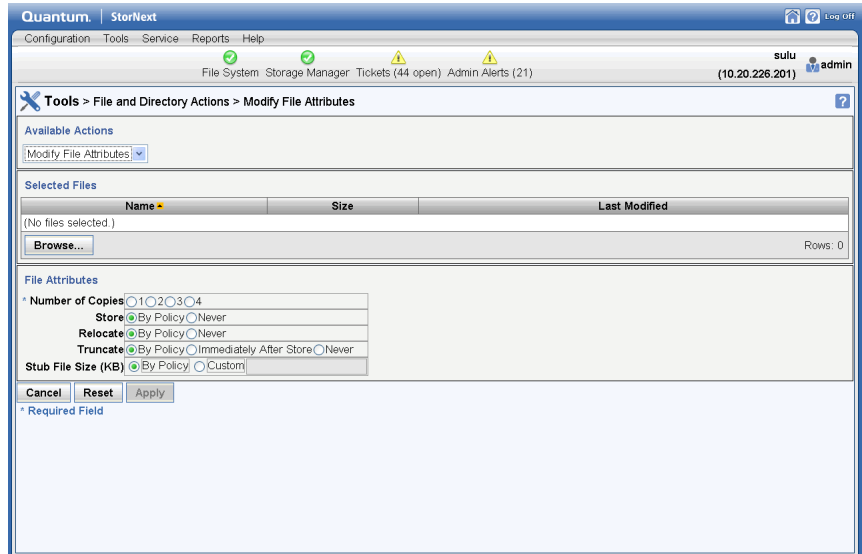
Modify File Attributes

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** page appears. (See [Figure 108](#).)

- 2 Choose **Modify File Attributes** from the **Available Actions** drop-down list.

Figure 116 Modify File Attributes Page




- 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** page appears. (See [Figure 108.](#))
- 2 Choose **View File Info** from the **Available Actions** drop-down list.

Figure 117 View File Info Page




- 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.

Assign Affinities

Choose this option to assign affinities to 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** page appears. (See [Figure 108.](#))
- 2 Choose **Assign Affinities** from the **Available Actions** drop-down list.

Figure 118 Assign Affinities
Page



- 3 Select the files you want to assign affinities to. If necessary, click **Browse** to navigate to the file location and then select the file.
- 4 Under the **Affinity Options** section, select the affinity you want to assign from the **Affinity to Assign** drop-down list.
- 5 Click **Apply**.
- 6 When the confirmation message appears, click **Yes** to proceed, or **No** to abort.

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 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 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
- **Manage Quotas:** Limit the amount of disk storage consumed on a per user, or per group basis across an entire file system, or within a designated directory hierarchy.

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](#)

Upgrade Firmware

The Firmware Upgrade option allows you to perform a firmware upgrade on StorNext M660, M440 and M330 Metadata Appliances and Pro Foundation. Upgrading the firmware also upgrades the StorNext software, if applicable.

Note: This menu option is only available on the StorNext M660, M440 and M330 Metadata Appliances. Use the StorNext GUI to perform all firmware upgrades and HA conversions.

Note: StorNext 5 Release 5.1.1 did not support StorNext M330 Metadata Appliance firmware upgrades.

Note: The StorNext M330 Metadata Appliance upgrade to StorNext 5 Release 5.2 and later is only permitted from StorNext 5 Release 5.1.

Upgrade Considerations

Before you begin the upgrade you should note the following considerations so you can plan accordingly:

Consider the following prior to upgrading:

- Not all StorNext releases may be upgraded to StorNext 5. As a result, an upgrade to the current version of StorNext may require multiple, incremental upgrades, depending on the version of StorNext currently installed.

For information about supported upgrade paths for StorNext 5, consult the *StorNext 5 Compatibility Guide*. If your system is running a StorNext release prior to the supported upgrade releases for StorNext 5, consult an earlier version of the *StorNext Compatibility Guide* that applies to your specific upgrade, and the dependencies for StorNext Clients in the environment.

- Firmware upgrade installation files must first be acquired from Quantum.

Obtain the Firmware Upgrade Files

To obtain the firmware files (both are required) you wish to install:

Note: The two files are large - around 2 GB total, so plan time to download the files for the upgrade.

- a Go to the CSWeb site and log in.
- b Navigate to the StorNext Products page for your appliance (on the lefthand side of the CSWeb site, look for the appropriate link under the StorNext Products section).

The first section of the page contains Downloads for the given appliance.

- c Scroll down to the Current Software section, and download both firmware image files.

- Firmware upgrade installation files, which contain the .fw suffix, must be uploaded to the system prior to beginning the upgrade process. Uploading the firmware upgrade files in a network with low latency should only take a matter of minutes. High network latency in your environment can slow the upload of these files onto the Metadata Appliance or Pro Foundation.
- When using the firmware upgrade process from the StorNext GUI, the license for the system will be automatically applied to the Metadata Appliance or Pro Foundation.

Upgrade Times

Some firmware upgrades apply a firmware update to both controllers in the metadata array(s), which will greatly increase the total upgrade time for the appliance.

This section explains two important pieces of information to help you through the upgrade process:

- [Firmware Upgrade Process and Time Estimates](#) — which provides an overview of the upgrade process, and provides guidelines on estimating the amount of time an upgrade will take, depending on the StorNext or StorNext 5 release you are upgrading from, and the StorNext or StorNext 5 release you are upgrading to.
- [Monitoring Upgrade Progress](#) on page 323 — which provides information on how to monitor the upgrade by viewing messages in various log files that are updated during the upgrade process.

See [StorNext Releases/Downtime Requirements](#) on page 321 to identify which upgrades include a firmware upgrade.

WARNING: In the StorNext GUI, do not deactivate the **Upgrade in Progress** status or start up services manually on either MDC node.

If StorNext services are manually started while the metadata array firmware is being upgraded, the MDC nodes will likely SMITH.

Firmware Upgrade Process and Time Estimates

There are several factors that affect the availability of the system and metadata operations during upgrades. StorNext appliance firmware upgrades (contained in the .fw file downloaded earlier in this process) always include an update to the StorNext software in the release, and can also includes updates to MDC node firmware and metadata array controller firmware. The way in which these updates are applied and the impacts to system availability vary, depending on which firmware upgrade being applied, as described here:

StorNext Releases/Downtime Requirements

Note: Metadata array operations and the StorNext file system are unavailable while metadata array controller firmware is installed. We refer to this suspension of operations as a **downtime upgrade**, since no array I/O may take place during this time.

Note: The upgrade of metadata array controller appliance firmware also, once installed, initiates a reboot of both MDC nodes prior to 5.1.1, requiring additional time. Plan upgrade times accordingly.

StorNext Releases requiring suspension of metadata operations during upgrades occur any time a newer release contains a newer version of metadata array controller firmware version than the version currently installed.

The following table contains examples of firmware upgrade versions. Upgrading StorNext Release with a newer firmware will be a **downtime upgrade**:

Table 2 StorNext Releases/
Firmware

StorNext Release*	Controller Array Firmware Version
5.3.x	08.20.09
5.2.2	08.20.09
5.2, 5.2.0.1, 5.2.0.2, and 5.2.1	08.10.13

StorNext Release*	Controller Array Firmware Version
5.1.1	08.10.13
5.1	07.84.46
5.0.1	07.84.46
5	07.84.46
4.7.1	07.84.46
4.7.0.1	07.84.46
4.7.0	07.84.46
4.6.1	07.80.55
4.6	07.80.55
4.3.3	07.80.55
4.3.2	07.80.55
4.3.1	07.80.55
4.3.0	07.80.55
4.2.2.0.1	07.75.17

* Note: Some StorNext releases are not supported upgrade paths to StorNext 5 releases. Please consult the StorNext 5 Compatibility Guides and/or older StorNext Compatibility Guides applicable for the StorNext Release you wish to upgrade to, in order to determine your particular upgrade path.

Example (based on the previous table): If you were upgrading the StorNext firmware from StorNext Release 4.7.1 to StorNext 5 Release 5.1.1, the array firmware would be upgraded from 07.84.46 to 08.10.13. This is a **downtime upgrade**.

Components and Upgrade Time Estimates

Component	Upgrade Time Estimate (approx. minimum)
MDC node	30 to 60 minutes per node*
Metadata Array	20 minutes**

* Reboot times could vary widely, depending on the size of the SAN in your environment, whether or not MDC node firmware needs to be upgraded, and the StorNext Release installed prior to the upgrade.

** Since metadata array operations, and the StorNext file system will be suspended during the array firmware upgrade, this is a **downtime upgrade**.

Note: See sections below for information on additional reboots that may be required for specific StorNext Releases.

Monitoring Upgrade Progress

You can monitor the upgrade process using the following log files:

Log File	Description
/var/log/DXi/upgrades/ NodeX.upgradeoutput.log	Upgrade output log for the main part of the upgrade (not including the NetApp FW updates).
/var/log/DXi/upgrades/ NodeX.upgrade_progress	Upgrade Progress Log for the main part of the upgrade (not including the NetApp firmware updates).
/var/log/DXi/baseos.log	Final upgrade progress at the end of the final boot of the second node (shows the NetApp firmware update has started).

Log File	Description
<code>/var/log/DXi/UpgradeArrays.log</code>	Contains the actual NetApp firmware upgrade progress. The firmware update command will not respond for 30-40 minutes so there is no specific change in progress noted until the controllers reboot.
<code>/var/log/messages</code>	Contains other boot messages like the application of firmware bundle updates and hardware detect script execution, both of which could take several minutes.

Each of these log files will continue to log upgrade status progress detail until the StorNext GUI is once again fully-accessible at the end of the upgrade.

Note: SNFS services are stopped prior to updating the controller firmware, so the StorNext GUI does not show the array(s) as available while the controller firmware is updated.

The end of the `NodeX.upgradeoutput.log` includes the following statement that indicates the array firmware update has started:

```
New NetApp FW needs to be activated. Bringing  
down the cluster.
```

The message below is added to the log after the array controllers reboot:

```
Successfully activated new NetApp FW on the  
array controllers.
```

These are the only messages provided on the upgrade progress.

Note: The array controllers will reboot after the firmware update is complete (about 20-30 minutes after the update begins). Messages about the SCSI driver and multi-path driver devices are removed and added are displayed during the upgrade.

These occurrences are normal since the controllers drop off the SAS BUS and are added again during the update. LUN access is unavailable during the controller reboot.

SNFS services will start up automatically after the upgrade is completed and all LUNs will be accessible on the metadata array.

For All Firmware Upgrades:

- Full access to all GUI features are not available during the upgrade. The GUI should not be started until the upgrade is complete.

For StorNext 5 and Later Upgrades:

Note: StorNext 5 Release 5.1.1 did not support StorNext M330 Metadata Appliance firmware upgrades.

- The upgrade is applied to each node in reverse order, beginning with the secondary node. When the secondary node has completed its upgrade, the system will failover from the primary to the secondary, and begin upgrading the primary system. While the node firmware upgrade is being applied, metadata operations are not interrupted. Access to a limited StorNext GUI is available on the MDC node acting as primary during the upgrade of the node acting as secondary.
 - Both MDC nodes will need to reboot after the metadata array firmware update is applied (if applicable).
-

Note: As of StorNext 5 Release 5.1.1, MDC node reboots after metadata array firmware updates are no longer required.

For Upgrades Prior to StorNext 5:

- Each StorNext firmware update requires a reboot of both nodes. In some cases, multiple reboots of the nodes are required.

- Each time the firmware update is done, the secondary node of the Metadata Appliance is left out of the HA configuration. As a result, you will need to convert the M660, M440 or M330 Metadata Appliance to an HA system after each upgrade in order to regain failover operations. After the HA conversion, both nodes will reboot, which can take 30 minutes or longer per node.
- The Storage Manager components will need to be restarted after the HA conversion is complete by clicking the **Start** button in the Storage Manager panel of the **Tools > System Control** page.
- All upgrades prior to StorNext 5 are **downtime upgrades**.


Upgrade Procedure

To upgrade the StorNext M660, M440 or M330 Metadata Appliance or Pro Foundation:

- 1 Download the required firmware file(s) from Quantum CSWeb for the StorNext release needed.
- 2 Log into the StorNext GUI.
- 3 Choose **Tools > Firmware Upgrade**.

The **Firmware Upgrade** page appears.

Figure 119 Firmware Upgrade Page



- 4 Do one of the following:
 - a Select **Auto Upload** to upload the file immediately after you select it.
 - b Do not select **Auto Upload**.

- 5 Click **Browse...**, and then navigate to the directory where the file resides. Firmware files are identifiable by the **.fw** extension.

Note: There are two **.fw** files required for updating firmware. The filenames are similar to QTM-DXISNA-upd-5.1.0.OS6-15147-15110.1of2.fw and QTM-DXISNA-upd-5.1.0.OS6-15147-15110.2of2.fw. Since it is a two-part upgrade, upload both files to the GUI. To begin activation of the upgrade, select either of the uploaded files and then click **Apply**. Both parts are applied to the system.

If you selected the **Auto Upload** option, the file is immediately uploaded. Proceed to [Step 7](#).

- 6 If you did not select the **Auto Upload** option and want to validate the file before uploading, click **Validate**. After a message informs you that the file is valid, click the **Upload** button located to the right of the **Browse...** button.

Note: Files are automatically validated after you click **Apply** ([Step 7](#)), but you won't receive a message telling you the file is valid.

- 7 Click **Apply** to begin the upgrade.

The green status indicator at the top of the page indicates upload progress, not the upgrade progress. To monitor upgrade progress, check the logs available under the Reports menu.

If Upgrading to StorNext 5 from StorNext 4.7.0:

The upgrade process is complete. There is no interruption of metadata operations during or after the upgrade.

If Upgrading to StorNext 5 from StorNext Releases Prior to StorNext 4.7.0:

After the upgrade to the primary MDC node completes, metadata operations will be interrupted for 30 minutes or more, and both MDC nodes will reboot, which could take an additional 30 minutes to complete, before you are able to log back in, so plan upgrade times accordingly.

If Upgrading to StorNext Releases Prior to StorNext 5:

After the upgrade to the primary MDC node completes, metadata operations will be interrupted for 30 minutes or more, and both

MDC nodes will reboot, which could take an additional 30 minutes to complete. It could be a long time before you are able to log back in, so plan upgrade times accordingly.

- 8 Convert the system to HA, according to [Converting to HA](#) on page 235. (This step is not necessary when upgrading to StorNext 5 releases).

GUI Feedback During Upgrades

There are some indications within the GUI that the system is being upgraded. Here are some notes about this visual feedback:

If Upgrading to StorNext 5 from StorNext 4.7.0 and later:

- For systems without a defined virtual IP, the secondary MDC nodes is upgraded first, followed by a fail over to the secondary MDC node, which takes the role of the primary MDC node. The GUI will run on the primary MDC node until the secondary MDC node completes its upgrade, and the system fails over to the secondary MDC node. At that point you will need to bring up the GUI for the secondary MDC node, while the primary MDC node completes its upgrade. At the end of updating the primary MDC node, the system will not automatically fail over again, it only fails over once.
- If a virtual IP is used, the GUI will need to be refreshed to display the GUI once the secondary MDC node upgrade and failover is complete. The assumes the role of the primary when the failover occurs. Once the original primary MDC node completes its upgrade, it will assume the role of the secondary MDC node.

If Upgrading to StorNext Releases Prior to StorNext 5:

- On the primary MDC node, the GUI will display different status messages throughout the installation, including messages that the system will reboot, and red icons indicating that the primary MDC node, secondary MDC node, File System, and Storage Manager are also disabled.
- The current user account will eventually time out and the GUI will stop functioning when the power to Metadata Appliance is removed during reboot. Status updates will cease and the GUI will not be fully-functional again until the system completes booting both MDC nodes.

Post-Upgrade Failover

If you desire to failover your system after the upgrade, see [Initiating a Graceful System Failover and Verification](#) on page 235.

Deleting Uploaded Files

Follow these steps to delete uploaded files you no longer need:

- 1 Log into the StorNext GUI.
- 2 Choose **Tools > Firmware Upgrade**.
- 3 The **Firmware Upgrade** page appears
- 4 Select from the list the file you want to delete, and then click **Delete**. (If you want to delete multiple files, you must delete them one at a time.)
- 5 When a confirmation appears, click **Yes** to proceed or **No** to abort.

When a message informs you that the file was successfully deleted, click **OK**.

StorNext 5 Release 5.1.1 Post-Upgrade Required - Critical Fix

Note: This section only applies for StorNext software upgraded to StorNext 5 Release 5.1.1. **DO NOT** perform this procedure if you are upgrading to a release other than StorNext 5 Release 5.1.1.

Note: Apply this fix only after the firmware upgrade has been applied. To upgrade the firmware, see [Upgrade Firmware](#).

WARNING: Always create a symbolic link after upgrading the firmware.

If you upgrade the appliance to StorNext 5 Release 5.1.1 and do not create the symbolic link, the root file system could fill completely or the CPU may run at 100%.

Determining if the Symbolic Link Exists

Before continuing to the next step, verify if the symbolic link to `/scratch/whisper` exists on **Node 2**:

Login to the Command Line

1 on .

stornextstornextstornext

Determine if the symbolic link exists

Perform the following steps on **Node 2**:

1 Enter the following command:

```
file /opt/graphite/storage/whisper
```

If the symbolic link exists, the following is displayed, and no further action is required:

```
/opt/graphite/storage/whisper:  
symbolic link to '/scratch/whisper'
```

2 If the symbolic link exists, disconnect from **Node 2**.
Otherwise, continue to next section.

Creating the Missing Symbolic Link

If the output above is not displayed, either the symbolic link to `/scratch/whisper` does not exist, or the `/opt/graphite/storage/whisper` directory does not exist. In this case, do the following:

1 On **Node 2**, enter the following to stop **carbon-cache**:

```
/opt/quantum/python27/bin/supervisorctl stop  
carbon-cache
```

The following is displayed:

```
carbon-cache: stopped
```

- 2 Create the symbolic link.

Note: If the `/opt/graphite/storage/whisper` directory exists, execute the following command to remove it:

```
rm -fr /opt/graphite/storage/whisper
```

- 3 Enter the following to make the `/scratch/whisper` directory:

```
mkdir -p /scratch/whisper
```

- 4 Enter the following to change ownership of the `/scratch/whisper` directory:

```
chown apache:apache /scratch/whisper
```

- 5 Enter the following to set permissions on the directory:

```
chmod 755 /scratch/whisper
```

- 6 Enter the following to create the symbolic link:

```
ln -s /scratch/whisper /opt/graphite/storage/whisper
```

- 7 Repeat [Step 2](#) through [Step 6](#) for the other node. (Open a separate SSH session for **Node 1**. See [Login to the Command Line](#) on page 330.).

- 8 On **Node 2**, enter the following to verify if the `/usr/cvfs/config/snstatd.conf` file is present:

```
file /usr/cvfs/config/snstatd.conf
```

If the file exists, the following is displayed:

```
/usr/cvfs/config/snstatd.conf: ASCII text
```

9 If the file does not exist, continue on to the next step.

If the file exists, enter the following to remove the file:

```
rm -f /usr/cvfs/config/snstatd.conf
```

10 Enter the following to determine if the **snstatd** daemon is running:

```
ps -C snstatd
```

If so, stop it by executing the following command:

```
kill snstatd
```

Note: **fsmpm** will restart **snstatd** daemon.


11 Repeat [Step 8](#) through [Step 10](#) for Node 1.

12 On Node 2 **ONLY**, restart **carbon-cache**:

```
/opt/quantum/python27/bin/supervisorctl start  
carbon-cache
```

Verify output displays that **carbon-cache** is restarted:

```
carbon-cache: started
```

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
- [Capture DSET](#): Obtain and preserve detailed information about the current state of your StorNext M660, M440 or M330 Metadata Appliance
- [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

Health Check

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
- **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 120 Health Check Screen

Test	Start	Finish	Status
<input type="checkbox"/> Archive	15-Jul-2010 10:25:27 CDT	15-Jul-2010 10:25:28 CDT	Success
<input type="checkbox"/> Config	15-Jul-2010 10:25:28 CDT	15-Jul-2010 10:25:29 CDT	Success
<input type="checkbox"/> Disk Space	15-Jul-2010 10:25:30 CDT	15-Jul-2010 10:25:32 CDT	Success
<input type="checkbox"/> Drive	15-Jul-2010 10:25:32 CDT	15-Jul-2010 10:25:33 CDT	Success
<input type="checkbox"/> Media	15-Jul-2010 10:25:33 CDT	15-Jul-2010 10:25:35 CDT	Success
<input type="checkbox"/> Policies	15-Jul-2010 10:25:35 CDT	15-Jul-2010 10:25:38 CDT	Success

- 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**.

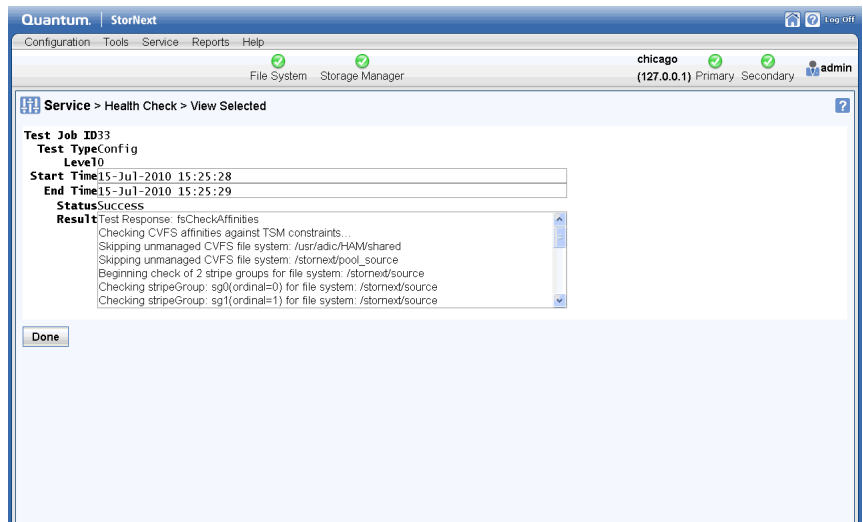
Viewing the Health Check Results

After a test has been run successfully (as indicated by a date and time in the respective **Start / Finish** columns, and **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 121 Health Check > View Selected Screen

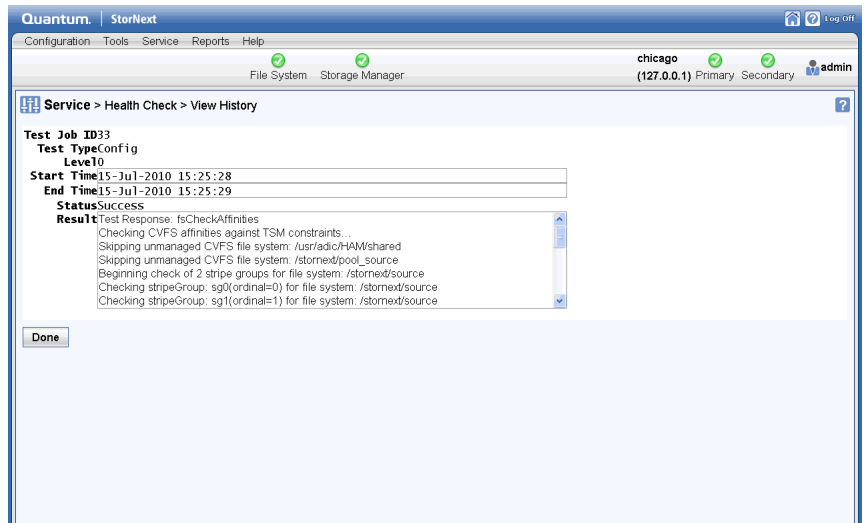


Viewing Health Check Histories

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 122 Health Check > View History Screen



- 3 When you are finished viewing history information, click **Done** to return to the **Service > Health Check** screen.

Capture State

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.

Note: If a snapshot file is not listed, then one will be created when the capture is selected.


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 123 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 123](#) on page 337.) 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.

Capture DSET

(This only applies to M660, M440 and M330 Metadata Appliances.)

The **Capture DSET** function uses the DSET (Dell System E-Support Tool) to create a log file which assists Quantum support personnel in identifying and diagnosing StorNext M660, M440 and M330 Metadata Appliance problems. The log file is a snapshot of system hardware status at the time when the report was generated.

Running **Capture DSET** creates a log file in this format:

```
dset-machinehostname-YYYYMMDDHHMMSS.zip
```


If desired, you can download or delete a captured DSET file.

Creating a Capture DSET File

Follow this procedure to create a Capture DSET log:

- 1 Choose **Capture DSET** from the **Service** menu. The **Service > Capture DSET** screen appears. Any previously captured DSET snapshots are shown.

Figure 124 Capture DSET Screen



- 2 Click **Capture**. The **Capture DSET Status** window is shown. The capture DSET file name appears in the list after the process completes.

Downloading a Capture DSET File

Follow these steps to download a capture DSET file:

- 1 Select the file you want to download, and then click **Download**.

- 2 When prompted, click the DSET file name to begin the download process.
- 3 Navigate to the location where you want to save the file.
- 4 After the file is saved, click **Done** to continue.
- 5 If you were instructed by Quantum Support personnel to send the DSET report, email the report (in .zip format) to the email address you were given by the Support representative.

Deleting a Previous Capture DSET File

Follow this procedure to delete an unwanted Capture DSET file.

- 1 If you have not already done so, choose **Capture DSET** from the Service menu. The **Service > Capture DSET** screen appears. (See [Figure 124](#) on page 339.) All previously captured DSET reports are shown.
- 2 Select the DSET 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**.

System Backup


The Service menu's Backup option allows you to perform a full or partial backup.

Note: StorNext must be configured with a library in order to perform a system backup, or an error will occur.

Note: Quantum recommends making two or more backup copies to minimize vulnerability to data loss in the event of hardware failure.

- 1 Choose **Backups** from the **Service** menu. The **Service > Backup** screen appears.

Figure 125 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**.

Admin Alerts


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 126 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

Tickets

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 127 Tickets Screen

Ticket	Status	Priority	Last Update	Component	Event
115	Open	high	2012-01-18 10:48:09 MST	Database component	Database corruption
114	Open	high	2012-01-18 10:48:09 MST	Database component	Database corruption
113	Open	high	2012-01-18 10:48:09 MST	Database component	Communication failure
112	Open	high	2012-01-18 10:48:09 MST	Database component	Operation failure
111	Open	high	2012-01-18 10:48:09 MST	Database component	Process/Task died (not restarted)
110	Open	middle	2012-01-18 10:48:09 MST	OS component	Operation failure
109	Open	high	2012-01-18 10:48:09 MST	OS component	System resource failure
108	Open	high	2012-01-18 10:48:09 MST	RAS component	Process/Task died (not restarted)
107	Open	high	2012-01-18 10:48:09 MST	Virtual Library Interface component	Communication failure
106	Open	high	2012-01-18 10:48:09 MST	Virtual Library Interface component	Process/Task died (not restarted)

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
 - **Component:** The name of the system component to which the event applies, such as "Database component" or "OS component."
 - **Event:** A short summary of the fault event that triggered creating the RAS ticket, such as "Data Corruption."
- 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 128 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
 - **Event Summary:** A brief description of event condition and the affected component.
 - **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.

Changing the Display View

There are three ways to display ticket information:

- **Page:** View tickets page by page, using the navigation controls at the lower right side of the table.
- **Scroll:** Use the scroll bar at the right side of the table to change the current view.
- **All:** Similar to Scroll mode, except the table expands to show all tickets. (In Scroll mode the table remains the same size.)

To change the current view, select Page, Scroll or All from the Table View pulldown field above the table on the right side of the screen.

Using Filter Options

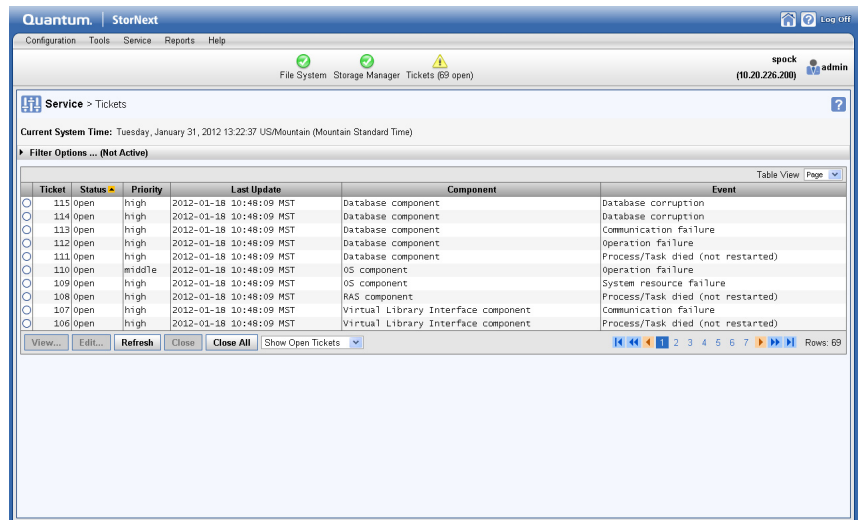
The filter options allow you to specify the ticket priority, component and event type for which you want to be notified. Unless you change filter options, the default is that you will received notifications for all

categories, components and events. (When everything is selected, Filter Options are Not Active.)

Follow this procedure to view or change the ticket filter options:

- 1 On the **Service > Tickets** screen, click the arrow to the right of the **Filter Options** heading to display the options.

Figure 129 Tickets > Filter Options Screen




- 2 Select or deselect the options according to your preference.
- 3 Click the arrow icon beside the Filter Options heading to collapse the list of options. Note that the Filter Options status changes to Active after you make changes.

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 130 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**.

Deleting Tickets

To delete tickets, perform one of the following procedures:

- a Select a ticket from the tickets table, and then click **Delete**. This procedure allows you to delete one ticket at a time.
- b Expand the **Filter Options** section, select the desired criteria, and then click **Delete Filtered**. This procedure allows you to delete the tickets that match the selected criteria in the **Filter Options**.

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 367.)


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 131 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 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](#). For additional information, see [FSM Failover In HA Environments](#).

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 550.

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

HA Reset 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 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*.

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, each of which must be sufficiently provisioned for the desired StorNext configuration. In addition, both MDCs must be running the same version of Linux. Variations in hardware provisioning, or software versions, could result in variations in observed performance characteristics between the two MDCs.
- 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 un-managed file system for use as the HA shared file system, which meets the following requirements:
 - The un-managed file system must be a file system that is not used for replication. (For more information about creating a file system, see [Step 5: File Systems](#).)

- The un-managed file system must be sufficiently provisioned for the desired StorNext configuration.
- The file system should not have quotas enabled. Enabling quotas on this file system can interfere with the proper functionality of Storage Manager and the HA infrastructure.

HA and LAN Clients

On a StorNext HA system using the StorNext LAN Client/Server 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 550.

In some cases the physical IP address must be included in the dpserver file in addition to the interface name. Note these conditions:

- 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. For additional information, see [StorNext LAN Clients in HA Environments](#).

StorNext LAN Clients in HA Environments

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 StorNext Gateway Server is configured after converting to HA, the file system(s) running as Gateway servers must be unmounted and mounted again to service StorNext LAN requests. 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. 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.

Converting to HA

Note: If you are upgrading to StorNext 5 or later from StorNext 4.7 or later, and have previously converted your system to HA configuration, the conversion process is unnecessary. Previously-converted systems will not be taken out of HA configuration, so the option to convert to HA is unavailable. However, for StorNext releases prior to 4.7, an HA conversion is necessary after StorNext software upgrades.

This section describes the configuration steps necessary to convert two StorNext MDC nodes into a High Availability MDC pair connected to a shared file system. Converting to HA consists of selecting the dedicated unmanaged StorNext file system for use as the controlling shared file system, and then instructing StorNext to convert each MDC node to operate as one MDC node of the HA pair. The following note and bullet items apply only to customer-supplied MDCs.

Note: The **Convert** menu option will be unavailable (grayed out) on the **Tools** menu if you have not specified a secondary system. If you have not already done so, specify a secondary system by using the Name Servers function. For more information, see [Name Servers](#) on page 60.

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 file system should not have quotas enabled. Enabling quotas on this file system can interfere with the proper functionality of Storage Manager and the HA infrastructure.

- The conversion process converts one MDC node at a time. The second MDC 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.

If Replication/Deduplication is configured, the following apply:

- 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.

The following bullet item applies to all MDCs used in an HA configuration, including StorNext Metadata Appliances:

- The UIDs for the **quantumdb** and **tdlm** users and the **adic** group **must** be identical on both nodes of an HA pair. If the UIDs are not identical, the MySQL database will not start (due to file permission errors), which in turn prevents storage manager from fully starting up. Quantum recommends creating the UIDs for the **quantumdb** and **tdlm** users along with the **adic** group on both nodes prior to running `install.stornext`.

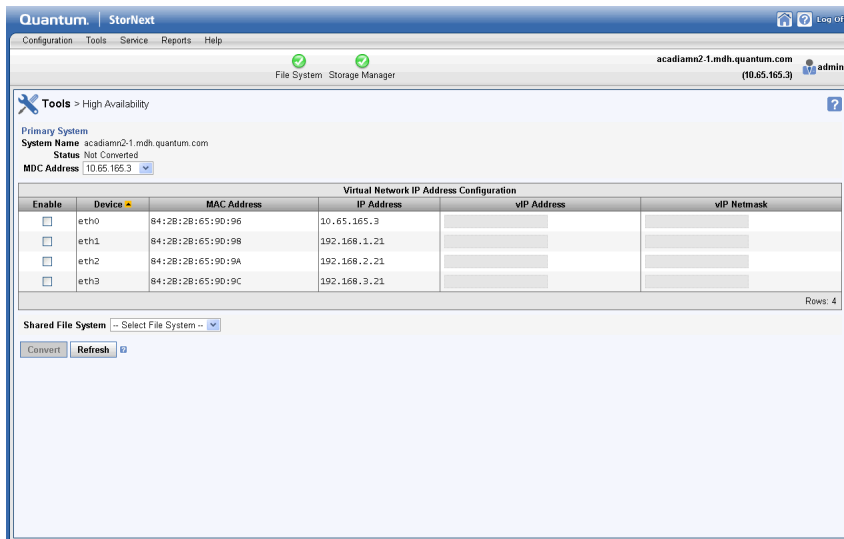
HA Conversion Procedure

Follow these steps to configure HA:

- 1 Choose **Tools > High Availability > Convert**. If the system has previously been converted to HA mode, continue to [Step 7](#). If this is the first time the system will be converted to HA, notice the Status for the **Primary System** is set to “Not Converted”.

The **High Availability** screen appears.

Figure 132 Tools > HA >
Convert (primary node not yet
converted)



- 2 At the **Shared File System** field, select the shared file system you want to use for the HA pair.

Caution: Once you convert a file system to HA you cannot undo the operation, so choose carefully when selecting the file system.

Note: The file system should not have quotas enabled. Enabling quotas on the file system can interfere with the proper functionality of Storage Manager and the HA infrastructure.

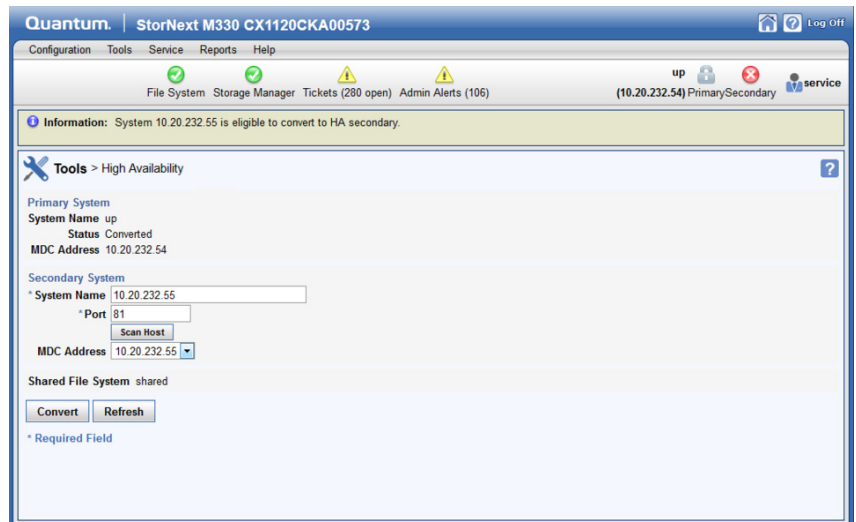
Note: MDCYou can set a virtual IP (VIP) (requires a netmask setting) for one or more Eth ports at this time. Use IP addresses assigned to the port used for the Management network on each MDC node for the VIP values. Do not assign VIPs for the IP addresses assigned to the ports used for the Metadata network.

- 3 For the **MDC Address**, select one IP address to be placed in the `ha_peer` file for use in administrative operations between the MDC nodes in the HA pair.

- 4 For **non-Lattus systems ONLY**: 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.)
- 5 Click **Convert** to convert the primary MDC node to HA and click **Yes** to confirm the conversion.
- 6 Once the primary MDC node has been converted, **Status** will change to **“Converted”**. Click **OK** to continue.
- 7 Enter the IP or DNS address of the secondary Node in the System Name field (this node must be on the same LAN as the primary MDC node).

Notice the Status for the Primary System is shown as **“Converted”**, and the note that the other node is available to be converted to HA secondary.

Figure 133 Tools > HA > Convert (primary node previously converted)



- 8 Click **Scan Host**. The system should resolve the secondary node - and the **MDC Address** drop-down list will auto-fill with the IP address of the secondary node. 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

MDC nodes. 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 the StorNext 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.

- 9 Click **OK** to convert the secondary node.
- 10 Storage Manager, may need to be started following the HA conversion if the system was in config mode at the time that HA conversion was initiated. To restart the Storage Manager components, click the **Start** button in the Storage Manager panel of the **Tools > System Control** page.

GUI Feedback During HA Conversion

There are some indications within the GUI that the system is being upgraded. Here are some notes about this visual feedback:

- After the StorNext upgrade has completed, and the HA configuration has been done, the GUI for the secondary MDC node provides a message stating it is not the primary MDC node and a link to launch the primary MDC node.

Caution: Do not login to the GUI of the secondary MDC node at any point during the upgrade/HA conversion process. System configuration for the system could be compromised.

- When you are able to log into the primary system/node, after accepting the EULA, the system will automatically display the **Tools > System Control** page. Click the **Start** button to restart the Storage Manager components.
- Wait until the system icons for both MDC nodes of the system as well as File System and Storage Manager are green, which indicates normal operation.

Post-Conversion Steps

If you are using the DDM feature, do the following:

- If you use the secondary MDC node as a DDM mover, make sure the file systems are mounted.
- 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.

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.

Caution: 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. Prepare your systems before entering HA Config mode.

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 134 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.

In a HA configuration, RAS messages are not generated for loss of SAN connectivity by the secondary system. A workaround is to activate some of the un-managed file systems on the secondary metadata controller. This allows RAS messages from the secondary MDC if there is SAN connectivity loss. Perform this action to also help with load balancing.

Troubleshooting HA

The Troubleshooting appendix in this guide contains simple troubleshooting procedures pertaining to HA. For more information, see [Appendix G, Troubleshooting](#).

For an in-depth look at HA systems and operation, see [Appendix C, High Availability Systems](#).



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 LAN clients, and client performance
- [The LAN Client Performance Report](#): View information about 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
- [The Hardware Status Report](#): View up-to-date information about the system board and network ports for both nodes in StorNext M660, M440 and M330 Metadata Appliances, plus storage arrays
- [The Gateway Metrics Report](#): View information and activity related to your gateways, clients, and file systems


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.
- **Shared Firmware Upgrade Logs:** Logs that track firmware upgrade status of both nodes.

Note: The **Shared Firmware Upgrade Log** only applies to the M-series appliance.

- **Local Firmware Upgrade Logs:** Logs that track firmware upgrade status of the current primary node.

Note: The **Local Firmware Upgrade Log** only applies to the M-series appliance.

- **Quota Logs:** Logs that track quota events and status, if this feature is enabled on a given file system.

Administrative Logs and Alerts

StorNext Storage Manager generates administrative logs for events which are worthy of notification. Such events generate e-mail notifications, while the rest are only displayed in the GUI.

There are a few notifications which are throttled, such that every occurrence does not generate a notification in order to prevent flooding you with these events repeatedly.

The default time between notification of these events is 8 hours. You can adjust the value by setting the **FS_ADMINLONG_MINTIME** system parameter in one of the following configuration files:

```
/usr/adic/TSM/config/fs_sysparm
```

```
/usr/adic/TSM/config/fs_sysparm_override
```

The following events are subject to throttling:

- The inability of a host to determine where to run a distributed data mover.
- A file system not being configured for distributed data movers.
- The inability to allocate a storage disk I/O stream to use for a distributed data mover.
- Encountering a duplicate entry in the database when storing a file.

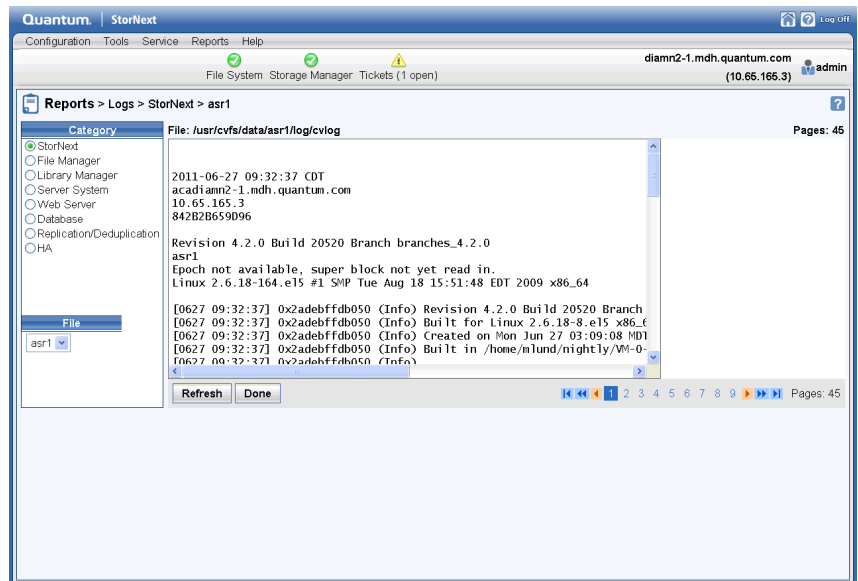
- Attempts to format a media which may already be formatted and actively used.
- **fsaddrelation** failure due to an inconsistency with policy attributes between replication/deduplication and **fspolicy**.

How to Access StorNext Log Files

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 135 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 366.

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 136 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 137 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 138 StorNext File
Browser



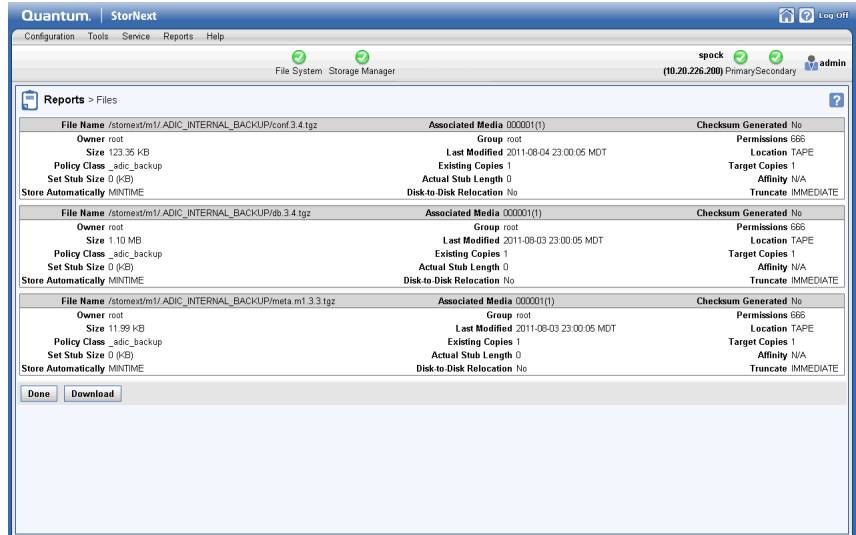
3 Do one of the following:

- 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 139 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.

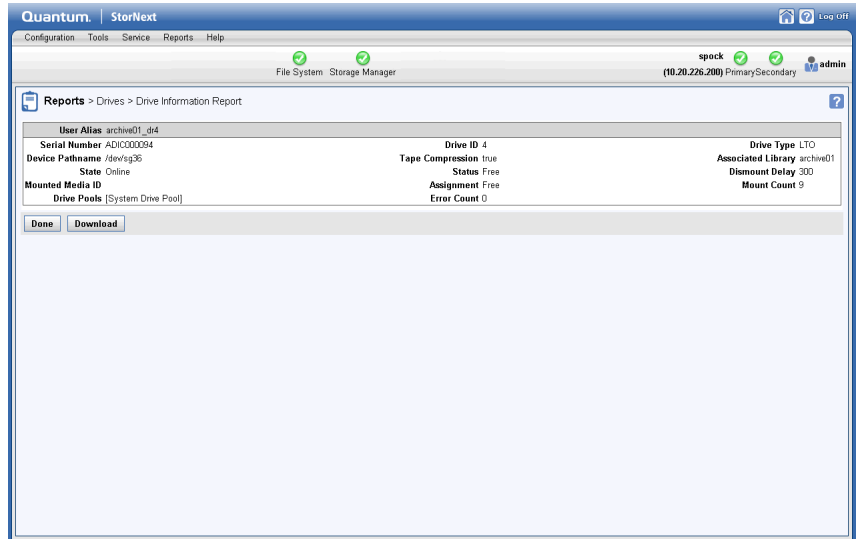
Figure 140 Drives Report

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 141 Drive Information Report



- 3 To download the report, click **Download**.
- 4 When you finished viewing report information, click **Done**.

The Media Report

The Media Report shows a list of details for all media (including cleaning 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 142 Media Report

Media ID	Library	Media Type	Media Class	Policy Class	Suspect	Write Protected	File Count	% Used	Copy	Mounted in Drive	Last Accessed
000000	archive01	LTO	F0_LTO_DATA	pc1	No	No	70	0.01%	2		2011-07-11 19:23:09 MDT
000001	archive01	LTO	F0_LTO_DATA	_ladic_backup	No	No	92	0.02%	1		2011-08-08 23:05:08 MDT
000002	archive01	LTO	F0_LTO_DATA	Scratch Pool	No	No	0	0%	0		2011-07-11 19:09:57 MDT
000003	archive01	LTO	F0_LTO_DATA	Scratch Pool	No	No	0	0%	0	archive01_dr2	2011-07-11 19:09:57 MDT
000004	archive01	LTO	F0_LTO_DATA	Scratch Pool	No	No	0	0%	0		2011-07-11 19:09:57 MDT
000005	archive01	LTO	F0_LTO_DATA	Scratch Pool	No	No	0	0%	0		2011-07-11 19:09:57 MDT
000006	archive01	LTO	F0_LTO_DATA	Scratch Pool	No	No	0	0%	0		2011-07-11 19:09:57 MDT
000007	archive01	LTO	F0_LTO_DATA	Scratch Pool	No	No	0	0%	0		2011-07-11 19:09:57 MDT
000008	archive01	LTO	F0_LTO_DATA	Scratch Pool	No	No	0	0%	0		2011-07-11 19:09:57 MDT
000009	archive01	LTO	F0_LTO_DATA	Scratch Pool	No	No	0	0%	0		2011-07-11 19:09:57 MDT
000011	archive01	LTO	F0_LTO_DATA	Scratch Pool	No	No	0	0%	0		2011-07-11 19:09:57 MDT
000012	archive01	LTO	F0_LTO_DATA	Scratch Pool	No	No	0	0%	0		2011-07-11 19:09:57 MDT
000013	archive01	LTO	F0_LTO_DATA	Scratch Pool	No	No	0	0%	0		2011-07-11 19:09:57 MDT
000014	archive01	LTO	F0_LTO_DATA	Scratch Pool	No	No	0	0%	0		2011-07-11 19:09:57 MDT
000015	archive01	LTO	F0_LTO_DATA	Scratch Pool	No	No	0	0%	0		2011-07-11 19:09:57 MDT
000016	archive01	LTO	F0_LTO_DATA	Scratch Pool	No	No	0	0%	0		2011-07-11 19:09:57 MDT
000017	archive01	LTO	F0_LTO_DATA	Scratch Pool	No	No	0	0%	0		2011-07-11 19:09:57 MDT
000018	archive01	LTO	F0_LTO_DATA	Scratch Pool	No	No	0	0%	0		2011-07-11 19:09:57 MDT
000019	archive01	LTO	F0_LTO_DATA	Scratch Pool	No	No	0	0%	0		2011-07-11 19:09:57 MDT

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 Select the media from the **Policy Class** list. This filter allows you to narrow down the media assigned to a particular policy class without having to sort and page through a potentially large set of media.

Note: The **N/A (Spans Multiple)** filter only applies to **Storage Disks** and **WAS Media**. If a storage disk or WAS media has not been assigned an explicit policy class, the storage manager system can store data from multiple policy classes on these media. This is denoted by the "N/A" designation. 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 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).
- 6 Click **Filter** to apply the filtering options.
- 7 To update the information displayed, you can click **Refresh** at any time.

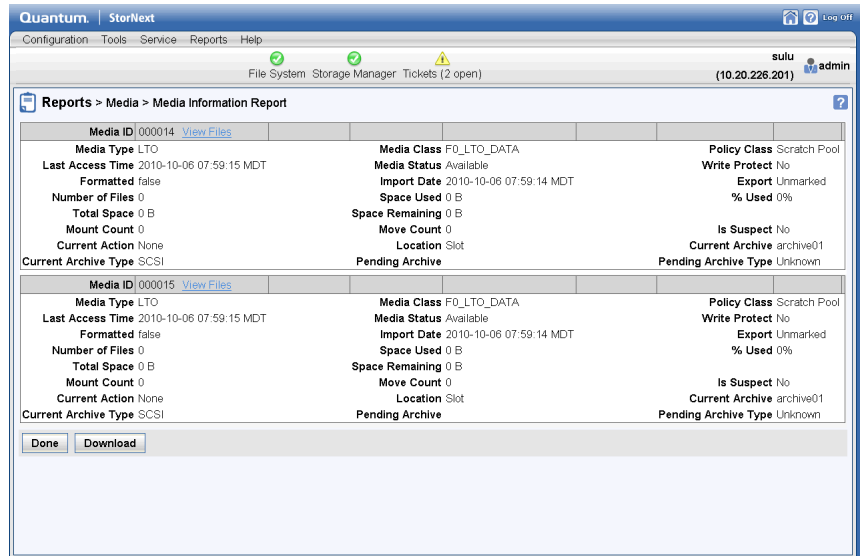
The Media Report provides the following information for each piece of media that meets the criteria you selected from the **Library**, **Media Class** and **Media ID** fields:

- **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
- 8 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 143](#).

Figure 143 Media Information Report



- 9 (Optional) Save the report output as a CSV file (Microsoft Excel format) by clicking **Download**.
- 10 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 144 Relations Report



- 2 When you are finished reviewing the report output, 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 145 File Systems Report

File System	Mounted on	Status
pool_target	/stornext/pool_target	Up
Stripe Group 1: sg0	Content: MJ	Status: Up
Total Space 99.98 GB	Depth 1	Realtime IO Limit 0/sec
Reserved Space 0 B	Breadth 4 MB	Realtime Bandwidth Limit 0/sec
Free Space 99.92 GB	Multipath Method Rotate	Realtime IO Commit 0/sec
Stripe Group 2: sg1	Content: U	Status: Up
Total Space 99.99 GB	Depth 1	Realtime IO Limit 0/sec
Reserved Space 4.13 GB	Breadth 2 MB	Realtime Bandwidth Limit 0/sec
Free Space 96.21 GB	Multipath Method Rotate	Realtime IO Commit 0/sec
target	/stornext/target	Up
Stripe Group 1: sg0	Content: MJ	Status: Up
Total Space 99.98 GB	Depth 1	Realtime IO Limit 0/sec
Reserved Space 0 B	Breadth 4 MB	Realtime Bandwidth Limit 0/sec
Free Space 99.42 GB	Multipath Method Rotate	Realtime IO Commit 0/sec
Stripe Group 2: sg1	Content: U	Status: Up
Total Space 199.98 GB	Depth 2	Realtime IO Limit 0/sec
Reserved Space 4.13 GB	Breadth 2 MB	Realtime Bandwidth Limit 0/sec
Free Space 196.81 GB	Multipath Method Rotate	Realtime IO Commit 0/sec

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 146 SAN Devices Report

Quantum | StorNext

Configuration Tools Service Reports Help

File System Storage Manager

portland (10.65.14.34) admin

Reports > SAN Devices

Current SAN Devices

Disks/LUNs						
Serial Number	Type	Label	Size	Status	Used	File System
60014380024D2E990001400000290000	GENERIC_209696735	HP_EVA_0001	99.99 GB	up	false	
60014380024D2E9900014000002C0000	GENERIC_209696735	HP_EVA_0002	99.99 GB	up	false	
60014380024D2E9900014000002F0000	GENERIC_209696735	HP_EVA_0003	99.99 GB	up	false	
60014380024D2E990001400000850000	GENERIC_209696735	HP_EVA_0004	99.99 GB	up	false	
60014380024D2E990001400000880000	GENERIC_209696735	HP_EVA_0005	99.99 GB	up	false	
60014380024D2E9900014000008B0000	GENERIC_209696735	HP_EVA_0006	99.99 GB	up	false	

Rows: 13

Libraries and Tape Drives (Some devices may not show up immediately. Click Refresh to update.)			
Serial Number	Product ID	Device Type	Device Path
HUI732085N	Ultrium 4-SCSI	Tape Drive	/dev/sg29
HUI7410GFw	Ultrium 4-SCSI	Tape Drive	/dev/sg28
QP0739BDC00082	SCALAR50	Tape Library	/dev/sg30

Rows: 3

Refresh Done

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.

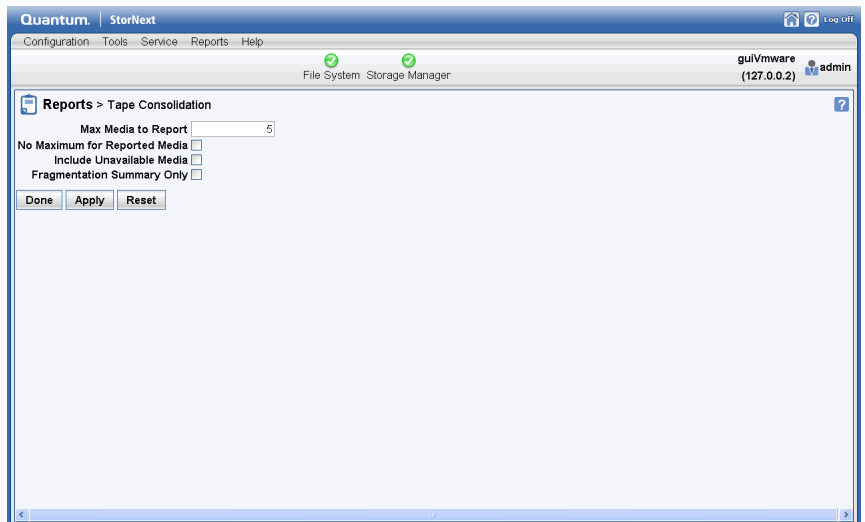
The Tape Consolidation Report

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 147 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 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 148 SAN and LAN Clients Report

Quantum | StorNext

Configuration Tools Service Reports Help

File System Storage Manager

portland (10.65.14.34) admin

Reports > SAN and LAN Clients

File System: pool_target Mounted on: /stormext/pool_target Status: ✔

SAN Clients 1 [10.65.14.34]

Distributed LAN Clients 0 []

Distributed LAN Servers(0)

Server	Listening Interface (IP:Port)	TCP Window Size	Transfer Buffer Size	Transfer Buffer Count	Server Buffer Count	Daemon Threads
--------	-------------------------------	-----------------	----------------------	-----------------------	---------------------	----------------

File System: target Mounted on: /stormext/target Status: ✔

SAN Clients 5 [10.65.14.34, 10.65.14.35, 10.65.174.233, 10.65.14.33, 10.65.167.243]

Distributed LAN Clients 2 [10.65.167.243, 10.65.174.233]

Distributed LAN Servers(1)

Server	Listening Interface (IP:Port)	TCP Window Size	Transfer Buffer Size	Transfer Buffer Count	Server Buffer Count	Daemon Threads
10.65.14.34	10.65.14.34:55661	0	0	0	0	0

Refresh Done Rows: 2

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.
- **LAN Clients:** The total number of StorNext LAN clients.
- **Gateway Servers:** The total number of gateway servers for the file system.
- **Server:** The names of the gateway servers.
- **LAN Clients:** The names of LAN clients.
 - **Listening Interface (IP:Port):** The IP address and port number through which the gateway server communicates with StorNext.
 - **TCP Window Size:** The TCP window size (in KB) used by the gateway server. (Default: 64)

- **Transfer Buffer Size:** The transfer buffer size (in KB) used by the gateway server. A larger buffer may increase performance for larger files. (Default: 256)
 - **Transfer Buffer Count:** The number of transfer buffers used by the gateway 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 gateway 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 StorNext 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 149 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	192.168.200.1:93983	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	10.65.186.252:44308	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 gateway 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 StorNext LAN client for the indicated file system and gateway server.
 - **Read Bytes:** The number of bytes read by the StorNext LAN client.
 - **Write Bytes:** The number of bytes written by the StorNext LAN client.
- 2 If desired, click **Refresh** to manually update (refresh) the report data.
 - 3 Click **Done** when you are finished viewing the report.

Replication Deduplication Reports

StorNext provides these reports that contain information pertaining to replication and deduplication:

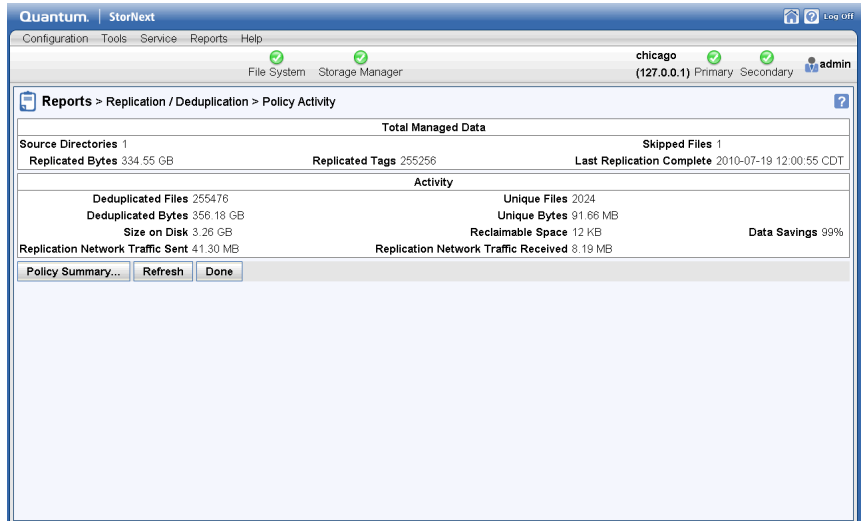
- **Policy Activity:** This report shows statistics related to replication and deduplication. This report also provides statistics such as savings realized by replication, and the current progress of ongoing replications.
- **Policy Summary:** This report shows information about replication storage policies 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 150 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 151 Replication/
Deduplication Policy Summary
Report

The screenshot shows the StorNext Reports interface. The breadcrumb path is 'Reports > Replication / Deduplication > Policy > Summary'. The report table has the following data:

Running Policy	File System	Source Directory	Replicated			Replication			Replication Network Traffic		
			Directories	Files	Skipped	Bytes	Average Rate	Estimated Completion	Last Completed	Sent	Received
rep1	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


Buttons below the table: 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 152 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 153 Replication/
Deduplication Policy
Completion Report



6 To view the Policy Activity report, click **Policy Activity**.

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 154 Distributed Data Mover Report

Reports > Distributed Data Mover						
Activity						
Host	Request ID	Device Alias	Run Time	Total Files	Files Copied	Files Failed
portland	2019715659	archive01_dr2	20:28:42	300	254	0
10.65.14.35	2019715659	archive01_dr2	20:28:42	300	254	0

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**.

The Hardware Status Report

The Hardware Status Report shows up-to-date information about the system board, network ports and fibre channel ports for both nodes in your HA system, plus storage arrays.

Note: This report is only visible on StorNext M660, M440 and M330 Metadata Appliances.

Use the following procedure to run the Hardware Status Report.

- 1 Choose **Hardware Status** from the **Reports** menu. The **Hardware Status Report** appears.

Figure 155 Hardware Status Report System Board Tab

Node 1				Node 2			
Component	Type	Value	Status	Component	Type	Value	Status
IPMI	IPMI	NA	Normal	IPMI	IPMI	NA	Normal
Ambient Temperature	Temperature	21 degrees C	Normal	Ambient Temperature	Temperature	21 degrees C	Normal
CMOS Battery	Alarm	NA	Normal	CMOS Battery	Alarm	NA	Normal
CPU1 VCORE PG	Alarm	NA	Normal	CPU1 VCORE PG	Alarm	NA	Normal
0.75 VTT CPU1 PG	Alarm	NA	Normal	0.75 VTT CPU1 PG	Alarm	NA	Normal
1.5V PG	Alarm	NA	Normal	1.5V PG	Alarm	NA	Normal
1.8V PG	Alarm	NA	Normal	1.8V PG	Alarm	NA	Normal
3.3V PG	Alarm	NA	Normal	3.3V PG	Alarm	NA	Normal
5V PG	Alarm	NA	Normal	5V PG	Alarm	NA	Normal
MEM PG	Alarm	NA	Normal	MEM PG	Alarm	NA	Normal
VTT PG	Alarm	NA	Normal	VTT PG	Alarm	NA	Normal
0.9V PG	Alarm	NA	Normal	0.9V PG	Alarm	NA	Normal
1.8 PLL PG	Alarm	NA	Normal	1.8 PLL PG	Alarm	NA	Normal
8.0 V PG	Alarm	NA	Normal	8.0 V PG	Alarm	NA	Normal
1.1 V PG	Alarm	NA	Normal	1.1 V PG	Alarm	NA	Normal

(At any time, you can click **Refresh** to manually update the report data.)


The first tab, **System Board**, shows the following information about the system board in both nodes:

- **Component:** The name of the component on the system board, such as the CMOS battery, IPMI, fans, and so on.
- **Type:** Indicates the category for the component, such as “Alarm” for the CMOS battery.

- **Value:** If applicable, indicates the current value for the component, such as “27 degrees” for the Ambient Temperature.
- **Status:** Indicates the current state of the component, such as “Normal.”

2 When you are finished viewing system board information, click the **Network Ports** tab. The **Network Ports** screen appears.

Figure 156 Hardware Status Report Network Ports Tab




The **Network Ports** tab displays the following information related to the network ports on both nodes:

- **Component:** The name of the port component, such as “eth1” for Ethernet port 1.
- **Value:** Indicates the correct value for the component, such as “1000 Mb/s” for the Ethernet port.
- **Status:** Indicates the current state of the component, such as “Up” or “Down.”

- 3 When you are finished viewing network ports information, click the **Fibre Channel Ports** tab. The **Fibre Channel Ports** screen appears.

Figure 157 Hardware Status
Report Fibre Channel Ports Tab




The **Fibre Channel Ports** tab displays the following information related to the fibre channel ports on both nodes:

- **Node name and status**
- **Unit name, location, port number and status**
- **Port name and status**

- When you are finished viewing fibre channel port information, click the **Storage Arrays** tab. The **Storage Arrays** screen appears.

Figure 158 Hardware Status Report Storage Arrays Tab



The **Storage Arrays** tab displays the following about your storage array:

- **Source:** The array node or cluster.
- **Array:** The name of the array.
- **Type:** The array manufacturer.
- **Status:** Indicates the current state of the component, such as "Attention" or "Normal".
- **Component Group:** The array component category, such as trays, controllers, volumes, drives, alarms and so on.
- **Count:** The current number of each component group in the array (for example, "2 Volumes").
- **Status:** The current operational status of the component group, such as "Normal".

Note: Information for the following headings appears after you select a component group.

- **Component:** The name or location indicator of the selected component group (such as "SAS Phy 1").

- **Location:** The current location of the selected component group.
- **Value:** The current value of the selected component group, such as "Up" or "Down".
- **Status:** The current operation status of the selected component group, such as "Normal".

The Gateway Metrics Report

The Gateway Metrics Report helps you monitor performance and throughput on your gateways, clients and file systems. Because you can see at a glance which gateways, clients or file systems are currently under or over-utilized, the Gateway Metrics Report is a useful tool for load balancing.

The Gateway Metrics Report consists of five main sections:

- Two graphs at the top of the screen which show at-a-glance information for your gateways, clients and file systems. These graphs change according to your selection. A large banner above the graphs provides a summary of total read and writer throughput for all gateways.
- A **Gateway Summary** section which shows a list of all gateways and statistics for each one.
- A **Client Summary** section which shows a list of all clients and statistics for each one.
- A **File System Summary** section which shows a list of all file systems and statistics for each one.
- A **Settings** section where you can enable or disable metrics tracking for one or more gateways.

By default, the summary sections show the average overall throughput, average read throughput and average write throughput for each specific gateway, client or file system listed. You can customize the information displayed by choosing additional columns or removing currently displayed columns as described in [Changing Displayed Columns](#) on page 408.


Note: You can collapse or expand any section by clicking the double arrow icon at the right side of the section title bar.

Use the following procedure to run the Gateway Metrics Report.

- 1 Choose **Gateway Metrics** from the **Reports** menu. The **Gateway Metrics Report** appears. By default, the initial view is in Summary mode and displays the Last Hour of aggregate throughput of all gateways. For information about viewing the report in Detail mode, see [Changing to Detail Mode](#) on page 404.

Chapter 10: StorNext Reports The Gateway Metrics Report

Figure 159 Gateway Metrics Report (Summary)



Note: The screen in [Figure 159](#) does not show all columns available on this report. See [Changing Displayed Columns](#) on page 408 for information on adding or removing columns from the report.

- 2 If desired, click **Refresh** to manually update (refresh) the report data.
- 3 You can change the information displayed in the graphs or view detailed information rather than a summary as described in the following sections.

Changing the Graph Display

You can use the radio buttons above the charts to change the level of detail shown on the graphs to display information from the **Day**, **Last 30 Days**, **Last 24 Hours**, or **Last Hour**.

You can also specify the time period for the graphs by using the **Select Time Period** field at the left of the radio buttons.

To change the time intervals used in the graphs displays, at the **Zoom** field click one of the following:

- **10m:** Displays graph information in ten minute increments. For example, from 7:00, 7:10, 7:20, 7:30 and so on.
- **1h:** Displays graph information in one hour increments. For example, from 1:00, 2:00, 3:00, 4:00 and so on.
- **2h:** Displays graph information in two hour increments. For example, from 1:00, 3:00, 5:00, 7:00 and so on.
- **6h:** Displays graph information in six hour increments. For example, from 12:00 (midnight), 6:00 a.m, 12:00 (noon), 6:00 p.m. and so on.
- **1d:** Displays graph information in one day increments.
- **30d:** Displays graph information in thirty day increments.

You can also manually compress or expand the displayed graph information by grabbing and dragging one of the handles at the far left and rights ends of the scroll bar below the graphs.

Note: On particularly large systems that generate a lot of metric information, if you choose the Last 30 Days option your system could reach the row limit for the Gateway Metrics report before the full 30 days can be displayed.

When you hover the mouse pointer over the graphs, an information box shows you the date and time at which a data point was captured for the displayed gateway, client or file system. You can view other data points by repositioning the mouse pointer over the graphs.

Changing to Detail Mode

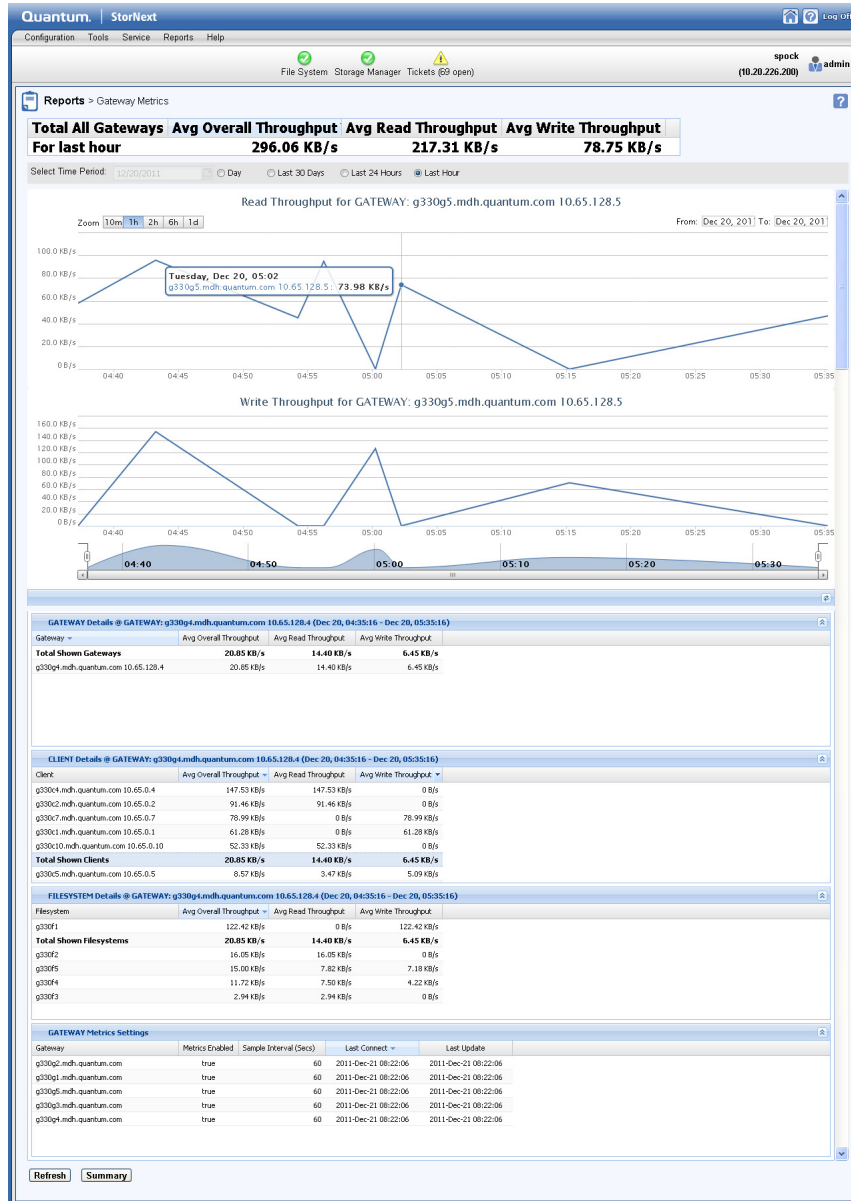
If desired, you can view the Gateway Metric report in Detail mode rather than Summary mode. To switch to Detail mode, double-click any gateway, client or file system name.

When you are in Detail mode, single-clicking the name of a specific gateway, client or file system updates the graphs to show detail for only the selected gateway, client or file system.

When you are viewing the Gateway Metrics Report in Detail mode you can return to Summary mode at any time by clicking the **Summary** button.

The Gateway Metrics Report in Detail mode view looks similar to the one shown in [Figure 160](#) (after double-clicking on a gateway name).

Figure 160 Gateway Metrics Report (Detail)



Note: The screen in [Figure 160](#) does not show all columns available on this report. See [Changing Displayed Columns](#) on page 408 for information on adding or removing columns from the report.

Viewing a Client or File System on a Gateway

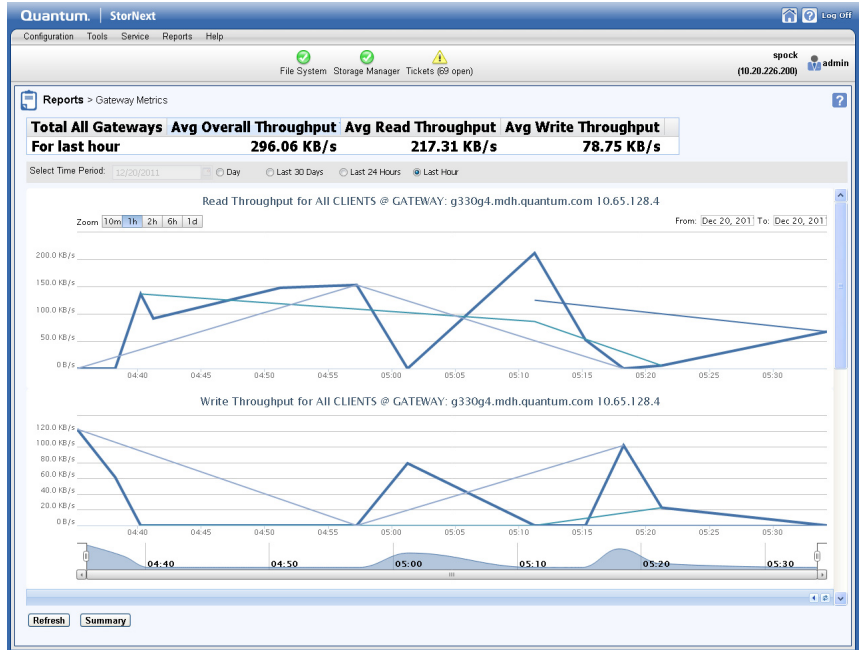
Another Detail mode display option is to first select a gateway by double-clicking the desired gateway name, and then doing one of the following:

- Single-click an individual client or file system name. The graphs update to show the selected client or file system on the selected gateway only.
- Single-click the heading **Total Shown Clients** or **Total Shown Filesystems**. The graphs update to show all clients or all file systems on the selected gateway.

Note: The Gateway Metrics Report charts will display up to 20 series. When you hover on a chart data point, the tool-tip popup will display the data point values, but the height of the chart window only allows for approximately 12 data values to be shown for a given sample. If not all data points are displayed in the popup, select the appropriate row in the gateway, client or filesystem table in order to individually display the desired series and examine the data values.


The graphs for a specific client on the gateway look similar to this (after single-clicking the client name in Detail mode):

Figure 161 Gateway Metrics Report (Client Detail)



The graphs for a specific file system on the gateway look similar to this (after single-clicking the file system name in Detail mode):

Figure 162 Gateway Metrics Report (File System Detail)



Changing Displayed Columns

In either summary or detail display mode, you can show fewer or more columns by clicking the down-arrow icon to the right of a column heading and then moving your mouse over the Columns heading in the dropdown.

Currently displayed columns are indicated by a green arrow inside the checkbox to the left of the column name. Check any additional columns you want to show, or deselect any currently displayed columns you want to hide.

Following are the available columns for the gateway, clients and file system sections:

- **Gateway** (in Gateway section): The name of the gateway.
- **Client** (in Client section): The name of the client.

- **File System** (in File System section): The name of the file system.
- **Average Overall Throughput**: The average overall data throughput including reading and writing per second.
- **Read Count**: The total number of read (I/O) operations.
- **Read Rate**: The number of reads (I/O) per second.
- **Read Bytes**: The number of bytes read.
- **Average Read Throughput**: The average read throughput per second.
- **Read Distribution**: The percentage of data read by the specific gateway, client or file system.
- **Write Count**: The total number of write operations.
- **Write Rate**: The number of writes per second.
- **Write Bytes**: The total number of bytes written.
- **Average Write Throughput**: The average write throughput per second.
- **Write Distribution**: The percentage of data written by the specific gateway, client or file system.
- **Utilization**: The currently selected Gateway Network Interface Card's maximum aggregate throughput.

Viewing and Changing Gateway Metrics Settings

The Gateway Metrics Settings section allows you to enable or disable gateway metric tracking for a single gateway, multiple gateways or all gateways. This section of the Gateway Metrics screen shows the following settings for all gateways:

- **Gateway**: The name of the gateway.
- **Metrics Enabled**: Shows whether metrics tracking is enabled (True) or disabled (False).
- **Sample Interval (Secs)**: The interval in seconds during which metrics are captured. For example, if the sample interval is 60 seconds (the default value), metrics will be sampled every minute.

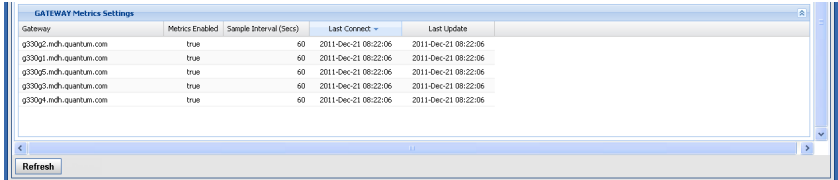
The Sample Interval value interacts with the currently selected time period (Day, Last 30 Days, and so on). The report data returned is based on the default sample interval of 60 seconds, so if you change the default sample interval to a value less than 60 seconds (effectively creating more sampled data), you could end up with less total data. For example, if the time period is set at Last 30 days and

you change the sample interval to 30 seconds, the reports would fill in 15 days rather than 30 days.

- **Last Connect:** The date and time when the last connection was established with the gateway.
- **Last Update:** The date and time when the statistics and graphs on the Gateway Metrics screen were last updated. Gateway metrics must be enabled (True) for a particular gateway to be included in the update.

Note: You can hide or show columns in the Gateway Metrics Setup section just as you can in the other sections.

Figure 163 Gateway Metrics Settings



Gateway	Metrics Enabled	Sample Interval (Secs)	Last Connect	Last Update
g300g2.mdh.quantum.com	true	60	2011-Dec-21 08:22:06	2011-Dec-21 08:22:06
g300g1.mdh.quantum.com	true	60	2011-Dec-21 08:22:06	2011-Dec-21 08:22:06
g300g5.mdh.quantum.com	true	60	2011-Dec-21 08:22:06	2011-Dec-21 08:22:06
g300g3.mdh.quantum.com	true	60	2011-Dec-21 08:22:06	2011-Dec-21 08:22:06
g300g4.mdh.quantum.com	true	60	2011-Dec-21 08:22:06	2011-Dec-21 08:22:06

Follow this procedure to enable or disable metrics:

- 1 Double-click the gateway whose metrics you want to enable or disable.
- 2 To enable metrics, click the checkbox to the right of the gateway name. To disable metrics, remove the checkmark.
- 3 If desired, specify a sample interval by either clicking the up or down arrows or typing a value in the field to the right of the checkbox. The valid range is between 10 seconds and 3600 seconds.
- 4 Save your changes by clicking **Update**, or click **Cancel** to abort.



Chapter 11

Wide Area Storage (Lattus)

StorNext Storage Manager (SM) now has a new storage destination for copies of managed files to go along with Tape and Storage Disk. The new destination is known as Wide Area Storage (Lattus) and only works with the Quantum Lattus product. Lattus devices provide Policy Based Reliability that gives the customer up to fifteen nines of durability that be configured to handle many component failures but remain functional with very little to no degradation in performance.

This chapter covers the following topics:

- [Wide Area Storage Features in the StorNext \(GUI\)](#) on page 413
- [Configuring Lattus Object Storage](#) on page 420
- [Setting Up Lattus Object Storage Destinations](#) on page 424
- [HTTPS Support for Lattus Object Storage](#) on page 435
- [Changes to Existing CLI Commands](#) on page 438
- [Other Changes and Considerations](#) on page 440
- [Wide Area Storage Segment Size](#) on page 440

Note: This section is specific to SNSM Wide Area Storage and its use of Lattus and does not cover the configuration of the hardware, etc. Those items are covered in separate documentation available online at: <http://www.quantum.com/lattusdocs>

Audience

This chapter is targeted at users who will be configuring the new Lattus destination or making use of managed files that have been stored to that destination.

Overview

When SM is managing user data it will make copies of those managed files on SM media. At policy creation time the classes created will have the storage destinations defined. The new Lattus destination is a set of external object storage devices that can be used for reliable long term data storage. Lattus devices function at a basic level in the same way as physical tape media or Storage Disk. You can add up to 32 independent Lattus devices and there is no software limit on the number of namespaces. The actual number of namespaces that can be configured is dependent on expected usage and performance requirements. The devices that make up the Wide Area Storage are referred to as Lattus devices. Refer to the *Lattus Service Reference Guide* and/or the *Lattus Installation Reference Guide* for more information on Quantum's Lattus devices.

A Lattus namespace is used like a Storage Disk or a Tape Media. In fact the namespaces are also referred to as Lattus Media. There can be many namespaces in a single Lattus device. A namespace is first created in the Lattus device, and later, made known to StorNext SM. StorNext SM moves data to a Lattus namespace for long-term retention in addition to or instead of tape and Storage disk. This enables users to leverage the extreme data reliability functionality of Lattus.

Lattus Media versus Storage Manager Media

Here are a few comparisons between Lattus devices and media (namespaces) as compared to other SM media:

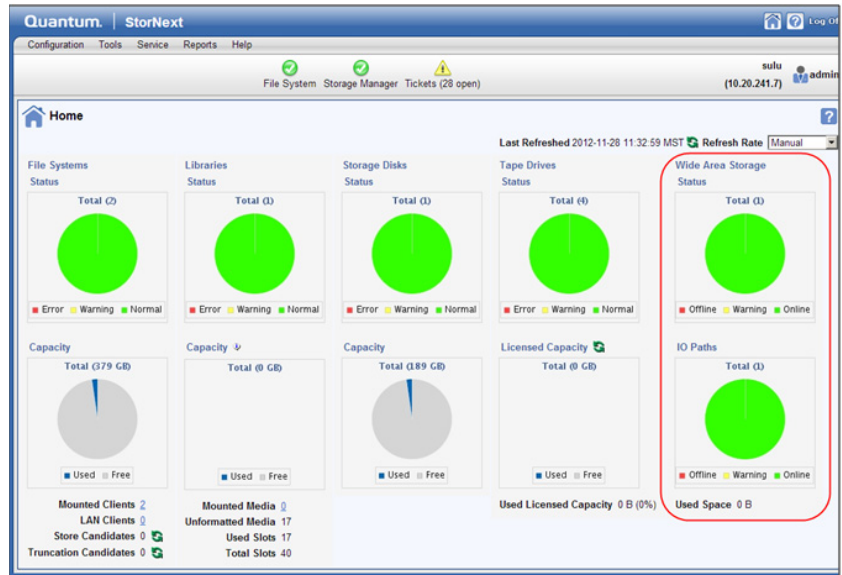
- A Lattus media belongs to no policy class. That is, multiple classes can have data on the same media. (This matches sdisk)
- Multiple stores can be occurring to the same media simultaneously. The number of streams to a device is configurable. (This matches sdisk)
- A copy number is assigned to each Lattus media (namespace) when it is configured. When files are stored, only copies of the indicated number will go to a specific media (namespace).

- This copy number can be changed by an administrator when the media (namespace) is blank.
- By definition, if files are to have multiple copies stored to Lattus media (namespace), then a Lattus media (namespace) per copy number will have to be configured.
- If there are multiple media (namespaces) configured for the same copy number, then at store time the media with the most available space will be selected for use.
- If a set of media become full, new ones can be configured and assigned whatever copy numbers are needed or desired. This is assuming there is another set of storage nodes with space available.

Wide Area Storage Features in the StorNext (GUI)

On the **Home** page of the GUI, the **Wide Area Storage, I/O Path** status and **Used Space** indicator are available for monitoring purposes (see [Figure 164](#) below).

Figure 164 Home Page




On the **Configuration > Licenses** page of the GUI, the **SNSM Wide Area Storage** license row is available for you to access Wide Area Storage feature (see [Figure 165](#) below).

Figure 165 Configuration > Licenses Page

Status	Feature	Limit	Expires	Description
<input checked="" type="checkbox"/>	SAN Client	65535	2012-12-13 23:59:59 MST	You must have a SAN client license for each client connected to StorNext over a SAN.
<input checked="" type="checkbox"/>	LAN Client	65535	2012-12-13 23:59:59 MST	You must have a LAN client license for each LAN client you use with StorNext. (In addition to any SAN clients).
<input checked="" type="checkbox"/>	Storage Manager	10 TB	2012-12-13 23:59:59 MST	A Storage Manager license provides full access to StorNext's Storage Manager Features that are not licensed separately.
<input checked="" type="checkbox"/>	Replication	unlimited	2012-12-13 23:59:59 MST	A Replication license is required if you want to use StorNext's Data Replication Feature.
<input checked="" type="checkbox"/>	Deduplication			A deduplication license is required if you want to use StorNext's Data deduplication (blockpool) feature.
<input checked="" type="checkbox"/>	Vaulting	unlimited	2012-12-13 23:59:59 MST	A vaulting license provides the ability to move seldom-used media to a manual archive vault, freeing room for media in the managed archives.
<input checked="" type="checkbox"/>	Storage Disk	unlimited	2012-12-13 23:59:59 MST	You must have a storage disk license to be able to configure and use StorNext storage disks.
<input checked="" type="checkbox"/>	Checksum	unlimited	2012-12-13 23:59:59 MST	A checksum license enables you to verify data integrity by ensuring that the checksum created when data was stored matches the checksum upon data retrieval.
<input checked="" type="checkbox"/>	Distributed Data Mover (DDM)	32	2012-12-13 23:59:59 MST	A license is required for using Distributed Data Mover if you plan to use additional machines besides the primary server to distribute data streams.
<input checked="" type="checkbox"/>	Failover (HA)	unlimited	2012-12-13 23:59:59 MST	A failover (High Availability) license is required if you plan to use StorNext's HA failover features.
<input checked="" type="checkbox"/>	Maintenance	0	2012-12-13 23:59:59 MST	A Maintenance license verifies that your site has purchased StorNext upgrade licenses, and is required for StorNext upgrades. It is also used at run time to verify the StorNext version in the software matches what was purchased.
<input checked="" type="checkbox"/>	Archive Conversion			An Archive Conversion license is required if you plan to read and migrate data from non-StorNext archive media into the StorNext file system.
<input checked="" type="checkbox"/>	Gateway Licensed WDC			A Gateway license allows a gateway to connect to StorNext servers that do not limit LAN clients connections. This license also allows the gateway to connect to StorNext servers regardless of File System license limits.
<input checked="" type="checkbox"/>	SNM wide Area Storage	unlimited	2012-12-13 23:59:59 MST	A SNM wide Area Storage license enables you to access wide Area Storage Features.

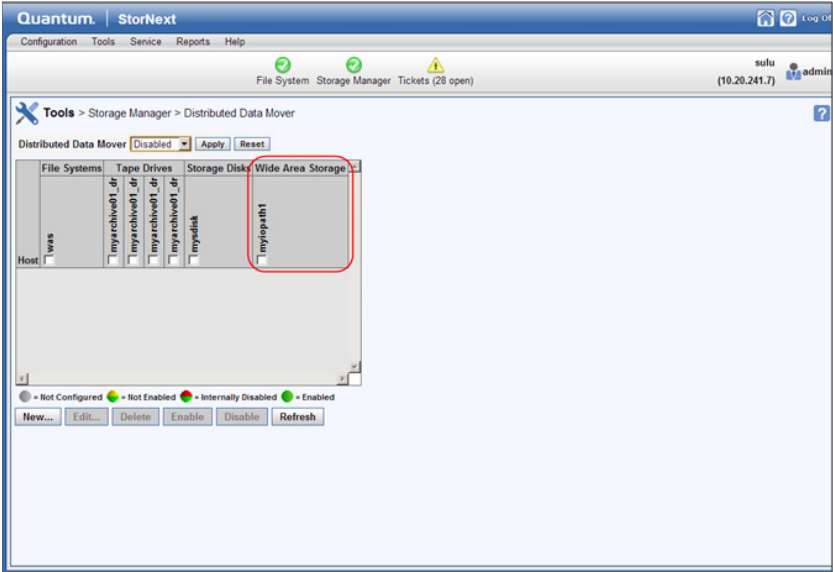
On the **Configuration > Storage Manager Policies** page of the GUI, support for LATTUS media type when configuring a SM policy is available (see [Figure 166](#) below).

Figure 166 Configuration >
Storage Manager Policies >
Edit Page



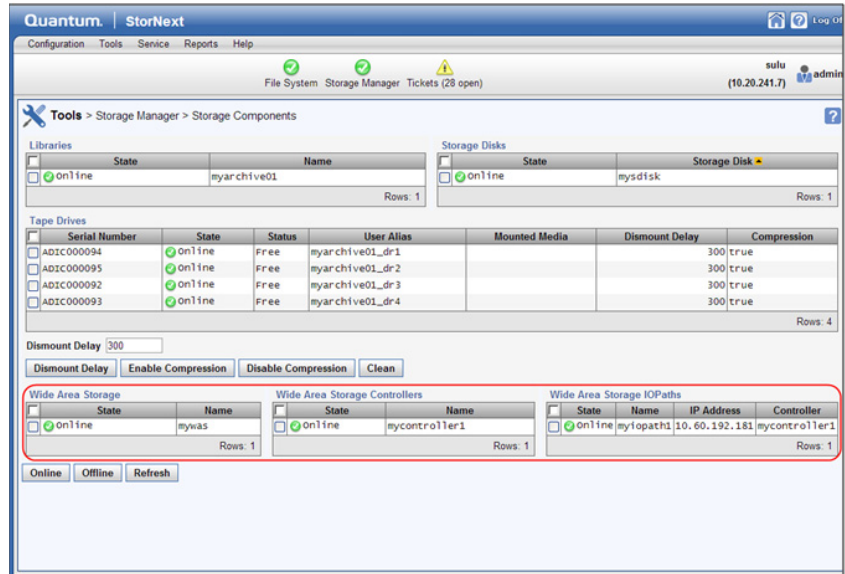
On the **Tools > Storage Manager > Distributed Data Mover** page of the GUI, the **Wide Area Storage** I/O path to DDM configuration is available (see [Figure 167](#) below).

Figure 167 **Tools > Storage Manager > Distributed Data Mover** Page



On the **Tools > Storage Manager > Storage Components** page of the GUI, you can toggle **WAS**, **WAS Controllers**, and **WAS I/O Paths** online/offline (see [Figure 168](#) below).

Figure 168 **Tools > Storage Manager > Storage Components Page**



On the **Tools > Storage Manager > Media Actions** page of the GUI, the **LATTUS** media type is available for you to remove, purge, assign policy, attributes and clean media operations (see [Figure 169](#) and [Figure 170](#) below).

Figure 169 Tools > Storage Manager > Media Actions (Media Selection) Page





Figure 170 Tools > Storage Manager > Media Actions (Available Media Actions list) Page



Configuring Lattus Object Storage

Within StorNext Storage Manager you must configure Tape Drives or Storage Disks before they can be used as storage destinations. The same is true for Lattus Object Storage. The Lattus Object Storage devices and media must be configured before they can be used by Storage Manager.

When configuring Lattus Object Storage for use, the typical order of events is:

- 1 Add the appliance.
- 2 Add the controller node.
- 3 Add the I/O path.
- 4 Add the namespace.

Note: Before you create a Lattus Object Storage device in Storage Manager, the namespaces, etc. you plan to use must reside in an existing Lattus Object Storage appliance prior to configuring in Storage Manager.

The `fsobjcfg` Command

The command used for configuring Storage Manager to make use of a Lattus Object Storage appliance is `fsobjcfg`. Via this command, the components that make up Storage Manager's view of Lattus Object Storage can be configured. These attributes include IP Addresses, TCP/IP port numbers, Network Protocol and namespace. These together with an object identifier, form the required URL to create/write/read/delete objects in the Lattus Object Storage appliance.

Refer to the *MAN Pages Reference Guide* for additional information on the `fsobjcfg` command, command syntax usage, and examples.

The usage output from the `fsobjcfg -h` command:


```
# fsobjcfg -h
Usage:
  Object Storage Report
  -----
  fsobjcfg [-l] [-F text | xml | json]

  Object Storage Appliance
  -----
  fsobjcfg -a -i ipaddress -p port_number [-e 'http' | 'https']
              [-U username -P password] [-B] [-v LATTUS]
              cloud_appliance_alias
  fsobjcfg -m [-i ipaddress] [-p port_number] [-e 'http' | 'https']
              [-U username [-P password]] [-B]
              cloud_appliance_alias
  fsobjcfg -d cloud_appliance_alias

  Object Storage Controller Node
  -----
  fsobjcfg -a -n controller_node_alias [-s streams] [-B] cloud_appliance_alias
  fsobjcfg -m [-s streams] [-B] -n controller_node_alias
  fsobjcfg -d -n controller_node_alias

  Object Storage IO Path
  -----
  fsobjcfg -a -o iopath_alias -i connection_endpoint [-B] [-e 'http' | 'https'] [-u 'VHOST' |
'PATH'] [-t 'AXR' | 'S3' ] -n controller_node_alias fsobjcfg -m -o iopath_alias [-i
connection_endpoint] [-B] [-e 'http' | 'https'] [-u 'VHOST' | 'PATH'] [-t 'AXR' | 'S3' ] -n
controller_node_alias
  fsobjcfg -m -o iopath_alias [-i connection_endpoint] [-B] [-e 'http' | 'https']
              [-u 'VHOST' | 'PATH'] [-t 'AXR' | 'S3' ] -n controller_node_alias
  fsobjcfg -d -o iopath_alias -n controller_node_alias

  Object Storage Namespace
  -----
  fsobjcfg -a -b namespace [-c copy] [-f media_id] [-U username -P password]
              [-t 'AXR' | 'S3' ] [-B] cloud_appliance_alias
  fsobjcfg -m [-b namespace] [-c copy] [-U username [-P password]]
              [-t 'AXR' | 'S3' ] [-B] -f media_id
  fsobjcfg -d -f media_id
  fsobjcfg -r -f media_id
```

How to Route Backup Files to Lattus for Easier Recovery

If you use Lattus Object Storage media for system backups, the procedure below instructs you how to create a special backup namespace and assign it to the backup policy.

Using the GUI, follow the procedure below to route backup files to Lattus Object Storage media and into backup-exclusive namespaces.

- 1 On the **Configuration** menu, click **Storage Destinations**, and then click the **Lattus Object Storage** tab.
- 2 Click **New...**
 - In the **Namespaces** section, click **Add** to create a Lattus Object Storage namespace for each copy number that is required for backup to Lattus Object Storage.
- 3 On the **Configuration** menu, click **Storage Manager Policies**, and then click the **Storage Manager** tab.
 - a Click the **_adic_backup Policy Class**, and then click **Edit...**
 - b In the **General** tab, for the **Default Media Type** option, click **Lattus** or **S3**.
 - c In the **Steering** tab, under the **Media Type** option, click **Lattus** or **S3** for the respective **Copy** numbers.
- 4 On the **Tools** menu, click **Storage Manager**, and then click **Media Actions**.
 - a In the **Available Media Actions** list, click **Assign Media to Policy Class**.
 - b In the **Media ID** column, click all the backup namespaces.
 - c In the **Assign Media to Policy Class Parameters** section, click **_adic_backup** in the **Destination Policy Class** list.

The MAX_STORE_SIZE System Parameter

Within Storage Manager, it is possible to set the maximum file size that will be stored automatically by the Storage Manager software. This is done by setting the system parameter **MAX_STORE_SIZE** in the file: `/usr/adic/TSM/config/fs_sysparm_override`. The file: `/usr/adic/TSM/config/fs_sysparm.README` contains a detailed description of all system parameters that can be adjusted including **MAX_STORE_SIZE**.

As indicated, this system parameter can be used to set the maximum size of files that are stored automatically. Specifically, this means that runs of **fspolicy** will recognize these files as being too large and will not store them. Additionally, they will be removed from the candidate lists so further policies will not even see these files as candidates. This is true of the **fspolicy** command whether it is run automatically by the Storage Manager software or run manually from the command line.

Note: Files larger than the maximum value can be stored manually by using the **fsstore** command; it does not check the system parameter.

When specifying the file size using **MAX_STORE_SIZE**, the value by default is the number of gigabytes (GB) in the file, if no suffix is given. That is, a value of 500 would be interpreted as 500 GB. Other suffixes are also available for specifying the size.

StorNext provides support for an object size of 16 TiB (bytes).

[Table 3](#) lists the supported suffixes for the **MAX_STORE_SIZE** parameter.

Table 3 Supported Suffixes

Unit of Measure	Value in bytes
B	Value in bytes
KB or KiB	Value (10 ³) or (2 ¹⁰)
MB or MiB	Value (10 ⁶) or (2 ²⁰)
GB or GiB	Value (10 ⁹) or (2 ³⁰)
TB or TiB	Value (10 ¹²) or (2 ⁴⁰)
PB or PiB	Value (10 ¹⁵) or (2 ⁵⁰)

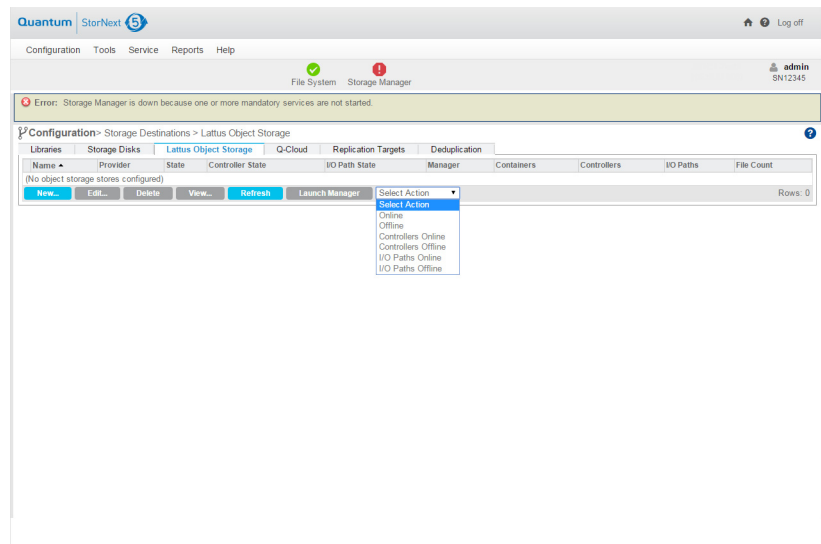
Refer to the `/usr/adic/TSM/config/fs_sysparm.README` file for the proper syntax of the **MAX_STORE_SIZE** system parameter.

Setting Up Lattus Object Storage Destinations

Within StorNext Storage Manager you must configure Tape Drives or Storage Disks before they can be used as storage destinations. The same is true for Lattus Object Storage. The Lattus Object Storage devices and media must be configured before they can be used by Storage Manager. See [Configuring Lattus Object Storage](#) on page 420.

On the **Configuration > Storage Destination** page of the GUI, the **Lattus Object Storage** tab is available for you to perform various functions (see [Figure 171](#) below).

Figure 171 **Configuration > Storage Destinations > Lattus Object Storage**



Note: Your GUI may differ slightly depending on your configuration.

Setting Up Lattus Object Storage Destinations on a StorNext Lattus Configuration

The Configuration menu's Storage Destinations option enables you to view, add, edit or delete Lattus Object Storage destinations.

Viewing Lattus Object Storage Destinations

Follow this procedure to view a list of currently configured Lattus Object Storage destinations.

- 1 Click **Storage Destinations** from the **Configuration** menu.
- 2 Click the **Lattus Object Storage** tab. Information for any previously configured Lattus Object Storage destinations is shown. For each configured destination, the screen displays the **Name**, **Provider**, **State** (online or offline), **Controller State**, **I/O Path State**, **Manager** host address, **Containers** count, **Controllers** count, **I/O Paths** count, and **File Count**.
- 3 Select the **Lattus Object Storage** destination whose information you want to view.
- 4 Click **View...**
- 5 When you are finished viewing library information, click **Done**.

Adding a New Lattus Object Storage Destination

Follow this procedure to add a new **Lattus Object Storage** destination.

Note: If you plan to use HTTPS, you must create or import a security certificate prior to creating a **Lattus Object Storage** destination. This applies to each Lattus controller configured for HTTPS. If you only intend to use HTTP, certificates are not needed. To **create** or **import** a Lattus Object Storage security certificate, see [Lattus Certificates](#) on page 272. To manage SSL certificates, click **Lattus Certificates** on the **Tools** menu. For additional information, see [Basic Secure Sockets Layer \(SSL\) Guidelines](#) on page 463, as it outlines some standard information about using private and public certificates.

- 1 If you have not already done so, click **Storage Destinations** from the **Configuration** menu.
- 2 Click the **Lattus Object Storage** tab.
- 3 Click **New...**
- 4 Enter the appropriate value into the following parameters:

Note: Parameters marked with an asterisk (*) are required.

Parameter	Description
*Name	Enter the name of the new Lattus Object Storage destination.
*Provider (<i>New in StorNext 5 release 5.1</i>)	Select Lattus .
*Manager Host	Enter the host address for the Lattus Object Storage manager host.
*Manager Port	Enter a decimal integer to specify the port number of the Lattus Object Storage manager GUI interface. The default port number is 80 .
Manager Protocol	Select the http or https protocol. (See https Note : on page 425.)
Authentication	Select if authentication is required for this configuration.
User Name	<p>Select a global username to be used for namespace permission, for this configuration. This parameter is mandatory if Authentication is set to "Enabled".</p> <p>Note: To prevent snpolicyd and Storage Manager from occasionally failing and perhaps taking down the iopath and namespace, do not configure User Name / Password for namespaces assigned to snpolicyd and Storage Manager usage.</p>

Parameter	Description
Password	<p>Select a global password to be used for namespace permissions, for this configuration. This parameter is mandatory if Authentication is set to "Enabled".</p> <p>Note: To prevent snpolicyd and Storage Manager from occasionally failing and perhaps taking down the iopath and namespace, do not configure User Name / Password for namespaces assigned to snpolicyd and Storage Manager usage.</p>

- 5 In the **Controllers** section, click **Add** or **Add Controller**, and then specify the following criteria:

Note: You must **create** or **import** a Lattus Object Storage security certificate prior to creating a **Lattus Object Storage Destination**. It applies to each Lattus controller configured for HTTPS. If you only intend to use HTTP, certificates are not needed. To **create** or **import** a Lattus Object Storage security certificate, see [Lattus Certificates](#) on page 272. To manage SSL certificates, click **Lattus Certificates** on the **Tools** menu. For additional information, see [Basic Secure Sockets Layer \(SSL\) Guidelines](#) on page 463, as it outlines some standard information about using private and public certificates.

Parameter	Description
Name	Enter the name of the controller.

Parameter	Description
Max Streams	The Max Streams value defines the number of concurrent I/O operations that can write concurrently to the controller. By default, the maximum number of streams is set to 48, or you can select another value from the Max Streams drop-down list.

6 In the **I/O Paths** section, click **Add** or **Add I/O Path**, and then specify the following criteria:

Parameter	Description
Name	Enter the unique name of the I/O path. If you do not have unique names, the "Already exists in the Tertiary Manager system. Duplicate component alias names are not allowed." error message is displayed.
Controller Name	Select the name of the controller associated with the new I/O path.
REST API Type (<i>New in StorNext 5 release 5.1</i>)	The StorNext GUI requires an API to query for namespaces/buckets on a particular host. This parameter specifies which API to use for the particular host. The available values are AXR and S3 .
URL Style (<i>New in StorNext 5 release 5.1</i>)	There are two ways to format the URL: <ul style="list-style-type: none"> • PATH • VHOST This parameter defines which style of URL to use.

Parameter	Description
Object Access Protocol (<i>New in StorNext 5 release 5.1</i>)	This specifies the network protocol to be used for the host. Select the protocol to be used for Lattus Object Storage object access. By default, the protocol is set to http . (See https Note : on page 425.)
Host[:Port] (<i>New in StorNext 5 release 5.1</i>)	Enter the connection endpoint address that contains host name or IP address with the optional port number separated by a colon ":". If the port number is not specified, the default (80 for http, 443 for https) is assumed. Connection endpoints must be unique.

7 In the **Containers** section, perform one of the following:

- a On the **Container Selection** list, click **Scan** or **Manual**. This parameter gives you the option to either scan the available container or enter the container name manually. If you select **Scan** and you need to specify user name and password, you can either use the credential specified for the manager host or a different credential. If you use a different credential, check the **Use different credentials** box, and enter user name and password. You are then presented with a pre-populated list of available containers. If you select **Manual**, you are presented with a text box to manually enter the name of the container. To view or add S3 buckets, on the **Tools** menu, click **S3 Buckets**.
- b Click **Add** or **Add Container**, and then specify the following criteria:

Parameter	Description
Container	Select (Scan mode) or enter (Manual mode) the appropriate container for this configuration.

Parameter	Description
Media ID	Enter the StorNext Media ID associated with the selected container. The Media ID must be unique.
REST API Type (<i>New in StorNext 5 release 5.1</i>)	The StorNext GUI requires an API to query for namespaces/buckets on a particular host. This parameter specifies which API to use for the particular host. The available values are AXR and S3 .
Authentication	Select if authentication needs be enabled for this container. If the authentication is disabled, the global username and password will be used, if applicable.
User Name	<p>Select a username to be used to access this container. This parameter is mandatory if Authentication is set to "Enabled". This selection overrides the global permissions settings.</p> <p>Note: To prevent snpolicyd and Storage Manager from occasionally failing and perhaps taking down the iopath and namespace, do not configure User Name / Password for namespaces assigned to snpolicyd and Storage Manager usage.</p>

Parameter	Description
Password	<p>Select a password to be used to access the container. This parameter is mandatory if Authentication is set to "Enabled". This selection overrides the global permissions settings.</p> <p>Note: To prevent snpolicyd and Storage Manager from occasionally failing and perhaps taking down the iopath and namespace, do not configure User Name / Password for namespaces assigned to snpolicyd and Storage Manager usage.</p>
Copy Number	<p>Select the copy number (1-4) assigned to the container. The copy number can be changed if no data has been written to the media.</p>

Note: If no data has been written to a controller, I/O path, or container, you can click **Delete** to remove the item, and then click **Apply** to save the changes

- 8 Click **Apply** to save your changes, or **Cancel** to exit without saving.
- 9 Repeat **Step 2** through **Step 7** to add additional Lattus Object Storage.

Note: The containers on Lattus share the same I/O paths and storage capacity. There is no advantage to define multiple containers for the same Policy Class and Copy number. Storage Manager selects the first available container that meets the policy class criteria for store operation.

- 10 Add a Storage Manager or replication storage policy (see [Step 7: Storage Policies](#) on page 69), and then continue with the standard procedures according to [The Configuration Wizard](#) on page 23 (see [Chapter 3, The Configuration Wizard](#)).

Editing a Lattus Object Storage Destination

Follow this procedure to edit an existing Lattus Object Storage destination.

- 1 If you have not already done so, click **Storage Destinations** from the **Configuration** menu.
- 2 Click the **Lattus Object Storage** tab.
- 3 Select the Lattus Object Storage destination whose information you want to edit.
- 4 Click **Edit...**
- 5 To edit a field, type directly in the field (for example, type a new name and IP address for an I/O path), or select another option from the drop-down list.

Note: To return to the last saved configuration for a controller, I/O path, or namespace, click **Reset**.

- 6 Click **Apply** to save your changes, or **Cancel** to exit without saving.
- 7 When a confirmation message appears, click **Yes** to proceed, or **No** to abort.
- 8 After a message informs you that the Lattus Object Storage destination was successfully modified, click **OK**.

Deleting a Lattus Object Storage Destination

Follow this procedure to delete an existing Lattus Object Storage destination.

- 1 If you have not already done so, click **Storage Destinations** from the **Configuration** menu.
- 2 Click the **Lattus Object Storage** tab.
- 3 Select the Lattus Object Storage destination you want to delete.
- 4 Click **Delete**.

- 5 When a confirmation message appears, click **Yes** to proceed with the deletion or **No** to abort.
- 6 After a message informs you that the storage disk was successfully deleted, click **OK**.

Performing Other Lattus Object Storage Destination Actions

Follow this procedure to launch the Lattus Object Storage manager GUI application.

- 1 If you have not already done so, click **Storage Destinations** from the **Configuration** menu.
- 2 Click the **Lattus Object Storage** tab.
- 3 Select the Lattus Object Storage destination.
- 4 Click **Launch Manager**. A new browser window appears, and displays the Lattus Object Storage manager GUI application login screen. If you entered a **User Name** and **Password** when you created the selected Lattus Object Storage destination, then the credentials are used as your login.

Note: If you are using Safari as your browser, you may have to enable pop-ups. See [How to Enable Pop-ups in Safari](#).

How to Enable Pop-ups in Safari

- 1 Open Safari if it is not already open.
- 2 On the **Safari** menu, click **Preferences**.
- 3 Click the **Security** heading.
- 4 Un-check (turn off) the box marked **Block pop-up windows** to allow pop-ups. Safari will then ask if you would really like to change the setting.
- 5 Click **OK**.
- 6 Close the **Preferences** window.
- 7 Shut down and restart Safari.

If you block pop-up windows, you might miss important information for a web page. For example, the **Launch Manager** might use a pop-up window to request your login credentials.

Changing the Current State of Lattus Object Storage Destinations, Controllers, and I/O Paths

You can also change the current state of Lattus Object Storage destinations, controllers, and I/O paths. To change the state, select the Lattus Object Storage destination, and then choose one of these options from the **Select Action** drop-down list:

Parameter	Description
Online	Select this option to set the Lattus Object Storage destination online.
Offline	Select this option to take the Lattus Object Storage destination offline.
Controllers Online	Select this option to set the controllers online.
Controllers Offline	Select this option to take the controllers offline.
I/O Paths Online	Select this option to set the I/O Paths online.
I/O Paths Offline	Select this option to take the I/O Paths offline.

Special Considerations for Multi-Geo Configurations

A Multi-Geo (multiple geographic) Lattus configuration consists of three sites configured under the same durability policy. It is likely that WAN (Wide Area Network) communication with remote sites will be slower due to higher latency in the WAN link. If you have significantly higher latency to the remote Lattus sites, it is recommended that only the I/O Paths to the local controller be **“Online”**.

Object Storage I/O Paths can be configured offline with the **Tools > Storage Manager > Storage Components** screen. Select the remote Object Storage I/O Paths and then click the Offline button ([Changing the Current State of Lattus Object Storage Destinations, Controllers, and I/O Paths](#) on page 434). The `fschstate(1)` command may also be used for this.

If the local Lattus controller is down, but the remote sites are still up, then you may want to change the local I/O Paths to the “Offline” state and the remote I/O Paths to “Online” state in order to continue using Lattus.

HTTPS Support for Lattus Object Storage

Storage Manager provides both HTTP and HTTPS support in this release.

The `fsobjcfg` command provides the capability to specify a list of connection endpoints for http or https. Each connection endpoint consists of an IP address (or a DNS hostname) and a port number. These connection endpoints must match the configuration of the Lattus system. Storage Manager can be configured to verify either PEER only, or both PEER and HOST when https is used.

You can provide either a file path name to a CA Certificate or a directory path name where CA Certificates are deposited. The CA Certificate file path name has no default. The default CA Certificate directory is `/usr/cvfs/config/ssl`. This directory path is automatically created on the MDCs. However, the Administrator is responsible for creating this directory on each DDM host.

Note: Beginning in StorNext 5 release 5.2, `/usr/cvfs/config/ssl` is no longer the default repository that is referenced by Storage Manager for SSL certificates when using HTTPS. The default Certificate file or repository will depend on the OS vendor. For additional information, see [HTTPS Default CA ROOT Certificate File or Path](#) and [How to Update Expired CA Root Certificates](#).

All of these attributes can be modified in the `fs_sysparm_override` file in `/usr/adic/TSM/config`, using the following parameters:

The FS_OBJSTORAGE_CACERT System Parameter

This parameter contains the full path name of a Certificate file. Both `FS_OBJSTORAGE_CAPATH` and `FS_OBJSTORAGE_CACERT` should not be set at the same time. If it is, `FS_OBJSTORAGE_CACERT` is used.

The FS_OBJSTORAGE_CAPATH System Parameter

This parameter contains the directory path where all the certificates reside.

The FS_OBJSTORAGE_SSL_VERIFY_PEERHOST System Parameter

The value assigned to this parameter determines Peer only or Peer and Host verification. A value of 1 means Peer only verification and a value of 2 will force Peer and Host verification.

Refer to the `/usr/adic/TSM/config/fs_sysparm.README` file for the proper syntax of these system parameters.

Future StorNext releases will provide GUI support for managing these attributes.

Configuring HTTPS on DDM Hosts

Configuring HTTPS SSL certificates on DDM Hosts is a manual process. The general high level process consists of the following three steps:

- 1 Create the directory to hold the certificates.
- 2 Move the Public Certificates to the directory from Step 1.
- 3 Run `c_rehash` on the directory from Step 1.

For the entire StorNext Cluster, the directory that holds the SSL certificates is identified by the system parameter `FS_OBJSTORAGE_CAPATH`. By default, this system parameter is set to `/usr/cvfs/config/ssl`. The script `c_rehash` is installed on each DDM Hosts as `/opt/quantum/openssl/bin/c_rehash`. `c_rehash` first deletes all existing symbolic links, and then creates new symbolic links for all files containing the file extension `.pem`.

Note: StorNext only supports certificates in `.pem` format. For additional information, see [Basic Secure Sockets Layer \(SSL\) Guidelines](#) on page 426, as it outlines some standard information about using private and public certificates.

Examples

Below are examples of the following:

- [Initial Setup of the SSL Directory and Certificates on the DDM Host](#)
- [Updating an Existing Certificate on the DDM Host](#)
- [Adding a New Certificate on the DDM Host](#)

Initial Setup of the SSL Directory and Certificates on the DDM Host

- 1 Log into the DDM Host.
- 2 Create the directory to hold the certificates.

```
mkdir /usr/cvfs/config/ssl
```

- 3 Move the Public Certificates to:
 - a The directory created in Step 2

```
scp root@MDC-host:/usr/cvfs/config/ssl/*.pem /usr/cvfs/config/ssl
```

- b Or via any other mechanism that you prefer to load the DDM Host /usr/cvfs/config/ssl with all your certificates.

- 4 Run `c_rehash` on the directory created in Step 2.

```
/opt/quantum/openssl/bin/c_rehash /usr/cvfs/config/ssl
```

Updating an Existing Certificate on the DDM Host

```
rm /usr/cvfs/config/ssl/cert.pem
```

```
scp root@MDC-host:/usr/cvfs/config/ssl/cert.pem /usr/cvfs/config/ssl
```

```
/opt/quantum/openssl/bin/c_rehash /usr/cvfs/config/ssl
```

Adding a New Certificate on the DDM Host

```
scp root@MDC-host:/usr/cvfs/config/ssl/cert.pem /usr/  
cvfs/config/ssl  
/opt/quantum/openssl/bin/c_rehash /usr/cvfs/config/ssl
```

HTTPS Support for Q-Cloud Archive™ Object Storage

Q-Cloud Archive™ destinations are only accessible via HTTPS. These destinations requires Verisign Root Certificates.

CA Root Certificates are packaged and distributed by OS vendors. By default, these certificates already exist on your system. The location of these certificates may be different, depending on your OS vendor.

For additional information, see [HTTPS Default CA ROOT Certificate File or Path](#) and [How to Update Expired CA Root Certificates](#).

Changes to Existing CLI Commands

Some existing SM commands had to be updated for use with the new Lattus destination. Most of the updates to these commands are related to the object IDs that are now assigned to files that are stored to Lattus. It is the object ID that is used to identify a file segment in the Lattus. With this object ID and the namespace (determined via the media ID) a file can be accessed directly in the Lattus. The following commands have new options related to the new storage destination.

The `fsfileinfo` Command

The `-o` option was added to this command. When the new option is specified the object IDs for the file are displayed. Along with the object ID, the copy number, segment offset, and length are also displayed.

Note: For a multi-segment file, all object IDs are displayed. Also, if there are old versions of the file, only the object IDs for the current version are displayed.

The `fsmedinfo` Command

This command was updated so that when the `-l` option is specified the object IDs are listed with the rest of the segment information. In addition the new options `-s` and `-e` were added and can be used in combination with the `-l` option. When the `-s` option, starttime, or the `-e` option, endtime, is specified that will limit the file segments that are reported. By default when the `-l` option is used all file segments on a media are reported. If the starttime and/or the endtime are provided then only the segments with a time in the indicated time range are reported.

The `fsmedread` Command

This command had a new option `-u` added for reading a file from Lattus media. With the `-u` the URL of the file on the Lattus media is specified and will be used for reading the file.

The `fsmedscan` Command

This command was updated to fail if an attempt is made to run against a Lattus media. Due to the nature of the media and files stored there, no scanning of the media can be done to determine contents.

The `dm_info` Command

This command is an administrative tool and should not normally be used without Quantum assistance. The command output was updated to display any object IDs associated with the file if present. Additionally, the `-o` option was added and when it is used the command will only report the object ID information.

The `dm_util` Command

This command is an administrative tool and should not normally be used without Quantum assistance. With this command it is possible to update information stored in the extended attributes of a file on disk. When the `-u` option is used with an attribute type and value, that

attribute will be updated for the file. The new attribute type `objid` was added and when it is specified the object IDs in the inode for the file can be updated. In addition to updating the object ID values, it can be used for deleting object ID information.

Other Changes and Considerations

- StorNext provides support for HTTP connections to the Lattus Object Store, and HTTPS connections.
- StorNext provides support for an object size of 16 TiB (1024^4 bytes).
 - Support for 16 TiB objects (refer to the AmpliStor 3.0.3 Release Notes).
 - Support for 16TiB objects from StorNext to AmpliStor requires special guidelines. There are a couple of potential issues you could encounter if the guidelines are not followed.
 - For large object support on AmpliStor requires a policy with 256 MiB superblocks. AmpliStor does not support an object size larger than 16 TiB and 65536 superblocks.
 - For large object support in StorNext to archive an object 16 TiB in size requires a setting in `usr/adic/TSM/config/fs_sysparm_override` and setting **MAX_STORE_SIZE**. For additional information, refer to the `/usr/adic/TSM/config/fs_sysparm.README` file for the proper syntax of the **MAX_STORE_SIZE** parameter. See [Wide Area Storage Segment Size](#) on page 440 for additional information.

Wide Area Storage Segment Size

When a file is stored to WAS Storage, the segment size can impact how the file's content is populated. For large files, if the segment size is

configured, the file is broken down into multiple segments and each segment is stored as an object in WAS storage.

Segment size should be configured if very large files exist, since the WAS storage could place a limit on the max object size it can store. For example, Lattus 3.0.0 has a limit of 16 TB for each object. If no segment size is configured, uploading a file whose size is larger than 16 TB to Lattus storage (version 3.0.0) will fail.

snpolicyd

For snpolicyd, there are two parameters in different configuration files. The parameter `max_seg_size` in `was.conf` applies to all namespaces that are associated with the same WAS config ID. While the snpolicy user-defined policy also has the policy parameter `was_seg_size`. If `was_seg_size` is defined in a policy, its value is used as the segment size regardless if `max_seg_size` is defined in WAS config. This offers the flexibility to specify different segment sizes for different policies on the same WAS config.

Note: The segment size is always rounded up to next power of 2.

For Lattus, the segment size is pre-configured to 64 GB. Once the segment size is configured and applied to files that are uploaded to WAS storage, the size cannot be changed. Otherwise, the objects for files that were uploaded before cannot be retrieved and cannot be deleted.

There is a trade-off to selecting a larger or smaller segment size. Below are some of the advantages and disadvantages for different segment sizes. The selection relies on the system configuration, workload characteristics, application requirements, and other parameters.

- A larger segment size reduces the number of segments for a large file. An object ID is assigned for each segment and an entry for each object ID is added to the file's metadata. Therefore, a larger segment size will reduce the number of metadata entries, thus reducing the size of metadata consumed. It also reduces the space to store and transfer such metadata which is contained in the manifest file that snbackup saves to Lattus. Additionally, it could be easier for third-party applications to operate if there are fewer objects per file (assuming the object IDs are also exported to the third-party).

- On the other hand, with a smaller segment size, it could benefit significantly if there is an instability issue in the network or storage system. Each read/write failure will waste all of the effort for uploading/retrieving an object and have to retry. An object of larger size is more likely to result in a failure. A smaller size could speed up the completion time of uploading a large file, since `snpolicyd` starts to upload a segment once a full segment is written to a staging file system. If a large file contains only one large segment, the upload starts while the whole file is written to the staging file system. A smaller segment size also benefits the retrieval of content from WAS storage, if the file is truncated and partial retrieval is needed. The minimum retrieval is an object (size of a segment) so if just one object is needed to retrieve, a smaller segment size will be much faster and requires fewer resources.

Storage Manager

For Storage Manager, `MED_SEG_OVER_LATTUS` is the system parameter which controls the segment size for files targeted for storage in Lattus media. The default size is 128 GiB (137,438,953,472 bytes). Storage Manager will segment files greater than `MED_SEG_OVER_LATTUS`.

There is another system parameter, `MAX_STORE_SIZE`, which limits the maximum size of a file that can be stored to 2 TB (2,000,000,000,000 bytes).

The considerations described for `snpolicyd` should be used when deciding upon a segment size, except for the following three considerations, with respect to small segment sizes:

- Storage Manager does not use segment size to determine when to start storing files.
- Storage Manager does not support partial file retrieve to the original file so this would not be a consideration. Although, a related performance consideration would be that Storage Manager will store and retrieve segments concurrently, which can reduce the overall storage and retrieval time.
- In the event of an error, recovering from a smaller segment size can also be beneficial.



Chapter 12

Customer Assistance

More information about StorNext is available on the Quantum Service and Support website at <http://www.quantum.com/ServiceandSupport>. The Quantum Service and Support website contains a collection of information, including answers to frequently asked questions (FAQs).

StorNext Upgrades

To request a StorNext software upgrade for non-Quantum MDCs , visit <http://www.quantum.com/ServiceandSupport/Upgrade/Index.aspx>. To request a StorNext software upgrade for StorNext Appliances, open a support ticket at: <https://onlineservice.quantum.com/>. For further assistance, or if training is desired, contact the Quantum Technical Assistance Center.

Contacts

Quantum company contacts are listed below.

Quantum Home Page

Visit the Quantum home page at:

<http://www.quantum.com>

Comments

To provide comments or feedback about this document, or about other Quantum technical publications, send e-mail to:

doc-comments@quantum.com

Getting More Information or Help

StorageCare™, Quantum's comprehensive service approach, leverages advanced data access and diagnostics technologies with cross-environment, multi-vendor expertise to resolve backup issues faster and at lower cost.

Accelerate service issue resolution with these exclusive Quantum StorageCare services:

Quantum. Global Services

- **Service and Support Website** - Register products, license software, browse Quantum Learning courses, check backup software and operating system support, and locate manuals, FAQs, firmware downloads, product updates and more in one convenient location. Benefit today at:

<http://www.quantum.com/ServiceandSupport/Index.aspx>

- **eSupport** - Submit online service requests, update contact information, add attachments, and receive status updates via email. Online Service accounts are free from Quantum. That account can also be used to access Quantum's Knowledge Base, a comprehensive repository of product support information. Sign up today at:

<https://onlineservice.quantum.com/>

For further assistance, or if training is desired, contact the Quantum Customer Support Center:

United States	1-800-284-5101 (toll free) +1-720-249-5700
EMEA	+800-7826-8888 (toll free) +49-6131-3241-1164
APAC	+800-7826-8887 (toll free) +603-7953-3010


For worldwide support:

<http://www.quantum.com/ServiceandSupport/Index.aspx>

Worldwide End-User Product Warranty

For more information on the Quantum Worldwide End-User Standard Limited Product Warranty:

<http://www.quantum.com/serviceandsupport/warrantyinformation/index.aspx>



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*.

Gateway Server/Client Network and Memory Tuning

Using the Gateway Server and Client feature places significant additional demands on network capacity and system memory. Before creating and using a Gateway Server and Client, review the following information:

- [Gateway Server and Client Network Tuning](#)
- [Gateway Server Memory Tuning](#)

Note: For additional information about LAN client and server performance tuning, see the *StorNext File System Tuning Guide*.

Gateway Server and Client Network Tuning

Due to significant demands placed on the network, the following network issues can occur when using Gateway 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 Gateway Server. In this case a 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 Gateway Server, the 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 Gateway Server and client feature, and will drop packets. This causes TCP retransmissions, resulting in a significant performance loss. This can be observed as an increase in the Segments Retransmitted count in `netstat -s` output during LAN client write operations and Gateway Server read operations.

To address this issue, edit the `/usr/cvfs/config/dpserver` configuration file and reduce the Gateway 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.

Gateway Server Memory Tuning

The minimum amount of memory required for a Gateway Server depends on the configuration.

- **Windows.** In addition to the minimum OS memory requirements, an additional 1GB of memory must be available for each file system the Distributed LAN Gateway will serve.
- **Linux.** For a Linux Gateway Server, use the following formula:

Required memory = 1GB +
(# of file systems served
* # of NICs on the Gateway Server used for
Distributed LAN traffic
* server buffer count
* transfer buffer size)

For example, consider a Linux Gateway 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 `sndpdcfg` 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 Gateway Servers contain a minimum of 2GB of RAM.

Configuring LDAP

This section 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.

Configuring a Data Partition for Use with Spectra Logic T-series Tape Storage Libraries

Configuring the Partition


Use the following steps to configure a data partition for use with StorNext.

Important: Make sure that the exporting QIP and the drives in the partition are connected and configured on the same network as the StorNext server.

- 1 Log in as a user with superuser or administrator privileges.
To change user types, select **Security > Switch User** from the toolbar menu, and then log in as a user with the required privileges.
- 2 Configure the library network settings as described in the *Configuring Network Settings* section of the user's guide for the library.
- 3 Create or modify the partition to be used with StorNext. Refer to the *Using Partitions* section in the library user's guide for detailed information about configuring and using partitions.

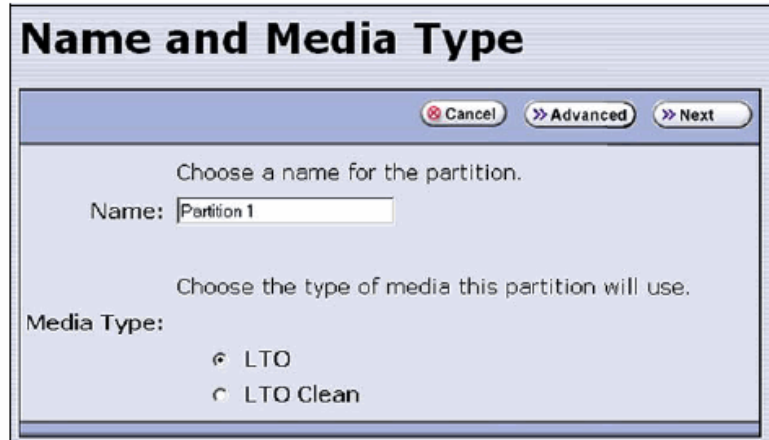
Note: StorNext treats each data partition as an independent library.

- a From the toolbar menu, select **Configuration > Partitions** to display the Shared Library Services screen.



The Shared Library Services screen.

- b Select **New** to create a new partition or **Edit** to modify an existing partition. The Name and Media Type screen appears.

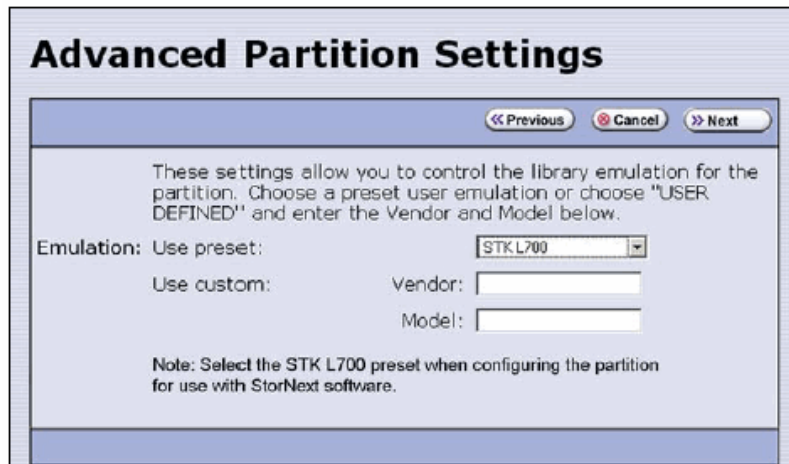


The dialog box is titled "Name and Media Type". At the top right, there are three buttons: "Cancel", "Advanced", and "Next". The main content area contains the following text and controls:

- Text: "Choose a name for the partition."
- Label: "Name:" followed by a text input field containing "Partition 1".
- Text: "Choose the type of media this partition will use."
- Label: "Media Type:"
- Two radio button options: "LTO" (which is selected) and "LTO Clean".

The Name and Media Type screen.

- c Enter a name for the partition and select the type of media that it will use.
- d Select **Advanced** to display the Advanced Partition screen.



The dialog box is titled "Advanced Partition Settings". At the top right, there are three buttons: "Previous", "Cancel", and "Next". The main content area contains the following text and controls:

- Text: "These settings allow you to control the library emulation for the partition. Choose a preset user emulation or choose 'USER DEFINED' and enter the Vendor and Model below."
- Label: "Emulation: Use preset:" followed by a drop-down menu showing "STK L700".
- Label: "Use custom:" followed by "Vendor:" and a text input field.
- Label: "Model:" followed by a text input field.
- Note: "Note: Select the STK L700 preset when configuring the partition for use with StorNext software."

Specify the emulation mode.

- e Select **STK L700** from the **Use preset** drop-down list.

- f Click **Next**. The Robotic Control Path screen displays listing the QIPs and/or direct-attach drives currently installed in the library.

This is the beginning of the series of configuration screens. The sequence of steps and screens matches those described in the *Configuring a New Data Partition* section of the library user's guide.

Refer to the documentation at <http://docs.sun.com/app/docs/prod/L700.tape> for information about the StorageTek L700 library.

Basic Secure Sockets Layer (SSL) Guidelines

If you are working on a Lattus system that already has an existing SSL certificate, this section outlines what you need to do to get the public portion of that certificate onto a StorNext MDC to be used for secure https transfers.

Note: StorNext only supports certificates in PEM format.

This section provides guidelines on how to use the PEM (Privacy Enhanced Mail) file that already exists on your Lattus system. A typical PEM file will look like the server.pem illustrated in [Example of a server.pem](#) [File](#) on page 464.

- The PEM file is a clear text file which contains both a private and public SSL certificate.
- The private portion of the PEM file begins with the text "-----BEGIN RSA PRIVATE KEY-----" and ends with the text "-----END RSA PRIVATE KEY-----". Below is an example of a PEM file containing 4 public certificates and 1 private certificate.


```
-----BEGIN CERTIFICATE-----  
(Your Primary SSL certificate: your_domain_name.crt)  
-----END CERTIFICATE-----  
-----BEGIN CERTIFICATE-----  
(Your Intermediate certificate: SomeCA.crt)  
-----END CERTIFICATE-----  
-----BEGIN CERTIFICATE-----  
(Your Root certificate: SomeRoot.crt)  
-----END CERTIFICATE-----  
This is a Certificate with Private and Public keys:  
-----BEGIN RSA PRIVATE KEY-----  
-----END RSA PRIVATE KEY-----  
-----BEGIN CERTIFICATE-----  
-----END CERTIFICATE-----
```

- The private portion of the PEM file should **NEVER** be transferred off the Lattus system in any format for use by a StorNext MDC, as the private portion of the PEM file is never needed by a StorNext MDC. This helps to ensure that the security on the Lattus system remains intact and is not jeopardized.

After you have identified where the PEM file is located, perform **Step 1** through **Step 5** below to create a public SSL certificate for use on a StorNext MDC:

- 1 Assume the name of your self-signed certificate is `server.pem` and that it contains both a private and public certificates. If your `server.pem` file only contains 1 public certificate, run the following command from a terminal to create a `public.pem` certificate file and then proceed to **Step 3**.

```
openssl x509 -in server.pem -out public.pem
```

- 2 If your `server.pem` file contains multiple public certificates, perform **Step 2(a)** through **Step 2(d)**:

- a Issue the following command on the terminal to make a copy of your private key (this file will become your public key):

```
cp server.pem public.pem
```

- b Open the `public.pem` file with your text editor of choice:

```
vi public.pem
```

- c In the above example of the `.pem` file, delete the lines beginning with text “-----BEGIN RSA PRIVATE KEY-----” and ending with “-----END RSA PRIVATE KEY-----”, inclusive.

Caution: The `public.pem` file should **NOT** contain any blank lines. If you edit the file, please verify there are no blank lines in the code. Blank lines in the `public.pem` file is not supported by the API used to import the file.

- d Save this `public.pem` file – the resulting file should look like the example in [Example of a public.pem File](#) on page 465.
- 3 Transfer the `public.pem` file to a place where the MDC’s GUI can access it.
- 4 On the **Tools** menu of the StorNext GUI, click **Lattus Certificates**. The **Tools > Lattus Certificates** page appears (see [Figure 101](#) in [Lattus Certificates](#) on page 272).
- 5 On the **Tools > Lattus Certificates** page, click **Import...** The **Import A Certificate** dialog box appears.
- 6 In the **Import A Certificate** dialog box, click **Choose File** to select a file to import. The **Open** dialog box appears. Alternatively, click **Close** to cancel the import.
- 7 In the **Open** dialog box, navigate to the `public.pem` certificate file you want to import, and then click **Open**.

Note: Public Certificate files uploaded through the GUI are placed in the following directory: `/usr/cvfs/config/ssl`

If the import is successful, the **Information** notification at the top of the **Tools > Lattus Certificates** page displays, "Certificate public.pem uploaded successfully."

Example of a server .pem File

```
-----BEGIN RSA PRIVATE KEY-----
izfrNTmQLnfsLzi2Wb9xPz2Qj9fQYGgeug3N2MkDuVHwpPcghkHkJgCQuuvT+qZI
MbS2U6wTS24SZk5RunJIUkItRKEwWMS28SLGfkDs1bBY1SPa5smAd3/q10ePi4ae
dU6YgWuDxzBAKEKVSUu6pA2H0dyQ9N4F1dI+F8w9J990zE93EgyNqZFBBa2L70h4
M7DrB0gJBWmDUmoxGnun5glLiCMo2JrHZ9RkMia1lS1sHMhELx2UA1P8I1+0Mav8
iM1HGyUW8EJy0paVf09MPpceEcVwDBEX0+G4UQ10551GTFtOSRjcd8U+GkCzka9W
/SFqrSGe3Gh3SDa0w/4JEMAjWPDLiCglwh0rLI04VwU6AxzTCuCW3d1ZxQsU6VFQ
PqHA8ha0UATZIrP3886PBThVqALBk9p1Nqn51bXLh13Zy9DZIVx4Z5Ioz/EGuzgR
d68VW5wybLjYE2r6Q9nHpitsZ4ZderwjIZRes67HdxYFw8unm4Wo6kuGnb5jSSag
vwBxKzAf30mn+J6IthTJKuDD13rKZGMcRpQQ6VstwhYt1TahQ/qfJUWpJpCU5ML
9LkgVwA8Ndi1wp1/sEpe+U1L16L6v09jUHcueWN7+zSUOE/cDSJyMd9x/ZL8QASA
ETd5dujVIq1INL2vJKr1o4T+i0RsnPfFiqFmBK1Fqww/SKzJeChdyEtpa/dJMrt2
8S86b6zEmkser+SDYgGketS2DZ4hB+vh2ujSxmS8Gkwrn+BfHMzkbti081WbGw0l
eM1tfdFZ6wMTLkxRhBkBK4JiMiUMvpERyPib6a2L6iXTFH+3RUDS6A==
-----END RSA PRIVATE KEY-----
-----BEGIN CERTIFICATE-----
MIICMzCCAzygAwIBAgIJALiPnVsvq8dsMA0GCSqGSIb3DQEBBQUAMFxCzAJBgNV
BAYTA1VTMQwwCgYDVQIEwNmb28xDDAKBgNVBACATA2ZvbzEMMAoGA1UEChMDZm9v
MQwwCgYDVQQLEwNmb28xDDAKBgNVBAMTA2ZvbzAeFw0xMzAzMTkxNTQwMT1aFw0x
ODAzMTgxNTQwMT1aMFxCzAJBgNVBAYTA1VTMQwwCgYDVQIEwNmb28xDDAKBgNV
BACATA2ZvbzEMMAoGA1UEChMDZm9vMQwwCgYDVQQLEwNmb28xDDAKBgNVBAMTA2Zv
bzCBnzANBgkqhkiG9w0BAQEFAAOBjQAwYkCgYEAzdGfXi9CNbMf1UUcvDQh7MYB
OveIHyc0E0KIbhjK5FkCBU4CiZrbfHagaW7ZEcn0tt3Evpb0Mxxc/ZQU2WN/s/WP
xph0pSfsfFsTKM4RhTWD2v4fgk+xZiKd1p0+L4hTtpwnEw0uXRvd0ki6muwV5y/P
+5FHUe1dq+pgTcgzuK8CAwEAAMPMA0wCwYDVRR0PBAQDAGlkMA0GCSqGSIb3DQEB
BQUAA4GBAJiDAATy0mQQeuxWdzLRzXmjvdSuL9GoyT3BF/jSnpxz5/58dba8pWen
v3pj4P3w5Do0so0rzKzY2jEsEit1VM2mLSbQpMM+MUVQCQoiG6W9xucFuxSrWpIS
pAqEAuV4DNoxQKKWmhVv+J0ptMWD25PnpXeq5sXzghfJns1J1QND
-----END CERTIFICATE-----
```

Example of a public.pem File

```
-----BEGIN CERTIFICATE-----
MIICMzCCAzygAwIBAgIJALiPnVsvq8dsMA0GCSqGSIb3DQEEBBQUAMFMxCzAJBgNV
BAYTA1VTMQwwCgYDVQQIEwNmb28xDDAKBgNVBAMTA2ZvbzEMMAoGA1UEChMDZm9v
MQwwCgYDVQQLEwNmb28xDDAKBgNVBAMTA2ZvbzAeFw0xMzAzMTkxNTQwMT1aFw0x
ODAzMTgxNTQwMT1aMFMCzAJBgNVBAYTA1VTMQwwCgYDVQQIEwNmb28xDDAKBgNV
BAMTA2ZvbzEMMAoGA1UEChMDZm9vMQwwCgYDVQQLEwNmb28xDDAKBgNVBAMTA2Zv
bzCBnzANBghkqhkiG9w0BAQEFAA0BjQAwYgYkCgYEAzdGfXi9CNbMf1UUcvDQh7MYB
OveIHyc0E0KIbhjK5FkCBU4CiZrbfHagaW7ZEcN0tt3EvpbOMxxc/ZQU2WN/s/wP
xph0pSfsfTKM4RhTWD2v4fgk+xZiKd1p0+L4hTtpwnEw0uXRvd0ki6muwV5y/P
+5FHUe1dq+pgTcgzuK8CAwEAAMPA0wCwYDVR0PBAQDAGLkMA0GCSqGSIb3DQEB
BQUAA4GBAjiDAAtY0mQQeuxWdzLRzXmjvdSuL9GoyT3BF/jSnpzz5/58dba8pWen
v3pj4P3w5Do0so0rzkZy2jEsEit1VM2mLSbQpMM+MUVQCQoiG6w9xucFuxSrWpIS
pAqEAuV4DN0xQKKWmhVv+J0ptMWD25Pnpxeq5sXzghfJns1J1QND
-----END CERTIFICATE-----
```

Name Limitations

[Table 4](#) contains words that cannot be used as names for any part of a StorNext system.

Table 4 Name Limitations

Active	Enabled	JournalSize	Regular
Affinity	Exclusive	KiloInteger	Rotate
AllocationStrategy	Exit	Letter	Rtios
AttrTokenSize	Extended	Log	RtiosReserve

Appendix A: Operating Guidelines
Name Limitations

B	FileName	Long	Rtmb
Bits	Fl	MaxConnections	RtmbReserve
Brls	ForcePerfectFit	MaxLogs	RtTokenTimeout
BrlTime	ForceStripeAlignment	MaxLogSize	S
BufferCacheSize	Format	MaxMBPerClientReserve	Save
BufferPoolSize	FrBlocks	Mbufs	Sb
BWMFields	FsBlockSize	MbufSize	Sectors
Config	GigaInteger	MediaType	SectorSize
Custom	GlobalSuperUser	MegaInteger	Sg
Cvfsdb	Help	MetaData	Show
DataMigration	HexDigit	MirrorGroup	Static
DataMigrationThreadPoolSize	HexInteger	MirrorReadMethod	StaticNodes
Debug	Icb	MultiPathMethod	Status
Delim	Iel	Name	Sticky
DeviceName	Inode	No	String
Digit	InodeCacheSize	Node	StripeBreadth
Dir	InodeDeleteMax	Ntsd	StripeClusters
DirCacheSize	InodeExpandInc	OnDisk	StripeGroup
DirFDCacheSize	InodeExpandMax	OpHangLimitSecs	TeraInteger
DirWarp	InodeExpandMin	Poke	ThreadPoolSize
Disabled	Integer	Q	Type
Disk	IoHangLimitSecs	Quit	WhiteSpace
DiskType	Journal	Quotas	WindowsSecurity
DosExtraChar	JournalLcBufNum	Raw	Write
Dump	JournalLcBufSize	Read	Yes

Ports Used By StorNext

The following table lists ports that are used by StorNext and its ancillary components.

For a thorough explanation of StorNext's port selection algorithm, consult the **fsports(4)** man page.

Port	Protocol	StorNext Use	Notes
81	TCP	GUI (Java)	User starts at port 81, redirected to 443
443	TCP	GUI (Java)	
1527	TCP	GUI (Java connection to derby db)	
3307	TCP	GUI (Java connection to MySQL)	
1062, 1063	TCP	Blockpool	Both ports if HA primary
14500	TCP	snpolicyd	
5164	TCP	fsmpm	This is the TCP port for the StorNext file system alternate portmapper. See the fsports(4) man page for information on changing the default setting.
5189	TCP	HA Manager	

Port	Protocol	StorNext Use	Notes
Various	TCP	fsm, fsmppm	These ports are used to for metadata exchanges between client hosts and FSM processes on the MDCs and for additional exchanges between StorNext components on hosts within the cluster. See the fsports(4) man page for information on setting up a range of ports that will be used.
Various	UDP	fsmppm	These ports are used to exchange heartbeat messages between the client hosts and the coordinator hosts. See the fsports(4) man page for information on setting up a range of ports that will be used.
20566	TCP	MySQL	Only used internally on an MDC.

Port	Protocol	StorNext Use	Notes
60001, 60002 ...	TCP	ACSLs Tape Libraries	Not used by StorNext, but related
61776	TCP	SNAPI	The SNAPI port number can be changed within the <code>snapi.cfg</code> file.

Log Rolling and Disk Space Health Check

Log Rolling

StorNext automatically rolls logs via a scheduled cron job to better manage the file system space used by the rolled log files. An additional log rolling cron job will now run frequently to check for and roll “runaway” log files. Additional functionality includes:

- Compression of rolled log files
- Configuration parameters specifically for space management
- Efficient log rolling performance

The configuration parameters include:

- **CRITICAL_FILL_LEVEL**: If the file system exceeds this fill level (default: 98%), StorNext will remove rolled files to recover space.
- **CLEANUP_MIN_SIZE**: StorNext will remove rolled files during file system space recovery only if they meet this minimum size (default: 10 MB).

All log rolling configuration parameters are contained in the `sn_log_update.cfg` configuration file located in the following directory:

```
/usr/adic/util/sn_log_update.cfg
```

Note: StorNext log rolling will not guarantee that the file system will not fill up, given that other non-StorNext files using the same file system may accumulate and fill up the file system. Additionally, the probability exists that a StorNext log may grow at an extraordinarily rapid rate that exceeds the ability for the automatic log rolling to keep up.

StorNext uses a timestamp for the filename extension, instead of a sequential numeric count extension. The rolled files are compressed; a ".gz" extension is appended to the filename. The complete filename format for rolled log files is illustrated in the following example:

```
tac_00.08:08:2012:13:00:01.gz
```

You may also configure the log rolling cron job to back up rolled log files to a managed file system. A new storage policy class will be needed, which will require additional media or sdisk space. The size of the data stored will depend on system activity levels.

The instructions below will back up all rolled logs. Additionally, the MSM and TSM tac log backups will be compressed to approximately 1/20th of their original size before being stored.

Identify the managed file system containing the .ADIC_INTERNAL_BACKUP directory (typically /stornext/snfs1). See the BACKUPFS environment variable in /usr/adic/TSM/config/fs_sysparm to determine the file system name. In the instructions below, change "/stornext/snfs1" as necessary to the name of the managed file system containing the backups.

- 8 If this is a High Availability (HA) configuration, execute the following command on both MDCs.

```
# mkdir -p /stornext/snfs1/.SNSM_LOG_BACKUP/`/bin/  
hostname`
```

- 9 Execute the following two commands on the active MDC only.

```
# fsaddclass -d 1 -f i -m 5 _snsm_log_backup  
# fsaddrelation /stornext/snfs1/.SNSM_LOG_BACKUP -c  
_snsm_log_backup
```

- 10 Save the existing tdlm crontab so it may be restored if an error occurs while updating the crontab.

```
# /usr/bin/crontab -l -u tdlm > /tmp/  
crontab.tdlm.save
```

- 11 Execute the following command to edit the tdlm crontab on the active MDC.

```
# /usr/bin/crontab -e -u tdlm
```

- 12 Append the following text to the end of the existing `sn_log_update` entry. It is all one continuous line. Note that it begins with a space, and there is a space preceding every hyphen, every occurrence of `/usr/adic`, and every occurrence of `/stornext/snfs1`.

```
-s /stornext/snfs1/.SNSM_LOG_BACKUP/`/bin/  
hostname`; /usr/adic/gui/bin/cmdwrap -  
NO_END_OF_FILE /bin/gzip /stornext/snfs1/  
.SNSM_LOG_BACKUP/`/bin/hostname`/?SM/logs/tac/  
tac_00.???:??:20??:??:??:??
```

A complete **crontab** command is illustrated in the following example:

```
0 1,7,13,19 * * * /usr/adic/gui/bin/cmdwrap -  
NO_END_OF_FILE /usr/adic/util/sn_log_update /usr/  
adic -s /stornext/snfs1/.SNSM_LOG_BACKUP/`/bin/  
hostname`; /usr/adic/gui/bin/cmdwrap -  
NO_END_OF_FILE /bin/gzip /stornext/snfs1/  
.SNSM_LOG_BACKUP/`/bin/hostname`/?SM/logs/tac/  
tac_00.???:??:20??:??:??:??
```

Disk Space Health Check

As part of its suite of health checks, StorNext Storage Manager periodically runs the "Disk Space" health check to find all file systems in use by StorNext that are running out of space. This check includes all file systems that are accessible from `/usr/adic`, including those reached by symbolic link. This check does not include any SNFS file systems configured by the user.

The Disk Space health check and the **CRITICAL_FILL_LEVEL** parameter for `/usr/adic/util/sn_log_update.cfg` can be used together to determine desired notification behavior. The system will send an admin alert warning if the Disk Space health check exceeds the health check's warning threshold. When the `sn_log_update` **CRITICAL_FILL_LEVEL** triggers, it will cause Storage Manager logs to start being automatically removed from the system until the percentage of the file system in use falls below the **CRITICAL_FILL_LEVEL**. The Disk Space health check generates a RAS event if the health check's fail threshold is exceeded. If users continually run with low disk space conditions, they may want to consider dropping the **CRITICAL_FILL_LEVEL** to match the Disk Space warning threshold.

By default, the Disk Space health check will report warnings if a file system is above 95% disk usage, and it reports failures if a file system is above 99% disk usage. If desired, the user may override these defaults by editing the `/usr/adic/TSM/config/filelist` configuration file. Find the `health_check` entry with the `checkDiskSpaceTsm` command. Append the `-w` option to the command to override the percentage threshold for warning messages. Append the `-f` option to the command to override the percentage threshold for failure messages. For example:

```
health_check : 0 : Disk Space : checkDiskSpaceTsm -w  
96 -f 98 : 0
```

In the example above, a user has configured the Disk Space health check to issue file system space usage warning messages at 96% usage and failure messages at 98% usage.

General Operating Guidelines and Limitations

[Table 5](#) lists updated information and guidelines for running StorNext, as well as known limitations.

Table 5 Operating Guidelines and Limitations

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. This issue also affects Samba servers running on Linux. 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>
Linux; RedHat Enterprise Linux 4, 5, and 6; SUSE Linux Enterprise Server 10, and 11	StorNext File System	<p>Many versions of Linux run a cron script nightly to build a database used by the <code>slocate</code> command. If StorNext file systems are mounted, they are traversed by this cron job which can have a dramatic impact on the performance of other applications currently using these file systems. Perform the following steps (based on Linux version) to prevent the cron script from traversing StorNext file systems.</p> <p>RedHat Enterprise Linux 4, 5, and 6</p> <p>Add “<code>cvfs</code>” to the list of file system types to be skipped. This is usually done by modifying the “<code>PRUNEFS</code>” line in the <code>/etc/updatedb.conf</code> file to read:</p> <pre>PRUNEFS="cvfs sysfs selinuxfs usbdevfs devpts NFS nfs nfs4 afs sfs proc smbfs cifs autofs auto iso9660 udf"</pre> <p>SUSE Linux Enterprise Server 10, and 11</p> <p>The optional “<code>findutils-locate</code>” package is used to build the <code>slocate</code> database. The default behavior is to disable building the database. If enabled, to prevent <code>cvfs</code> file systems from being scanned, add “<code>cvfs</code>” to the list of file system types to be skipped. This is usually done by modifying the “<code>UPDATEDB_PRUNEFS</code>” line in the <code>/etc/sysconfig/locate</code> file to read:</p> <pre>UPDATEDB_PRUNEFS="cvfs"</pre>

Operating System	Feature or Category	Description
Linux	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
	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, Red Hat 6, SuSe 10 and SuSe 11.</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>
	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 File System Tuning Guide</i>.</p>

Operating System	Feature or Category	Description
Linux	HA	<p>Changing the haFsType 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 HaShared or HaManaged 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</p> <pre data-bbox="661 579 1305 708">/usr/adic/install/.snsm_ha_configured touch file (\$SNSM_HA_CONFIGURED environment variable). No conversion is done or necessary for SNFS only (HaUnmanaged) configurations.</pre> <p>If the conversion is not successfully completed, the HaManaged FSMs will not start, and the HaShared FSM will cause an HA Reset when it is stopped.</p> <p>To remedy this situation, edit every FSM configuration file to set its haFsType parameter to HaUnmonitored, then run the following commands to avoid the HA Reset in this special case only:</p> <pre data-bbox="661 961 1243 1031">touch /usr/cvfs/install/.vip_down_hint service cvfs stop</pre>
	Subtree Check option	<p>Subtree Check Option in NFS No Longer Supported</p> <p>Although supported in previous StorNext releases, the <code>subtree_check</code> 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>
	FQDN	<p>SuSe Linux distributions automatically associate the FQDN of the local machine with the address 127.0.0.2 in the <code>/etc/hosts</code> 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 <code>/etc/hosts</code> file.</p>

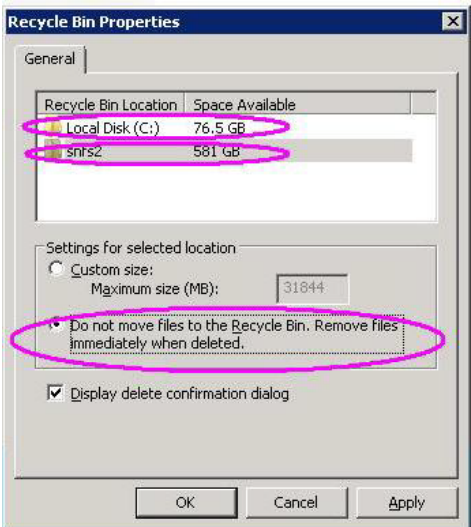
Operating System	Feature or Category	Description
Linux	Software Firewalls	<p>Software firewalls such as “iptables” on Linux and Windows Firewall can interfere with the proper functioning of StorNext and result in unexpected errors unless specifically configured for use with StorNext.</p> <p>Quantum strongly 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 File System Tuning Guide</i>.</p>
	cpuspeed Service	<p>cpuspeed, an external Linux service on recent Intel processors, is not correctly tuned to allow StorNext to take advantage of processor speed. Suse systems may also be impacted, as may AMD processors with similar capabilities.</p> <p>On processors with a variable clockspeed (turbo boost), the cpuspeed service on Redhat controls the actual running speed of the processors based on system load.</p> <p>A workload such as a heavily used FSM and probably Storage Manager does not register as something which needs a faster cpu. Turning off the cpuspeed service has been shown to double metadata performance on affected hardware.</p> <p>Looking at the reported CPU clock speed by doing <code>cat /proc/ cpuinfo</code> while the system is under load shows if a system is impacted by this issue.</p>

Operating System	Feature or Category	Description
Linux	Power-on Diagnostics	<p>During testing a Quantum PX502 library running Red Hat 6.1 did not finish power-on diagnostics. When the same test was run on a PX502 library running either Red Hat 5.X or SuSE 10 / 11, power-on diagnostics completed and the system initialized without any issues.</p> <p>The workaround for this issue is to disconnect the SAN from the library running Red Hat 6.1. If the library powers on while the SAN is disconnected from the library controller, the library finishes its power-on diagnostics and performs an audit of the library. Subsequently reconnecting the Red Hat 6.1 server to the SAN (library ready) causes the library to perform a new physical audit of the library.</p> <p>Note: Testing was performed on the Red Hat 6.1 system which did not have StorNext loaded or running.</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>
	StorNext File System	<p>Virtual Hard Disk (VHD) files are not supported on a StorNext file system.</p>

Operating System	Feature or Category	Description
Windows	StorNext Security	<p>In StorNext releases prior to 3.5, the StorNext Windows client attempted to keep the UNIX uid, gid and mode bits synchronized with similar fields in the Windows security descriptor. However, these Windows and UNIX fields were often not synchronized correctly due to mapping and other problems. One consequence of this problem was that changing the owner in Windows incorrectly changed the UNIX uid and file permissions and propagated these errors into sub-directories.</p> <p>In release 3.5, the StorNext Windows client sets the UNIX uid, gid and mode bits only when Windows creates a file. The StorNext Windows client will no longer change the Unix uid, gid or mode bits when a Windows user changes the Windows security descriptor or Read-Only file attribute.</p> <p>If you change the UNIX mode bits and the file is accessible from Windows, you must change the Windows security descriptor (if Windows Security is configured On) or Read-Only file attribute to ensure the change is reflected on both Windows and UNIX.</p>

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 (see http://support.microsoft.com/kb/263253).</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.

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	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.

Operating System	Feature or Category	Description
All	Tape drives	<p>StorNext's configuration of the tape drives within an ACSLS library can get out of sync with the ACSLS server's configuration of the tape drives. If this occurs, then when StorNext sends a request(s) to the ACSLS server for a specific drive it could be sending the wrong location to the drive.</p> <p>The problem occurs if location modification occurs with the tape drives in the ACSLS library such that what StorNext thinks are the correct ACSLS locations no longer match what the ACSLS server knows about. This can happen by one of the following library maintenance type activities:</p> <ol style="list-style-type: none"> 1 If a tape drive is replaced with a new tape drive. 2 If two tape drives are swapped within the library. 3 If new tape drive(s) are added into location(s) that causes ACSLS to assign different locations to the previously existing drives. 4 If a tape drive(s) are removed from location(s) that causes ACSLS to assign different locations to the remaining drives. <p>Workaround:</p> <p>Contact Quantum Technical Support.</p> <p>There is currently no automated way to update StorNext to re-sync itself with the tape drive changes done within the library. It requires knowledge of what library scenarios were done in order to make the correct changes within StorNext. The necessary changes can be done to update StorNext's library/tape drive configuration to match what ACSLS knows of, but this is a manual process, and requires detailed knowledge of what needs to be updated in order to accomplish the procedure.</p>
	Tape drives	<p>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.</p>

Operating System	Feature or Category	Description
All	Tape drives	<p>Tools outside of StorNext that issue a st command to rewind tapes may result in data loss.</p> <p>Workaround(s):</p> <ol style="list-style-type: none"> 1 Do not run any tools that could possibly issue a rewind. <ul style="list-style-type: none"> • It may not be possible to determine this ahead of time. • Gathering hardware info is imperative. 2 Rename the /dev/st* devices per "Product Alert 16". Product Alerts are available on the Quantum Service and Support website at www.quantum.com/ServiceandSupport. <ul style="list-style-type: none"> • They get created on the next reboot. 3 Stop SNSM before running any tools. <ul style="list-style-type: none"> • Stopping TSM could take 10-15 minutes if very busy. • Not optimal to stop/start TSM all the time.
	Cluster-Wide Central Control	<p>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 resides under /usr/cvfs/config on an nss coordinator server.</p> <p>Currently the nss coordinator server capable of parsing this xml file must be on the Linux platform.</p>

Operating System	Feature or Category	Description
All	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.
	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>
All	HA	<p>On HA systems only:</p> <p>The <code>/usr/cvfs/config/ha_peer</code> 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 <code>ha_peer</code> 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 <code>ha_peer</code> file is removed for any length of time while StorNext is running, the <code>snhamgr(1)</code> 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 <code>snhamgr</code> itself. If this occurs, restore the <code>ha_peer</code> file to the IP address of the peer MDC, and then restart the HA Manager service by running the following command: <pre>service snhamgr restart</pre> </p>

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 <code>255.255.255.255</code> 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 (<code>255.255.255.255 == 255.255.255.255</code>).</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>.


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 "cvadmin -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. 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	Quotas	Quotas can only be enabled or disabled by modifying the Quotas parameter of the file system config file. The CLI <code>cvadmin "quotas"</code> command will still return the quota state, but it cannot change it.
	fsretrieve	If you run multiple fsretrieve 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. 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 <code>% fsretrieve -R .</code>
All	Stripe group expansion	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.
	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. For additional information, see StorNext LAN Clients in HA Environments.

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	<p>Note the following recommendations and limitations for using DXi as a virtual tape library for StorNext:</p> <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	<p>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.</p> <p>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>.</p>

Operating System	Feature or Category	Description
All	Converting file systems	StorNext does not currently support converting from a managed file system to an unmanaged file system.
	JournalSize Setting	<p>The optimal settings for <code>JournalSize</code> are in the range between 16M and 64M, depending on the <code>FsBlockSize</code>. 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.</p> <p>For more information about <code>JournalSize</code>, refer to “<code>FsBlockSize</code>, Metadata Disk Size, and <code>JournalSize</code> Settings” in Appendix A of the <i>StorNext User’s Guide</i>.</p>
	Truncation	<p>StorNext does not currently support running two truncation policies for different policy classes at the same time.</p> <p>Workaround:</p> <p>To avoid errors, do not run two truncation policies for different classes at the same time.</p>

Appendix A: Operating Guidelines
General Operating Guidelines and Limitations



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 in Replication Versus Copies and Versions in Storage Manager

Snpolicy supports multiple copies (instances) of files and directories on the target, while TSM supports both multiple copies and multiple versions of files and directories. 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 (`rep_copies`) 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

There is a special case where policy parameter `rep_copies` is set to 0. In this configuration, the target creates a target directory for the first namespace replication. It then operates on the same target directory for all subsequent replications on the same source directory.

Note: Currently there are some limitations to this setting:

1. If a file is removed on the source, the target will not remove the file accordingly at next namespace replication time.
 2. If a file is renamed on the source, at the next namespace replication, though the new name will appear in the target directory, but the old file name still stays.
-

Context 3

`snpolicyd` does not support multiple versions. Though we support multiple copies, so a file can appear in directories of multiple copies, however, they are all linked to the same inode. This means whenever a change is made on the source and is populated to target by replication, all copies of the file will be impacted on the target. For instance, if a source file has its file owner changed, and the change is populated to the target through replication, all previous copies of this file will have the new owner instead of the prior owner.

Context 4

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 `host1`, 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 5

Storage Manager also supports multiple 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 policy. (Actually, 4 is the default maximum number of copies. The number can be greater than 4.) Each of these copies is a copy of the same file contents. Different from snpolicy copies, the set of files and directories are the same in different copies.

Context 6

Unlike snpolicyd, Storage Manager supports multiple versions. 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 7

File Recovery. When a file is removed by accident from a snpolicy source directory, before a new namespace replication occurs, or if the policy is configured with multiple rep copies, the file still exists in rep copies on target (if new namespace replication does not occur, the file exists in each rep copy, otherwise, the file exists in all rep copies except rep copy 0) and can be copied back to restore. On the other hand, 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 snpolicy 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 `exporttrepcopy` 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 <code>fsaddclass</code> and <code>fsmodclass</code>	Yes. Use the <code>snpolicy</code> command.

Storage Policy Option	Storage Manager Policy	snpolicyd Policy
<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.
<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 -dumpopol/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:

```
total 4160
```

```

lrwxrwxrwx 1 root      root      72 Jan 26 10:13
0x1a4062.0 -> /snfs/sn2/
__CVFS_Handle.00047DA10CB84CF8000E0000000000000000000000000000000
1A412B
-rwxr-xr-x 2 testuser root 1388936 Jan 26 10:11
0x1a4065.0
-rw-r--r-- 2 testuser root 1430896 Jan 26 10:11
0x1a4066.0
-rw-r--r-- 2 testuser root 1397888 Jan 26 10:12
0x1a4067.0
drwx----- 2 root      root      2051 Jan 26 10:13 copies

```

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
```

Appendix B: Additional Replication and Deduplication Information
Deduplication Overview

```
        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 dumpool` option, you will see `dedup=on` in the output when the policy has deduplication enabled.

Deduplication Modification Time

Note that in the "`snpolicy -dumpool`" 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 -dumppl" 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 "fFlags:" 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.

Also note that when Storage Manager processes each store candidate, it needs to obtain its parent directory information. If the parent directory is removed, the candidate is discarded. A store candidate event is generated in the first namespace replication after the file has been changed on the source side. As a result, if the store candidate is discarded during store processing, unless the file on the source is changed again, all subsequent namespace replications will not generate

a store candidate event for this file. This could lead to file candidates not being stored promptly. The file needs to become a store candidate again by rebuilding store candidate list before it can be stored to media by Storage Manager.

A store candidate parent directory could be removed on the target if the interval between two consecutive namespace replications is too short. For each namespace replication, the target retains a maximum number of rep copies designated by replication parameter **rep_copies** or **Copies to Keep on Target** labeled by the GUI. By default, it is **1**. When the number of rep copies exceeds the number designated by **rep_copies**, the oldest namespace is removed. So if **rep_copies** is **1**, then only 1 namespace is maintained, the next namespace replication will cause the prior namespace to be removed. Thus, if the store candidates are generated from the first namespace replication, when **rep_copies** is **1** and the store operation has not been finished yet, its parent directory can be destroyed by the next namespace replication.

As a result, to ensure that legitimate store candidates are not discarded, configure the interval of scheduled namespace replication to be long enough or configure enough replication copies to allow the newly generated store candidates to finish storing before its namespace is destroyed by subsequent replication.

If, for example, a store candidate has been discarded due to the parent directory missing, then you can detect the files that were discarded and add them back to the store candidates list by performing the following:

- 1 Run `snetadump -a FsName` to apply restore journal to metadump
- 2 Run `fspolicy -b -y mnt_pnt` to add the discarded files back as store candidates

To detect whether store candidates were discarded, check the TSM tac log file for log messages, such as:

```
Mar 11 20:54:47 sjoshi-rh62-2 sntsm fspolicy[13776]:  
E1201(8)<1107031798>:msa2dmi1222: dm_get_fileattr failed, ino:  
3429418 gen: 0  
errno: (2) No such file or directory  
Mar 11 20:54:47 sjoshi-rh62-2 sntsm fspolicy[13776]:  
E1200(7)<1107031798>:mda4str1954: parent stat failed ino:  
3429418 gen: 0  
errno: 2
```


Alternatively, check whether the file was stored long after the namespace replication was finished by running `/usr/adic/TSM/bin/fsfileinfo -o file`. Running the command will display whether the file was stored to tape or object IDs, if it was stored to WASTorage.

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 `snpolicyd.out.1`, `snpolicyd.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 a HA System](#)
- [Moving a HA Shared File System to a New Raid](#)

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. Un-managed 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.

Caution: Operating two (or more) MDCs configured to serve any StorNext file system is only supported when HA protection is configured for the file system. Otherwise, it is possible for metadata to become corrupt if there are simultaneous failures, delays or excessive loads in the LAN and SAN networks.

HA Internals: HAmon 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 HAmon timers puts an upper limit on the range of timing for ARB brand updates. Brand updates and HAmon 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 HAmom 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.

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 HAmon 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 HAmon 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 HAmon 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 FSMPPM 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 FSMPPM, 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 FSMPPM 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 HAmon 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 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-settable 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.

Caution: If a secondary MDC in **locked** mode is rebooted or powered down while the primary MDC is in **config** or **single** mode, **snhamgr** may detect that the HA cluster is in an invalid state. If it does, it will attempt to safeguard the HA cluster by stopping StorNext on the primary MDC and putting it into **default** mode. To return the HA cluster back to the **config** state, check to ensure that the secondary MDC is either powered off and **peerdown** is set on the primary MDC or that the secondary MDC is running and its **snhamgr** mode is set to **locked**. Once this is verified, restart StorNext on the primary MDC and then set its **snhamgr** mode back to **config**.

- 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 FSM 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. The commands are listed in the table below and the command syntax is as follows, where *cmd* is the command in the table:

```
snhamgr cmd
```

Command	Description
status	Return cluster modes and operating statuses. All commands return status; this one does nothing else.
stop	Safely stop both servers in the cluster without incurring a HA reset. The secondary server is placed in locked mode, which stops StorNext on that server, then the primary server is placed in config mode and stopped, and then both servers are put in default mode with StorNext stopped.
start	Stop each server when there is a need and transition both servers to default mode, then bring up the local server first followed by the peer server so that the local server becomes primary and the peer server becomes secondary. Note: Running <code>service cvfs start</code> is the preferred method of starting, so Quantum recommends using this command rather than using other methods. Likewise, use <code>service cvfs stop</code> to stop StorNext.
config	First, check that the peer server is in locked or peerdown mode. Then, place the local server in config mode. The command must be run on the primary server, or either server when CVFS is stopped on both.

Command	Description
clear	Remove the file referenced by the HA_IDLE_FAILED_STARTUP environment variable. Run this after correcting any conditions that caused the previous failure of StorNext startup scripts.
force smith	Trigger an immediate HA reset if the local server is in default mode. This command is meant for use in health-monitoring scripts. The command is two words to make accidental firing less likely.
peerdown	Certify that the peer server is powered off. This mode is used when the peer server is powered down. In the event that the peer returns to service and begins to communicate, the assertion that the peer is down becomes false. Immediate action may be taken by the local server to transition itself to a safe operating mode, which could trigger an HA reset. The best practice is to power off the server or uninstall StorNext before setting peerdown mode, and to unset the mode before powering on the server.
peerup	Undo the peerdown mode. The command will fail if the local mode is config or single . Run this command before powering on the peer server. The local server will assume the peer is in default mode until the peer starts snhamgr_daemon communications.

Command	Description
mode= <modeval>	Set the mode of the local server to modeval . The mode is stored on the local server so that it persists across reboots. config Set the mode to config . The peer server must be in locked or peerdown mode. default Set the mode to default . The peer server can be in any mode. locked Set the mode to locked . The peer server can be in any mode. single Set the mode to single . The peer server must be in locked (recommended) or peerdown mode.

See the *Man Pages Reference Guide* for additional details on the commands.

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/.snm_ha_configured
```

The existence of that file enables the running of the snhamgr 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 .snm_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.

Caution: Operating two (or more) MDCs configured to serve any StorNext file system is only supported when HA protection is configured for the file system. Otherwise, it is possible for metadata to become corrupt if there are simultaneous failures, delays or excessive loads in the LAN and SAN networks.

Figure 172 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 540.

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 HAmom 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.

Configuration Changes

When making configuration changes that require an FSM to be stopped, the primary HA server must be placed in **config** mode. When configuration or software changes are made, a StorNext HA cluster may need to be placed in **config** mode with only one server running. 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.

Many file system configuration changes require an active FSM (file system manager process for each file system) to be stopped and restarted on the primary HA node. This can trigger a standby FSM on the secondary HA node to take over control of the file system, likely using the old/original configuration.

Examples of changes that require an HA downgrade:

- Removing file systems (and adding new file systems).
- Changing name servers (fsnameservers file).
- File system configuration file changes 'config/*.cfg' and 'config/*.cfgx'.
 - Stripe group changes.
 - Adding stripe groups.
 - Marking stripe groups **OFF**.
 - Adding LUNs to stripe groups (bandwidth expansion).
- Changes in the GUI panels under **Configuration > File Systems > Edit > file-system-name** that match a parameter found in *.cfgx cause an FSM restart.
- Associating/disassociating affinities with a stripe group requires a *.cfgx change. Using **cvaffinity** to connect an affinity with a file or directory does not change *.cfgx.
- Additional guidance can be found within snfs_config(5) and snfs.cfgx(5) in the *MAN Pages Reference Guide*.

Note: License changes do not cause an FSM restart, nor do Policy class changes.

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 `fsmppm` process, so there would not be an entry in the file when an `fsmppm` 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 peerdownt with the 'snhamgr peerdownt' 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.

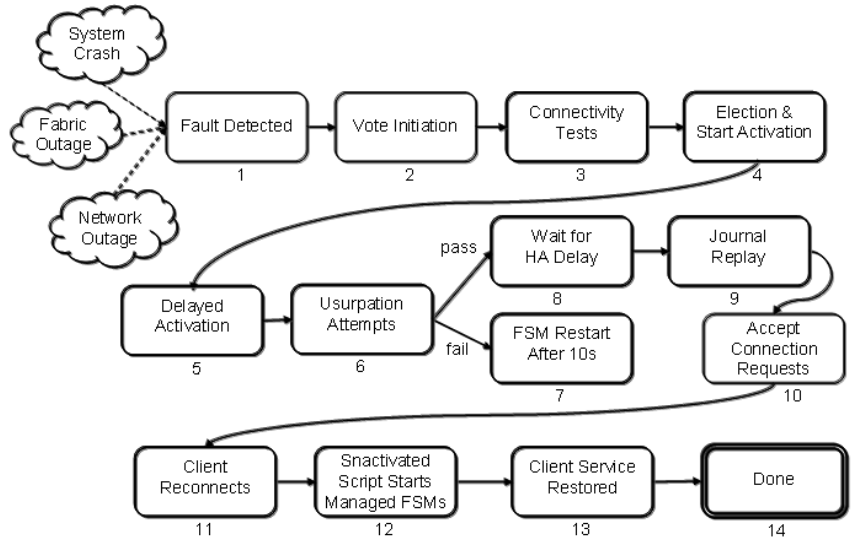
The StorNext system exclusively opens files on the HaShared file system, but assumes that only the Primary MDC does this and waives the delay for that one file system. Quantum advises against running user processes other than StorNext processes on HA MDCs for performance, reliability and availability 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 173 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 T_{3.5} 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 HAmon 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 FSMPM 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 a HA System

To replace a HA server, refer to [Replacing an MDC in an HA environment](#) on page 684 in [Appendix J, Repairing and Replacing StorNext Metadata Servers](#).

Moving a HA Shared File System to a New Raid

Caution: This procedure was created for **Quantum Professional Services**. It is highly recommended that **only** Quantum Professional Services or a certified Quantum partner perform this task. Failure to correctly implement this procedure could lead to configuration, file system and/or database problems, which may require a Professional Services engagement to resolve.

This procedure requires downtime because all StorNext file systems will be inaccessible. The pathnames of Linux commands used below are those from RedHat Enterprise Linux 5.

The checks for filesystems using **cvadmin** assumes that the HA MDCs are the only users of the fsnameservers. If the fsnameserver has been offloaded to another system, it may show filesystems that are not relevant to the HA MDCs.

- 1 Create a new unmanaged FS to be used for the new HaShared FS and mount it.

Note: While its called the HaShared FS, it contains much more than HA information. This has the MySQL database, metadumps, storage manager configuration and more (if configured).

- 2 Run **cvgather** to save original config, just in case. Save it in a non-StorNext location.
- 3 Uninstall StorNext on the secondary system.

Note: This will automatically stop StorNext on the secondary system. HA conversion will be re-done after the primary system has been modified.

Use the StorNext 4.7.0 (or later) installer from the StorNext distribution (DVD or download) to uninstall StorNext. Go to the distribution directory and execute the following command:

```
./install.stornext -remove
```

- 4 After this step, the primary will have **filesystem** only running.

On Primary server:

a `service stornext_web stop` or `/etc/init.d/stornext_web stop`

b `service cvfs stop` or `/etc/init.d/cvfs stop`

c `/usr/adic/DSM/bin/snhamgr status`

Expect to see this:

LocalMode=default

LocalStatus=stopped

RemoteMode=default

RemoteStatus=unknown

d `/usr/adic/DSM/bin/snhamgr peerdwn`

The secondary should have been uninstalled. If that was not done properly, this may fail.

Expect to see this:

LocalMode=default

LocalStatus=stopped

RemoteMode=peerdwn

RemoteStatus=unknown

e `/usr/adic/DSM/bin/snhamgr mode=config`

Expect to see this:

LocalMode=config

LocalStatus=stopped

RemoteMode=peerdwn

RemoteStatus=unknown

f `/usr/adic/DSM/bin/DSM_control startonly`

Starts only file system components.

- 5 After this step, only the old and new HaShared filesystems should be running.
 - a On the primary MDC unmount all managed file systems, if any.
 - b `/bin/mount`
Ensure no managed file systems are mounted.
 - c Stop all managed file systems on the primary MDC using `/usr/adic/DSM/bin/cvadmin`
Unmounting and stopping the managed filesystems is important because active filesystems will be writing to restore journals, which reside on the HaShared filesystem. It is not required that the unmanaged filesystems be unmounted and stopped, but recommended.
 - d Unmount all remaining StorNext filesystems except old and new HA filesystems.
 - e `/bin/mount`
Ensure file systems are unmounted.
 - f Stop all remaining StorNext filesystems except old and new HA filesystems using `/usr/adic/DSM/bin/cvadmin`.
 - g `/usr/adic/DSM/bin/cvadmin -e select`
Only the old and new HA file systems should be listed.
- 6 After this step, the data is copied from the old to the new HaShared filesystem.
 - a `cd /usr/adic/HAM/shared`
 - b `/usr/bin/find . | /bin/cpio -pdum /<new HA shared mount point>`
- 7 Stop all StorNext operations.
 - a `cd /`
Somewhere out of the StorNext directory or your login will be fuser'ed and you will be logged out.
 - b Unmount the old and new HaShared filesystems.
 - c `service cvfs stop` or `/etc/init.d/cvfs stop`
- 8 After this step, the new HaShared filesystem will be used and StorNext will be running.

a `/usr/adic/DSM/bin/sncfgedit -n <new ha shared fsname>`

Change HaUnmanaged to HaShared

Note: This is case sensitive.

b `/usr/adic/DSM/bin/sncfgedit -n <old ha shared fsname>`

Change HaShared to HaUnmanaged

Note: This is case sensitive.

c `mv /usr/cvfs/config/<old HA FS>.cfgx /usr/cvfs/config/<old HA FS>.cfgx.save`

d Remove the old HA file system name from `fsmlist`. Add the new file system name, if it is not already there.

e `vi /etc/fstab`

Change file system name for `/usr/adic/HAM/shared` to the new file system name.

f Also in `/etc/fstab`

Remove pre-existing mount for the new HaShared file system.

g `vi $SNSM_HA_CONFIGURED` file (`/usr/adic/install/.snsn_ha_configured`) and change it to the new file system name.

h `vi /usr/adic/HAM/.SHARED_FSNAME` and change it to the new file system name.

i `service cvfs start` or `/etc/init.d/cvfs start`

Starting StorNext is important to do before changing back to default mode. The system is currently in config mode. Changing to default mode before starting may cause updates to be overwritten by HA.

9 Update HA mode on the primary.

a `/usr/adic/DSM/bin/snhamgr mode=default`

Setting mode to default is an implicit stop of services, so it will need to be started again.

Expect to see this:

LocalMode=default

LocalStatus=stopped

RemoteMode=peerdown

RemoteStatus=unknown

b snhamgr peerup

Expect to see this:

LocalMode=default

LocalStatus=stopped

RemoteMode=default

RemoteStatus=unknown

c service stornext_web start

d service cvfs start

e /usr/adic/util/syncha.sh -primary

Should have been done automatically, but repeating it is safer.

f /bin/mount

Ensure all StorNext file systems are mounted and accessible.

10 To bring the Secondary online, perform the following:

a If needed, power up secondary.

b Reinstall StorNext 4.7.x on the secondary by using the installer from the StorNext distribution (DVD or download). Go to the distribution directory and execute the following command:

```
./install.stornext
```

c Replace the existing /usr/cvfs/config/license.dat file on the secondary with the license.dat file from the primary.

d Using the StorNext GUI on the primary, convert to HA. Make certain the GUI displays the new HaShared file system name during this step.



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 the 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)
- `fsfiletapeloc` (getFileTapeLocation)
- `fsmedinfo` (getMediaReport)
- `fsmedlist` (getMediaInfo)
- `fsqueue` (getSMQueue)
- `fsrelocate` (relocateFiles)

- fsretrieve (doRetrieve)
- fsrmdiskcopy (doTruncate)
- fsstate (getDriveReport)
- fsstore (doStore)
- fswascfg (getWasConfiguration)
- vsmove (doMediaMove)

The new fsxsd and vsxsd commands provide 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 MySQL 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, MySQL 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.

Parameters

Parameter	Req / Opt	Num	Description
password	Required	1	This is the password stored in the /usr/adic/.snapipassword file on the server. Example: password=stornext
requestID	Required	1	The request identifier of the request to be cancelled. Operations in progress can be retrieved by running the getSMQueue API. Example: requestID=12345
format	Optional	1	Output format: json, text, or xml. The default output format is text. Example: format=xml

Num is the number of arguments: 1 = only one allowed, N = multiple allowed, the parameter name is repeated for multiple entries, for example, x=1&x=2&x=3.

Example

```
http://host:81/axis2/services/stornext/doCancel
?password=stornext
&requestID=12345
&format=json
```

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.

Parameters

Parameter	Req / Opt	Num	Description
password	Required	1	This is the password stored in the /usr/adic/.snapipassword file on the server. Example: password=stornext
mediaIDs	Required	N	Specifies one or more media to be moved. A valid medium identifier may contain up to 16 alphanumeric characters, including spaces. Leading and trailing spaces are not permitted. The number of media that can be specified is restricted by the CLI software. Currently, the maximum allowed number is 64. Example: mediaIDs=1234
archiveName	Required	1	Specifies the name of the archive to which the specified media are to be moved. Valid archive names may contain up to 16 alphanumeric characters, including spaces. Leading and trailing spaces are not permitted. Example: archiveName=lto
interactive	Optional	1	
remoteHost	Optional	1	The host name of the Media Manager server. The default host name is the host name of the computer where this web service command is issued. Example: remoteHost=hostname

Parameter	Req / Opt	Num	Description
format	Optional	1	Output format: json, text, or xml. The default output format is text. Example: format=xml

Num is the number of arguments: 1 = only one allowed, N = multiple allowed, the parameter name is repeated for multiple entries, for example, x=1&x=2&x=3.

Example

```
http://host:81/axis2/services/stornext/doMediaMove
?password=stornext
&mediaIDs=1234
&mediaIDs=5678
&archiveName=lto
&format=json
```

The doRetrieve API

Use the doRetrieve API to retrieve or recover a file from media and place it on disk. By default, the primary copy of a file is retrieved. Running this API is equivalent to running the fsretrieve command.

Parameters

Parameter	Req / Opt	Num	Description
password	Required	1	This is the password stored in the /usr/adic/.snapipassword file on the server. Example: password=stornext

Parameter	Req / Opt	Num	Description
files	Optional	N	The path and file name of the file(s) to retrieve. The full path name starting from the root directory is required as input to the command. Each file path must specify a file in a migration directory. Example: files=/stornext/dir/a
updateATime	Optional	1	Updates the access time of the requested files. Example: updateATime=1
copy	Optional	1	Used to retrieve a specific copy of filename if one exists. A number from 1 to N. Example: copy=2
newFileName	Optional	1	The new path and file name into which to retrieve the file. The location specified for the new file must be a local file system. Retrieval to an NFS-mounted file system is not permitted. Example: newFileName=/dir/new/filename

Parameter	Req / Opt	Num	Description
starByte	Optional	1	The startByte must be less than endByte, and both must be within the byte range of the file. The byte range is inclusive. To retrieve a single byte, the startByte is equal to the endByte. If the startByte and endByte are specified, the newFileName must be specified. Otherwise, the command is rejected. The byte range is zero relative; therefore a specified byte range must be zero to the end byte minus 1. Example: startByte=123
endByte	Optional	1	See the description for startByte Example: endByte=456
directory	Optional	1	The directory from which to start the recursive retrieve. All files from the specified directory and any subdirectories will be retrieved. Depending upon the number of files in the directory and subdirectories, running this option may use extensive Tertiary Manager resources. Example: directory=/stornext/managed/pooldir/mp3

Parameter	Req / Opt	Num	Description
topPriority	Optional	1	Specifies top priority and will cause all files for the retrieve request to be placed at the top of the retrieve queue. Example: topPriority=1
format	Optional	1	Output format: json, text, or xml. The default output format is text. Example: format=xml

Num is the number of arguments: 1 = only one allowed, N = multiple allowed, the parameter name is repeated for multiple entries, for example, x=1&x=2&x=3.

Exceptions

Parameter	Exception	Description
copy	startByte	If you specify a copy, you can NOT specify a startByte.
directory	copy	If you specify a directory, you can NOT specify a copy.
newFileName	files	If you specify a newFileName, you can NOT specify more than one (1) files.
newFileName	directory	If you specify a newFileName, you can NOT specify a directory.
startByte	newFileName	If you specify a startByte, you MUST specify a newFileName.

Examples

Restore a file and give it a new file name.

```
http://host:81/axis2/services/stornext/doRetrieve  
?password=stornext  
&files=/stornext/managed/a  
&newFileName=/stornext/managed/b  
&format=json
```

Restore the second (2nd) copy of a file and give it a new file name.

```
http://host:81/axis2/services/stornext/doRetrieve  
?password=stornext  
&files=/stornext/managed/a  
&copy=2  
&newFileName=/stornext/managed/a2  
&format=xml
```

Restore a portion of a file and give it a new file name.

```
http://host:81/axis2/services/stornext/doRetrieve  
?password=stornext  
&files=/stornext/managed/a  
&newFileName=/stornext/managed/a-123-456  
&startByte=123  
&endByte=456  
&format=xml
```

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.

Parameter	Req / Opt	Num	Description
password	Required	1	This is the password stored in the /usr/adic/.snapipassword file on the server. Example: password=stornext
files	Required	N	The names of the file(s) on disk to store to media. The full path name starting from the root directory is required as input to the command. Multiple filenames must be separated by commas. Example: files=/stornext/dir/a
mediaType	Optional	1	Defines the type of medium to be used for storage. Depending on the type of platform used, the following media types are supported by Tertiary Manager software: AIT, AITW, LTO, LTOW, SDISK, 3590, 3592, 9840, 9940, T10K, DLT4, DLT2 (SDLT600 media) Example: mediaType=LTOW

Parameter	Req / Opt	Num	Description
copies	Optional	1	<p>Number of copies of the file(s) to be stored. The value is the total number of copies, including the primary copy of the file. This number cannot exceed the number of copies defined in the policy class <code>maxcopies</code> parameter. To store more copies of the file than are specified in the default <code>copies</code> parameter in the policy class, the file's <code>copy</code> attribute must be modified. Use <code>fschfiat filename -c</code> to change this attribute. Then the <code>-c</code> option can be used with the <code>doStore</code> API to store additional copies of the file. If the number of copies stored is less than the number specified by the policy class definition or by the fschfiat command, the remaining copies are stored when the storage policy is applied.</p> <p>Example: <code>copies=2</code></p>
retention	Optional	1	<p>The file retention policy for the filename specified. The files can be truncated immediately (i) or at policy application time (p) once all file copies are stored on a medium. If the <code>retention</code> option is not used, the file retention policy will be specified by the policy class definition.</p> <p>Example: <code>retention=i</code></p>

Parameter	Req / Opt	Num	Description
drivePool	Optional	1	Media Manager drive pool group used to store the file specified. The drive pool must be defined in Media Manager software. If the drivePool option is not used, the default drive pool group will be specified by the policy class definition. The special "_" character is permitted to identify the drive pool group. Example: drivePool=pool3
minSize	Optional	1	Minimum file size in bytes to be stored. Files larger than or equal to the specified minSize will be stored. Files with a size less than specified minSize will not be stored. Example: minSize=123456
runTime	Optional	1	Maximum allowable time in hours for the command to finish. This command normally runs until it completes. This option can be used to limit how long it should remain active. If the store has not completed in the specified amount of time, then any outstanding activity will be cancelled. Example: runTime=2
format	Optional	1	Output format: json, text, or xml. The default output format is text. Example: format=xml

Num is the number of arguments: 1 = only one allowed, N = multiple allowed, the parameter name is repeated for multiple entries, for example, `x=1&x=2&x=3`.

Examples

```
http://host:81/axis2/services/stornext/doStore
?password=stornext
&files=/stornext/managed/a
&format=json
```

The doTruncate API

The doTruncate API allows a user to remove the copy/copies of the specified file(s) from disk after the file is copied to media. The cleanup policy is the system administration method of routinely freeing disk space by removing files after being stored on media.

The doTruncate API is used by users to maintain a desired level of disk space by truncating individual files. Files specified for removal from disk with the doTruncate API command must have an exact copy on media.

Suggested uses for the **fsrmdiskcopy** command include the following:

- After viewing a migrated file. Viewing the file caused it to be retrieved to disk. If the file is not modified, the disk copy can be removed.
- After storing the file to medium with the **fsstore** command without using the option to immediately truncate the file from disk.
- Between the application of the storage and truncation policies by the system administrator.

Running this API is equivalent to running the **fsrmdiskcopy** command.

Parameter	Req / Opt	Num	Description
password	Required	1	This is the password stored in the /usr/adic/.snapipassword file on the server. Example: password=stornext
files	Required	N	The names of the file(s) on disk to store to media. The file path(s) must also be in a migration directory. The full path name starting from the root directory is required as input to the command. Multiple filenames must be separated by commas. Example: files=/stornext/dir/a
format	Optional	1	Output format: json, text, or xml. The default output format is text. Example: format=xml

Num is the number of arguments: 1 = only one allowed, N = multiple allowed, the parameter name is repeated for multiple entries, for example, x=1&x=2&x=3.

Example

```
http://host:81/axis2/services/stornext/doTruncate
?password=stornext
&files=/stornext/managed/a
&format=json
```

The getDriveReport API

The getDriveReport API is a user command that can be executed when Tertiary Manager software is active or nonactive. The getDriveReport API reports the state of the Tertiary Manager software and/or all storage subsystems and drive components configured in the Quantum storage subsystem.

Submitting the getDriveReport API with the componentAlias option generates a report for a single Quantum component, i.e. drive(s), drive identifier(s), and Media Manager system(s).

Running this API is equivalent to running the fsstate command.

Parameter	Req / Opt	Num	Description
password	Required	1	This is the password stored in the /usr/adic/.snapipassword file on the server. Example: password=stornext
componentAlias	Optional	1	The alias for storage subsystem and drives. The system administrator configures the possible values for component aliases during system configuration or by using the fsconfig command. Example: componentAlias=sdisk

Parameter	Req / Opt	Num	Description
format	Optional	1	Output format: json, text, or xml. The default output format is text. Example: format=xml

Num is the number of arguments: 1 = only one allowed, N = multiple allowed, the parameter name is repeated for multiple entries, for example, x=1&x=2&x=3.

Examples

Status of the 'sdisk' component.

```
http://host:81/axis2/services/stornext/getDriveReport
?password=stornext
&componentAlias=sdisk
&format=xml
```

The getFileLocation API

The getFileLocation API reports the current location(s) of files, whether on disk, on a particular medium, or not in the SNMS system. It also shows file attribute information.

Running this API is equivalent to running the **fsfileinfo** command.

Parameter	Req / Opt	Num	Description
password	Required	1	This is the password stored in the /usr/adic/.snapipassword file on the server. Example: password=stornext

Parameter	Req / Opt	Num	Description
files	Required	N	The path name of at least one file is required. The full path name starting from the root directory is required. Example: files=/stornext/dir/a
checksum	Optional	1	If checksum was turned on for the file when stored and this option is specified, the checksum value generated for the file will be displayed. Example: checksum=1
showIds	Optional	1	Show whether the file has any objects stored to the Wide Area Storage. Example: showIds=1
format	Optional	1	Output format: json, text, or xml. The default output format is text. Example: format=xml

Num is the number of arguments: 1 = only one allowed, N = multiple allowed, the parameter name is repeated for multiple entries, for example, x=1&x=2&x=3.

Examples

Basic file information.

```
http://host:81/axis2/services/stornext/getFileLocation
?password=stornext
&files=/stornext/managed/a
&format=xml
```

Information about two files, with checksum and object IDs.

```
http://host:81/axis2/services/stornext/getFileLocation
?password=stornext
&files=/stornext/managed/a
&files=/stornext/managed/b
&checksum=1
&showIds=1
&format=xml
```

The getFileTapeLocation API

The getFileTapeLocation API reports the location information of the file's on-tape copies. The report will list all the segments belonging to the specified file copy.

Running this API is equivalent to running the **fsfiletapeloc** command.

Parameter	Req / Opt	Num	Description
password	Required	1	This is the password stored in the /usr/adic/.snapipassword file on the server. Example: password=stornext
file	Required	N	The full path name starting from the root directory is required. Example: file=/stornext/dir/a
copy	Optional	1	The copy id to generate report for. If not specified, the information for the primary copy will be reported. Example: copy=2

Parameter	Req / Opt	Num	Description
format	Optional	1	Output format: json, text, or xml. The default output format is text. Example: format=xml

Num is the number of arguments: 1 = only one allowed, N = multiple allowed, the parameter name is repeated for multiple entries, for example, x=1&x=2&x=3.

Examples

Basic file information.

```
http://host:81/axis2/services/stornext/
getFileTapeLocation
?password=stornext
&file=/stornext/managed/a
&format=xml
```

The getMediaInfo API

The getMediaInfo API produces a list of media. The organization of the media list is defined by the use of options. If no options are used, the getMediaInfo API generates a short report that lists the total quantity of media in each policy class, including the general scratch pool.

Running this API is equivalent to running the **fsmedlist** command.

Parameter	Req / Opt	Num	Description
password	Required	1	This is the password stored in the /usr/adic/.snapipassword file on the server. Example: password=stornext

Parameter	Req / Opt	Num	Description
scratchPoolOnly	Required	1	Used to report on the blank media in the general scratch pool. Example: scratchPoolOnly=1
format	Optional	1	Output format: json, text, or xml. The default output format is text. Example: format=xml

Num is the number of arguments: 1 = only one allowed, N = multiple allowed, the parameter name is repeated for multiple entries, for example, x=1&x=2&x=3.

Examples

Media report, no scratch pool.

```
http://host:81/axis2/services/stornext/getMediaInfo
?password=stornext
&scratchPoolOnly=0
&format=xml
```

Media report, general scratch pool.

```
http://host:81/axis2/services/stornext/getMediaInfo
?password=stornext
&scratchPoolOnly=1
&format=xml
```

The getMediaReport API

The getMediaReport API produces either a short report or a long report on the specified media, based on the options entered. One or more media identifiers must be entered.

Running this API is equivalent to running the **fsmedinfo** command.

Parameter	Req / Opt	Num	Description
password	Required	1	This is the password stored in the /usr/adic/.snapipassword file on the server. Example: password=stornext
mediaIDs	Required	N	One or more media identifiers on which to report. Example: mediaIDs=sdisk
longReport	Optional	1	Produce the long form of the report that contains the same information as the short form, plus a list of the file segments on the medium. The pathname that is shown is the name of the file at the time the file was stored. If the file has been renamed since that time that will not be reflected in this report. If the parent or name of an individual file cannot be found, the getMediaReport API will indicate that fact but still report the key. Example: longReport=1
format	Optional	1	Output format: json, text, or xml. The default output format is text. Example: format=xml

Num is the number of arguments: 1 = only one allowed, N = multiple allowed, the parameter name is repeated for multiple entries, for example, x=1&x=2&x=3.

Examples

Media report, short form.

```
http://host:81/axis2/services/stornext/getMediaReport
?password=stornext
&mediaIDs=sdisk
&format=xml
```

Media report, long form.

```
http://host:81/axis2/services/stornext/getMediaReport
?password=stornext
&mediaIDs=sdisk
&longReport=1
&format=xml
```

The getSMQueue API

The getSMQueue API checks the request queue for the specified request identifier(s), filename(s), or media. Requests awaiting resources (drives and media) are displayed.

Issuing the getSMQueue API without any options will report all resource requests associated with storage subsystems, for example, drive-media mount requests.

Running this API is equivalent to running the **fsqueue** command.

Parameter	Req / Opt	Num	Description
password	Required	1	This is the password stored in the /usr/adic/.snapipassword file on the server. Example: password=stornext

Parameter	Req / Opt	Num	Description
report	Optional	1	The report type. Valid names are: file, media, moverhost, or moverrequest. Example: report=file
requestID	Optional	1	The request identifier of the request to be reported. Valid when the report type is file or media. Example: requestId=1234
file	Optional	N	The file option reports the current status and request identifier associated with a specified file or files. This method gets the active request identifier associated with a store or retrieve if a filename is known. However, because of potential dependencies, the original request identifier for The doStore API or The doRetrieve API issued by a particular user may not be discernible. Valid when the report type is file. Example: file=/stornext/dir/a
format	Optional	1	Output format: json, text, or xml. The default output format is text. Example: format=xml

Num is the number of arguments: 1 = only one allowed, N = multiple allowed, the parameter name is repeated for multiple entries, for example, x=1&x=2&x=3.

Examples

Check request queue for a few file.

```
http://host:81/axis2/services/stornext/getSMQueue  
?password=stornext  
&report=file  
&file=/stornext/managed/a  
&format=xml
```

Check request queue for a specific request ID.

```
http://host:81/axis2/services/stornext/getSMQueue  
?password=stornext  
&report=media  
&requestId=1234  
&format=xml
```

The getWasConfiguration API

The `getWasConfiguration` API reports configuration settings for Wide Area Storage components in the storage system. Wide Area Storage components are: **Appliances**, **Controllers**, **I/O Paths** and **Namespaces**. These components and their attributes, provide the addressing information required to form the URL to store and retrieve objects from the Wide Area Storage.

Running this API is equivalent to running the `fswascfg` command.

Parameter	Req / Opt	Num	Description
password	Required	1	This is the password stored in the <code>/usr/adic/.snapipassword</code> file on the server. Example: <code>password=stornext</code>

Parameter	Req / Opt	Num	Description
format	Optional	1	Output format: json, text, or xml. The default output format is text. Example: format=xml

Num is the number of arguments: 1 = only one allowed, N = multiple allowed, the parameter name is repeated for multiple entries, for example, x=1&x=2&x=3.

Example

```
http://host:81/axis2/services/stornext/
getWasConfiguration
?password=stornext
&format=xml
```

The relocateFiles API

Use the relocateFiles API to relocate a managed file from one disk affinity to another, or change the affinity association of a truncated file.

Running this API is equivalent to running the fsrelocate command.

Parameter	Req / Opt	Num	Description
password	Required	1	This is the password stored in the /usr/adic/.snapipassword file on the server. Example: password=stornext

Parameter	Req / Opt	Num	Description
affinity	Required	1	The destination affinity. This affinity must be defined for the file system in which the file resides. The file will be relocated to this affinity if it has not been truncated. If the file has been truncated, only the file's affinity association will change. Example: affinity=tier1
file	Required	N	One or more files to be relocated. The files must reside in a managed directory. Example: file=/stornext/dir/a
format	Optional	1	Output format: json, text, or xml. The default output format is text. Example: format=xml

Num is the number of arguments: 1 = only one allowed, N = multiple allowed, the parameter name is repeated for multiple entries, for example, x=1&x=2&x=3.

Examples

Relocate a single file.

```
http://host:81/axis2/services/stornext/relocateFiles
?password=stornext
&affinity=tier1
&file=/stornext/managed/a
&format=xml
```

Relocate multiple files.

```
http://host:81/axis2/services/stornext/relocateFiles
?password=stornext
&affinity=tier1
&file=/stornext/managed/a
&file=/stornext/managed/b
&format=xml
```

Examples

This section contains examples for Web Services URLs, sample XML output, sample JSON output, and sample text output.

Example: Web Services URLs

Following are example URLs which demonstrate how to use the output format parameter. The URLs are formatted as multiple lines to make them easier to read.

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>  
&updateATime=1  
&directory=/stornext/dir/  
&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>
</xml>
```



```

<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 MINTRUNCTIME, 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 5 minutes, so that neither the LoSpace nor the daily truncation will truncate any files for 5 minutes.

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 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 `mintruntime` 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 truncation 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 617.

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 619 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 effect 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 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 of security as it pertains to StorNext.

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. For additional information, see the heading [StorNext Security](#) on page 479.

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.

Note: For Xsan / Mac OS X clients to use ACLs when connecting to a StorNext MDC, the file system must be configured with the **enforceAcls** file system configuration option set to **Yes**.

On StorNext version 4.7.x Windows MDCs, this configuration must be added manually using a text editor such as Notepad. If configurations are saved using the **Windows StorNext File System Configuration** tool, the **enforceAcls** file system configuration option is removed and should be added manually.

The configuration file is located at `\Program Files\StorNext\config`.

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

Directories	Files
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

Windows	Mac OS X
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 "+ai" 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_ctl.xml` file (this file resides under `/usr/cvfs/config` on an nss coordinator server). 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/
ctl/v1.0">
  <securityControl fileSystem="snfs1">
    <controlEntry>
      <client type="host">
        <hostName value="192.168.230.132"/>
      </client>
      <controls>
        <mountReadOnly value="true"/>
        <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>
    </controlEntry>
  </securityControl>
</snfsControl>
```

```

        </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: Quotas can only be enabled or disabled by modifying the Quotas parameter of the file system config file. The CLI `cvadmin "quotas"` command will still return the quota state, but it cannot change it.

`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.

This appendix contains the following troubleshooting topics:

- [Troubleshooting StorNext File System](#)
- [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: 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 fsnameservers file on all clients and verify the fsnameservers file match what the MDCs display.

The fsnameservers file is located in the following directory, depending upon the product and operating system:

* For Windows StorNext File System: `C:\Program Files\StorNext\config`

* For Linux or UNIX: `/usr/cvfs/config`

- 3 Correct the fsnameservers file to resemble the following IP addresses:
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 UNIX-based and Windows-based clients, 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:\Program Files\StorNext\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 `/usr/cvfs/config` for UNIX-based clients and `C:\Program Files\StorNext\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 StorNext client.

On the StorNext client, only the client portion of the StorNext product suite is installed. Verify this by running the command `/usr/cvfs/bin/cvversions`. The following output appears:

```
qlha2:~ # cvversions
```

```
Server not installed.
```

```
File System Client:
```

```
Client Revision 4.2.0 Build 21233 Branch branches_4.2.0
```

```
Built for Linux 2.6.16.60-0.21-smp x86_64
```

```
Created on Thu Aug 4 04:11:01 MDT 2011
```

```
Built in /home/mlund/nightly/VM-0-SuSE100ES-26x86-64-
SP2/sn/buildinfo
```

```
Host OS Version:
```

```
Linux 2.6.16.60-0.85.1-smp #1 SMP Thu Mar 17 11:45:06
UTC 2011 x86_64
```

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 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.

Question: Why does data pulling not occur when changing the replication target from the replication policy?

Answer: The reason the data pulling does not occur after the change of target, is that the source directory had been replicated before. The files were in sync with the content in the old target. Thus each file had its flag replication in-sync (not stale). When replication occurs after the change of target, these files are not flagged as stale, so no data pulling occurs on the target, since only the files flagged as stale are pulled.

If an snpolicyd managed source directory has been replicated to one or more target hosts' managed directories, and a replication policy is subsequently changed to cause replication to new directories on the target hosts, an snpolicy **replicateforce** command must be run following the policy change(s) to ensure proper replication and Storage Manager processing of all files on the targets.

Question: Why are BLOBs not replicated if replication is enabled after files are deduplicated, and truncated? In other words, enabling

replication will only copy the metadata information, not the data or BLOBs.

Answer: Since the file contents are up-to-date for ingested files, files are not flagged as required to be pulled when replication occurs. As a result, no data, and blob tags are pulled on the target.

If an snpolicyd managed directory has a deduplication-enabled (no replication) policy associated, then when replication is later turned on for this policy, an snpolicy **replicateforce** command must be run immediately following the policy change in order to replicate content of the files under the directory to the new target.

Question: Why do the replication quiesce scripts not synchronize data on any clients that have open files?

Answer: To avoid this issue, close open files prior to running quiesce scripts.

Question: How can I delete a TSM relation point used for replication?

Answer: You can manually delete the relation point by running the command `rm -rf /snfs/sn2/tsm/.rep_private`, which empties the TSM relation point. When running this command, be aware that there may have been several targets being realized with the TSM relation point in question, so you should remove the directory `tsm_dir / .rep_private` only after the LAST target policy has been removed from the relation point.

Question: Why do I receive blockpool errors if a deduplication candidate is removed before blockpool processing is completed?

Errors such as the following may be sent to the syslog:

```
Oct  2 15:22:00 orleans Blockpool[16403]: E: [5]
(Store Local) Error storing file "/stornext/source/
__CVFS_Handle.000474F892EBB65E000E0000000000000000000
00292BF4". Error opening file "/stornext/source/
__CVFS_Handle.000474F892EBB65E000E000000000000000000
00292BF4". No such file or directory.
```

Answer: Errors such as these may appear serious, but there is no reason for concern. If you receive these errors, no action is required.

Question: Why do the default settings for snpolicyd cause memory starvation problems for small deduplication-enabled configurations (1TB deduplication capacity) when ingesting to or retrieving from the blockpool?

Answer: Quantum recommends changing the values of the parameters **ingest_threads** and **event_threads** to **4** (from their default values of **8**) in the StorNext Replication/Deduplication configuration file (/usr/cvfs/config/snpolicyd.conf).

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 559 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.

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 'fsmist'` command, which displays the local FSMPM's internal FSM and process table. A remote FSMPM's table can also be viewed with `'fsmist 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 FSMPM on an MDC with active HA monitored FSMs sends a request to its peer FSMPM 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 FSMPM. 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.

Question: The Secondary HA MDC system is in "locked" and "stopped" mode, as seen from the Primary HA MDC. How can the Secondary HA MDC be restored to the "default" mode of operation?

```
PrimaryMDC# /usr/adic/DSM/bin/snhamgr status
LocalMode=config LocalStatus=primary RemoteMode=locked
RemoteStatus=stopped
```

Answer: Follow these steps to restore a secondary HA MDC to the default mode of operation. If the Primary HA MDC has status `LocalMode=default LocalStatus=primary`, proceed to step 6.

- 1 Verify the file systems availability before exiting the Config mode.

```
PrimaryMDC# /usr/adic/DSM/bin/cvadmin
```

Verify that all the files systems listed in the `fsmlist` file are listed and have "*" displayed to signify the Primary MDC has activated its FSMs.

If not, within **cvadmin** run:

```
PrimaryMDC<snadmin> disks refresh
```

```
PrimaryMDC<snadmin> select
```

If any of the file systems are not listed or do not display as activated ("*"), resolve this before making any other changes to the HA modes.

- 2 Verify that the HaShared file system is mounted.

```
PrimaryMDC# /bin/mount | grep HAM
```

For example: `/dev/cvfsctl1_HAFS on /usr/adic/HAM/shared type cvfs (rw,sparse=yes)`

- 3 Verify that you can write and read to the HaShared filesystem.

```
PrimaryMDC# date > /usr/adic/HAM/shared/test_1.tmp
```

```
PrimaryMDC# cat /usr/adic/HAM/shared/test_1.tmp
```

You should see: the current date

```
PrimaryMDC# rm /usr/adic/HAM/shared/test_1.tmp
```

When you are able to successfully write and read to the file system, continue the steps below.

- 4 Set the HA mode back to default mode of operation.

Note: This step will restart StorNext, and may prevent clients from working for an extended period of time.

```
PrimaryMDC# cd /usr/adic/DSM/bin/
```

```
PrimaryMDC# ./snhamgr mode=default
```

```
LocalMode=default:LocalStatus=stopped:RemoteMode=locked:RemoteStatus=stopped
```

```
PrimaryMDC# /sbin/service cvfs start
```

(Numerous amounts of output)

- 5 Repeat **Steps 1, 2** and **3** to verify file system functionality. If this passes then continue to next step.

- 6 From the secondary MDC, change the state of the backup server **snhamgr** to **default** and start the StorNext software:

```
SecondaryMDC# snhamgr mode=default
```

```
LocalMode=default:LocalStatus=stopped:RemoteMode=default:RemoteStatus=primary
```

```
SecondaryMDC# /sbin/service cvfs start
```

(Numerous amounts of output)

- 7 Verify the status of the HA.

```
SecondaryMDC# cd /usr/adic/DSM/bin
```

```
SecondaryMDC# ./snhamgr status
```

```
LocalMode=default:LocalStatus=running:RemoteMode=default:RemoteStatus=primary
```

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 I/O 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: What should I do after a database (MySQL) crash?

Answer:

After a crash and restart, the next time Storage Manager initiates, the MySQL database will check internal logs to determine whether there were any unfinished transactions, and normally will automatically make corrections. If it cannot make necessary corrections, MySQL will not complete startup. In this case contact Quantum Technical Support.

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 disk.

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.

Question: How do I restart the StorNext GUI if it is inaccessible in a Web browser, with one of the following error messages displayed?

Firefox: Unable to connect. Firefox can't establish a connection to the server


Internet Explorer: Internet Explorer cannot display the webpage

Answer: If you encounter this condition, restart the StorNext GUI on the MDC server by doing the following:

1 Open a root UNIX shell window on the MDC.

2 Run the command `service stornext_web restart`

The "service" command will return before the service is ready to be accessed by a browser. Wait a few moments before trying to connect, and then retry if that fails.



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 StorNext Offline Notification feature can be installed on Windows SAN Clients. Refer to the *StorNext Compatibility Guide* for more information on Windows SAN clients.

Caution: The StorNext Offline Notification feature is intended for single user systems. Do not install it on systems where multiple users may be logged on at the same time. 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 174](#).

If you select this option, you must log in with the credentials for the Administrator account as shown in [Figure 175](#).

Appendix H: StorNext Offline Notification Installing the Notification Application

Figure 174 Run as Administrator

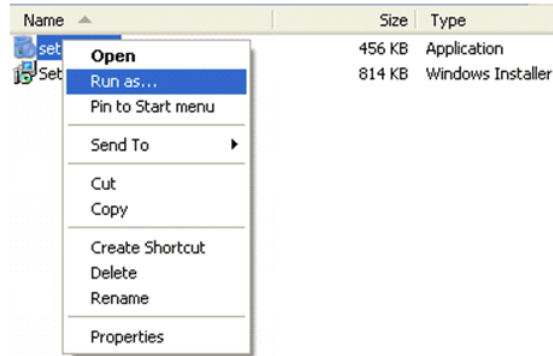



Figure 175 Logging in to the Administrator Account



- 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 4.0 as shown in [Figure 176](#).

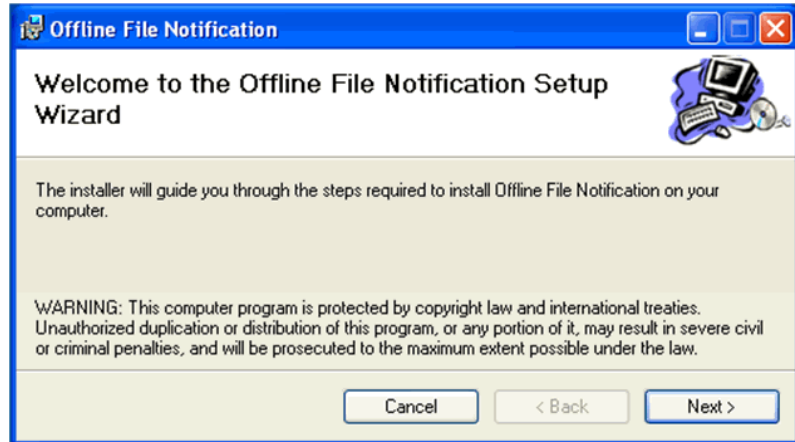
Figure 176 Installing the .NET Framework



- 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 177 Offline Notification Setup Wizard




- 5 Click **Next** to continue. The Quantum End User License Agreement window appears.

Figure 178 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.)

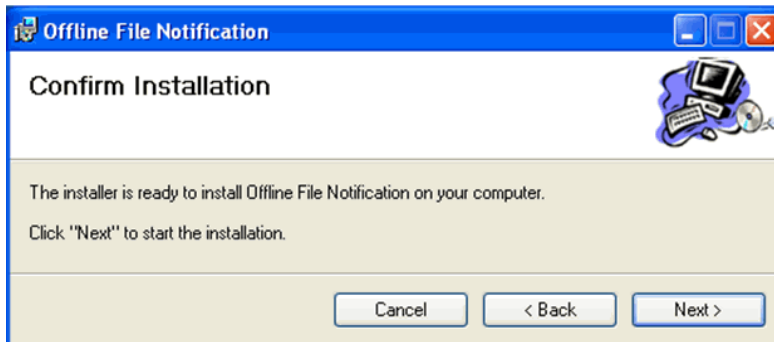
Figure 179 Select Installation Folder



- 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 180 Confirm Installation

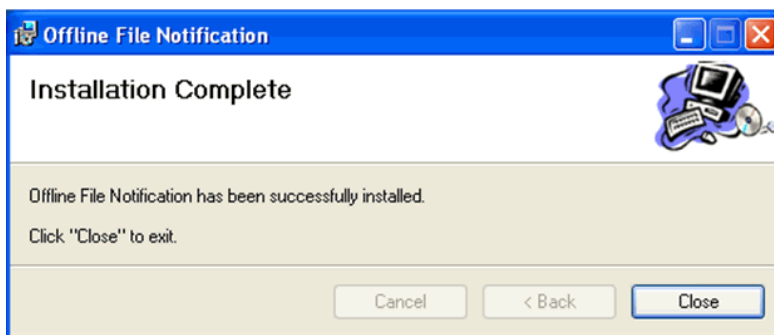


- 9 Before you proceed, make sure you have enough disk space to install the binaries. A minimum of 10MB is required.
- 10 Click **Next** to begin the installation.

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 181 Installation Complete



- 11 Click **Close** to exit the installation .

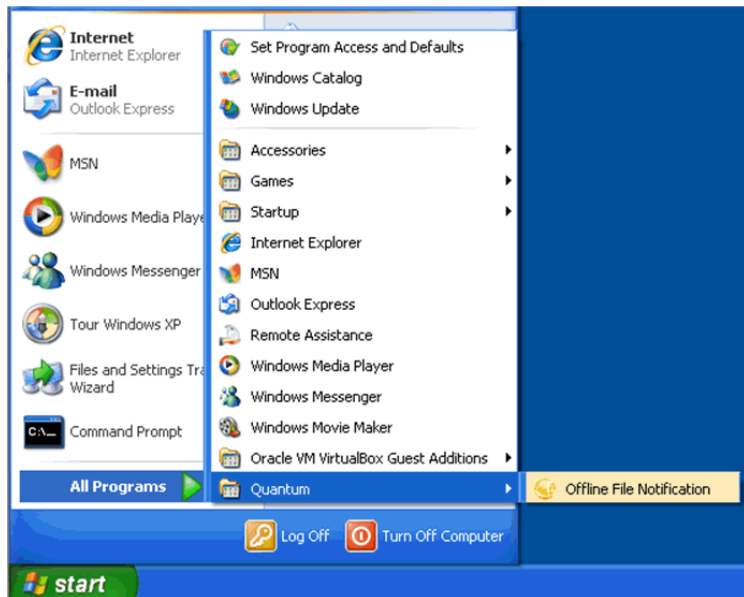
The next time you start your computer, the notification application automatically starts on login.

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 182 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:

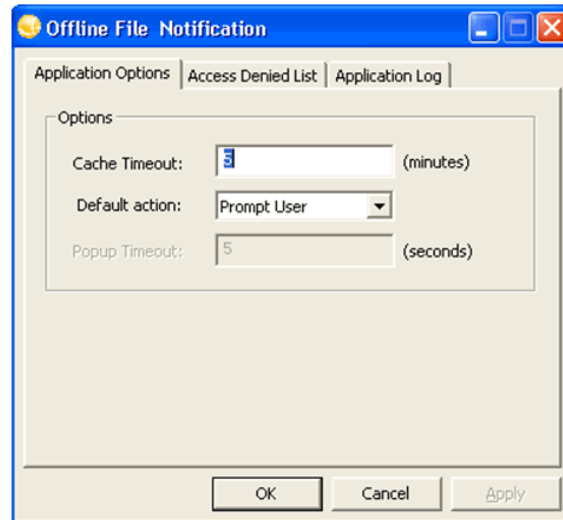
- **Preferences:** View and enter application options, view a blocked file list, and view the application log
- **About:** View the StorNext Offline Notification application's current version number and copyright date
- **Exit:** Terminate the application

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 183 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 184](#).) 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 184 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 185 Application log



Exiting Application Preferences

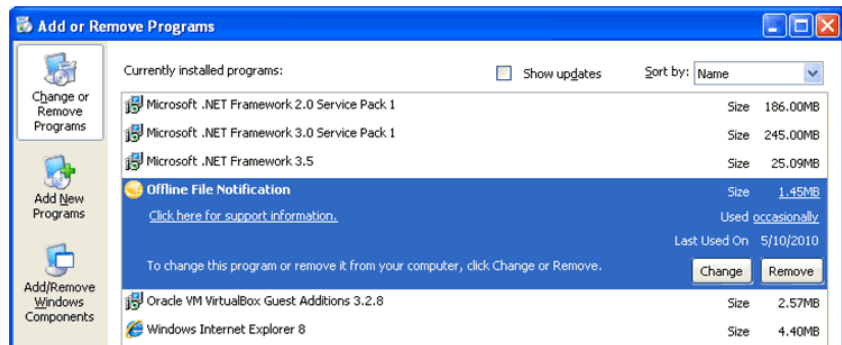
When you are finished viewing or setting preferences for the StorNext Offline Notification application, click **OK** to close the Preferences window.

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 186 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

As of StorNext 4.7.3, the RAS messages documentation is being maintained in the *StorNext RAS Events and FRU Reference Guide*.

For the latest RAS messages documentation, see the *StorNext RAS Events and FRU Reference Guide* available on-line at <http://www.quantum.com/sn5docs>.



Appendix J

Repairing and Replacing StorNext Metadata Servers

This appendix contains information on repairing and replacing Metadata Servers. For purposes of this chapter, we will use the Metadata Controller (MDC) when discussing the server. The term MDC includes both customer-configured metadata servers as well as StorNext M660, M440 and M330 Metadata Appliances.

MDCs may need to be replaced for various reasons. An MDC might be experiencing hardware failures, need to be moved to better performing hardware, or the operating system may need to be upgraded or reinstalled. This document provides procedures for replacing an MDC. It assumes that the existing MDC operational.

Note: When replacing Network Interface Cards, if you replace the MDC host network interface card which contains the MAC address matching the `cvfsid` value noted in the host's `license.dat` file, then a new `license.dat` file must be generated. This must also be taken into consideration when replacing the entire MDC host. Please contact Quantum Support (see [Chapter 12, Customer Assistance](#)) for an updated `license.dat` file.

Replacing an MDC in a non-HA environment (non-backup/restore method)

This procedure covers the steps necessary to move a non-HA StorNext MDC from a fully functioning StorNext environment to a potentially different system. This covers the cases of moving an MDC from one system to another and the case where a user wishes to reinstall their operation system (e.g. upgrading from RHEL 5 to RHEL 6).

The procedure in this case is to create a full StorNext backup (**snbackup** file) and restore only the configuration files on the new system. The database and metadump files will be packaged up by the user by way of tar files and unpackaged on the new system. This will allow users to skip creating new metadump files after extracting the database and metadump files in the new environment.

This method will demand that all managed file systems remain offline and unmounted until after all of the database and metadump files have been extracted onto the destination system.

- 1 While not required, it is recommended that all managed file systems be unmounted on all the clients and quiesced to help eliminate I/O errors on clients due the StorNext MDC being down.
- 2 Create a full backup using the **snbackup** command. Take note of the backup ID. This will be needed in the next step:

```
# snbackup
```

The **snbkpreport** can be used to determine the latest backup ID if necessary:

```
# snbkpreport
```

- 3 Copy backup files and manifests off the system onto the network or external file system. For the purpose of this procedure assume that **\$DESTDIR** is the location of an nfs share at **/net/share/migration**. Also assume the **\$ID** is the backup ID from the full backup in [Step 2](#). Use **showsysparm** to identify the mount point for the StorNext backup. There is a **meta.\$FSNAME.\$ID.tgz** file for each managed file system in the system where **\$FSNAME** is the file system name of the managed file system.

Note: The **db**, **meta** and **manifest** files are copied to **\$DESTDIR** in the event a full recovery if necessary later.

- ```
showsysparm BACKUPFS
mkdir $DESTDIR/snbackup
cp $BACKUPFS/.ADIC_INTERNAL_BACKUP/conf.$ID.*.tgz
$DESTDIR/snbackup
cp $BACKUPFS/.ADIC_INTERNAL_BACKUP/db.$ID.*.tgz
$DESTDIR/snbackup
cp $BACKUPFS/.ADIC_INTERNAL_BACKUP/meta.*.$ID.*.tgz
$DESTDIR/snbackup
cp /usr/adic/TSM/internal/status_dir/
snbackup_manifest $DESTDIR
cp /usr/adic/TSM/internal/status_dir/
device_manifest $DESTDIR
```
- 4 Stop StorNext. It is critical that all managed file systems be cleanly shutdown to avoid having to create a new metadump later:

```
service stornext_web stop
service cvfs stop
```
  - 5 Create a tar archive of database db and journal directories:

```
tar -zcvhf $DESTDIR/db_source.tar.gz /usr/adic/
mysql/db \
/usr/adic/mysql/journal
```
  - 6 Create a tar archive of database metadumps directory:

```
tar -zcvhf $DESTDIR/meta_source.tar.gz /usr/adic/
database/metadumps
```
  - 7 Copy `/etc/fstab` to the external file system (for later reference):

```
cp /etc/fstab $DESTDIR
```
  - 8 Remove StorNext with the `-remove` option to preserve log files:

```
install.stornext -remove
```
  - 9 Create tar archive of preserved files for later reference:

```
tar -zcvhf $DESTDIR/preserved_logs.tar.gz /usr/adic
```

- 10 (Optional) At this point, StorNext has been removed from the original MDC with the necessary state copied off to a safe location and the operating system can now safely be reinstalled if so desired.
- 11 Install StorNext on new destination using the identical version that the original MDC was running.

```
install.stornext
```

- 12 Restore the configuration files from the backup in [Step 2](#). You will be prompted for the backup ID. This can also be found in the **snbackup\_manifest** or derived from the **db.\$ID.0.tgz** file name where **\$ID** is the backup ID.

---

**Note:** The **snrestore** command will generate error messages that it is unable to read the manifest files. This is expected and can be safely ignored.

---

Be sure to use the **-c** modifier as shown here:

```
snrestore -c -r $DESTDIR/snbackup
```

- 13 Make sure that StorNext has stopped:

```
service cvfs stop
```

- 14 Restore database from tar backup created in [Step 5](#):

```
tar -zxvf $DESTDIR/db_source.tar.gz -C /
```

- 15 Restore metadump from tar archive created in [Step 6](#):

```
tar -zxvf $DESTDIR/meta_source.tar.gz -C /
```

- 16 Modify local config files as necessary. If the IP address for the MDC has changed, the **fsnameservers** file may need to be updated:

```
/usr/cvfs/config/fsnameservers
```

If the motherboard or network card(s) have changes, you may need to acquire a new license file (**/usr/cvfs/config/license.dat**) may be required from Quantum.

Following the **snrestore** in [Step 12](#), the file system configuration files are named **.backup\_\$(FSNAME).cfgx** where **\$(FSNAME)** is the file system name. These files will need to be renamed to the format **\$(FSNAME).cfgx** as shown here:

```
mv /usr/cvfs/config/.backup_$(FSNAME).cfgx /usr/cvfs/
config/$(FSNAME).cfgx
```

17 The mount points and fstab entries will need to be recreated. Refer to the `/etc/fstab` file copied to `$DESTDIR` in [Step 7](#).

18 Start StorNext:

```
service cvfs start
service stornext_web start
```

19 Verify that backups are working correctly by running a full backup.

```
snbackup
```

---

**Note:** If the backup complains about metadump failures due to missing restore journal sequence numbers, then a new **metadump** file will need to be made. To create a new **metadump** file, go to the next section, [Step 12](#).

---

20 Verify StorNext is functioning by storing and retrieving files.

21 StorNext should now be up and running and safe for clients to start accessing. However, it is possible that a new mapping file or event files may need to be generated for each file system. This happens as part of the rebuild policy. This step may be scheduled to run at a later time and is by default run once a week:

```
fspolicy -b -y /path/to/stornext/mount/point
```

## Replacing an MDC in non-HA environment (backup/restore method)

In the event that the metadump is damaged or a restore journal is missing, an MDC can be migrated to another system. If you need further assistance, contact Quantum support.

- 1 While not required, it is recommended that all managed file systems be unmounted on all the clients and quiesced to help eliminate I/O errors on the clients due the StorNext MDC being down.

- 2 Create a full backup using the `snbackup` command. Take note of the backup ID. This will be needed in the next step:

```
snbackup
```

The `snbkpreport` can be used to determine the latest backup ID if necessary:

```
snbkpreport
```

- 3 Copy backup files and manifests off the system onto the network or external file system. For the purpose of this procedure assume that `$DESTDIR` is the location of an nfs share at `/net/share/migration`. Also assume the `$ID` is the backup ID from the full backup in [Step 2](#). Use `showsysparm` to identify the mount point for the StorNext backup. There is a `meta.$FSNAME.$ID.tgz` file for each managed file system in the system where `$FSNAME` is the file system name of the managed file system.

```
showsysparm BACKUPFS
```

```
mkdir $DESTDIR/snbackup
```

```
cp $BACKUPFS/.ADIC_INTERNAL_BACKUP/conf.$ID.*.tgz
$DESTDIR/snbackup
```

```
cp $BACKUPFS/.ADIC_INTERNAL_BACKUP/db.$ID.*.tgz
$DESTDIR/snbackup
```

```
cp $BACKUPFS/.ADIC_INTERNAL_BACKUP/meta.*.$ID.*.tgz
$DESTDIR/snbackup
```

```
cp /usr/adic/TSM/internal/status_dir/
snbackup_manifest $DESTDIR
```

```
cp /usr/adic/TSM/internal/status_dir/
device_manifest $DESTDIR
```

- 4 Copy `/etc/fstab` to the external file system (for later reference):

```
cp /etc/fstab $DESTDIR
```

- 5 Remove StorNext with the `-remove` option to preserve log files:

```
install.stornext -remove
```

- 6 Create tar archive of preserved files for later reference:

```
tar -zcvhf $DESTDIR/preserved_logs.tar.gz /usr/adic
```

7 (Optional) At this point, StorNext has been removed from the original MDC with the necessary state copied off to a safe location and the operating system can now safely be reinstalled if so desired.

8 Install StorNext on new destination using the identical version that the original MDC was running:

```
install.stornext
```

9 Do a full restore of the **snbackup**. You will be prompted for the backup ID. This can be found in the **snbackup\_manifest** or derived from the **db.\$ID.0.tgz** file name where **\$ID** is the backup ID.

---

**Note:** The **snrestore** command will generate error messages that it is unable to read the manifest files. This is expected and can be safely ignored.

---

```
snrestore -r $DESTDIR/snbackup
```

10 Modify local configuration files as necessary.

If the IP address for the MDC has changed, the **fsnameservers** file may need to be updated:

```
/usr/cvfs/config/fsnameservers
```

If the motherboard or network card(s) have changes a new license file may be required from Quantum:

```
/usr/cvfs/config/license.dat
```

11 The **fsmpm** must be running and the **fsms** for the managed file systems must be stopped:

```
service cvfs stop
```

Comment out **fsmlist** and any **cvfs /etc/fstab** entries, and then start **cvfs** with the **stratonly** option:

```
DSM_control stratonly
```

12 Recreate the metadump for each managed file system where **\$FS\_NAME** is the name of the file system. After the new metadump has been created it must be applied and optimized for the **fspostrestore** process to work later.

```
snmetadump -dc $FS_NAME
```

```
snmetadump -a $FS_NAME
```

- 13 The mount points and **fstab** entries may need to be recreated. Refer to the **/etc/fstab** file copied to **\$DESTDIR** in [Step 4](#). Restore the **fsmlist** if it was modified in [Step 11](#).
- 14 Restart StorNext with the following commands:

```
service cvfs stop
service cvfs start
service stornext_web start
```
- 15 Synchronize all the managed file system with the database using the **fspostrestore** command. **\$FS\_MNT\_PT** is the mount point of the managed file system and **<YYYY:MM:DD:hh:mm:ss>** is the time from just before the backup was created. Use the **snbkpreport** command to determine when the backup was created:

```
fspostrestore -s <YYYY:MM:DD:hh:mm:ss> $FS_MNT_PT
```
- 16 Verify that backups are working correctly by running a full backup:

```
snbackup
```
- 17 Verify StorNext is functioning by storing and retrieving files.
- 18 StorNext should now be up and running and safe for clients to start accessing. However, it is possible that a new mapping file or event files may need to be generated for each file system. This happens as part of the rebuild policy. This step may be scheduled to run at a later time and is by default run once a week.

```
fspolicy -b -y /path/to/stornext/mount/point
```

---

## Replacing an MDC in an HA environment

This section describes how to replace the secondary HA server. This procedure may also be used to upgrade the operating system of the MDC (e.g. upgrading from RHEL5 to RHEL6). If you need to replace both the primary and secondary MDCs, then run through this procedure once, fail over so that NEW secondary MDC becomes primary, and then run through this procedure a second time.



Before beginning this procedure make sure you have obtained the proper licenses required for the new HA MDC. The current license should be sufficient if you are just upgrading the OS.

---

**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

---

**Note:** If you need to replace the system that is currently the primary MDC, stop StorNext on the primary. This will cause a failover to the secondary system.

---

- 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

- 8 Set the primary node to “Config” mode and the peer node to “Peerdown” mode by running the following commands:  

```
snhamgr peerdown
snhamgr mode=config
```
- 9 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
```
- 10 Change the `/usr/cvfs/config/ha_peer` file on the primary MDC to the new MDC IP address.
- 11 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.
- 12 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.
- 13 Push those changes to the synchronization mirror directory by running this command: `/usr/adic/util/syncha.sh -primary`
- 14 (Optional) Upgrade the Operating System of the new secondary server at this point.
- 15 Install StorNext 4.7.x build on the NEW secondary server by running this command: `install.stornext`
- 16 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.

---

- 17 In the StorNext GUI, go to the **Tools > High Availability > Convert** screen and convert the secondary MDC to HA.

## Post-Conversion Steps

- 18 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.
- 19 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.
- 20 Start and mount StorNext file systems on the clients, and then verify that all clients have full access
- 21 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.
- 22 If you conducted a failover to the secondary server, fail back to the original primary server.
- 23 Verify that all clients still have full access.
- 24 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.
- 25 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.
- 26 (Optional) If you need to replace both MDCs, fail the primary MDC over to the NEW secondary and then repeat this procedure.

Appendix J: Repairing and Replacing StorNext Metadata Servers  
Replacing an MDC in an HA environment